

THE MATHEMATICS OF POKER

By Chip and JoAnn Holden

Other ConJelCo titles:

- Cooke's Rules of Real Poker* by Roy Cooke and John Bond
Hold'em Excellence by Lou Krieger
How to Play Like a Poker Pro by Roy Cooke and John Bond
How to Think Like a Poker Pro by Roy Cooke and John Bond
Internet Poker by Lou Krieger and Kathleen Watterson
Mastering No-Limit Hold'em by Russ Fox and Scott T. Harker
More Hold'em Excellence by Lou Krieger
Real Poker II: The Play of Hands J 992-1999 (2nd ed.)
by Roy Cooke and John Bond
Serious Poker by Dan Kimberg
Stepping Up by Randy Burgess
The Home Poker Handbook by Roy Cooke and John Bond
Why You Lose at Poker by Russ Fox and Scott T. Harker
Winning Low-Limit Hold'em by Lee Jones
Winning Omaha/8 Poker by Mark Tenner and Lou Krieger
Winning Strategies for No-Limit Hold'em
by Nick Christenson and Russ Fox
Video Poker – Optimum Play by Dan Paymar

Software

- StatKing* for Windows

The Mathematics of Poker

Copyright © 2006 by Bill Chen and Jerrod Ankenman

All rights reserved. This book may not be duplicated in any way or stored in an information retrieval system, without the express written consent of the publisher, except in the form of brief excerpts or quotations for the purpose of review. Making copies of this book, or any portion, for any purpose other than your own, is a violation of United States copyright laws.

Publisher's Cataloging-in-Publication Data

Chen, Bill
Ankenman, Jerrod

The Mathematics of Poker
x,382p. ; 29cm.
ISBN-13: 978-1-886070-25-7
ISBN-10:1-886070-25-3

I Title.

Library of Congress Control Number: 2006924665

First Edition
5 7 9 8 6

Cover design by Cat Zaccardi
Book design by Daniel Pipitone
Digital production by Melissa Neely

ConJeiCo LLC
1460 Bennington Ave
Pittsburgh, PA 15217
[412] 621-6040
<http://www.conjelco.com>

Errata, if any, can be found at <http://www.conjelco.com/mathofpoker/>

Acknowledgments	5
Foreword	6
Introduction	8
<hr/>	
Part I: Basics	
Chapter 1 Decisions Under Risk: Probability and Expectation.....	15
Chapter 2 Predicting the Future: Variance and Sample Outcomes.....	24
Chapter 3 Using All the Information: Estimating Parameters and Bayes' Theorem...	34
<hr/>	
Part II: Exploitive Play	
Chapter 4 Playing the Odds: Pot Odds and Implied Odds.....	47
Chapter 5 Scientific Tarot: Reading Hands and Strategies.....	59
Chapter 6 The Tells are in the Data: Topics in Online Poker.....	70
Chapter 7 Playing Accurately, Part I: Cards Exposed Situations.....	74
Chapter 8 Playing Accurately, Part II: Hand vs. Distribution.....	85
Chapter 9 Adaptive Play: Distribution vs. Distribution.....	94
<hr/>	
Part III: Optimal Play	
Chapter 10 Facing The Nemesis: Game Theory.....	100
Chapter 11 One Side of the Street: Half-Street Games.....	110
Chapter 12 Headsup With High Blinds: The Jam-or-Fold Game.....	121
Chapter 13 Poker Made Simple: The AKQ Game.....	138
Chapter 14 You Don't Have To Guess: No-Limit Bet Sizing.....	146
Chapter 15 Player X Strikes Back: Full-Street Games.....	155
Appendix to Chapter 15 The No-Limit AKQ Game.....	167
Chapter 16 Small Bets, Big Pots: No-Fold [0,1] Games.....	174
Appendix to Chapter 16 Solving the Difference Equations.....	191
Chapter 17 Mixing in Bluffs: Finite Pot [0,1] Games.....	194
Chapter 18 Lessons and Values: The [0,1]Game Redux.....	212
Chapter 19 The Road to Poker: Static Multi-Street Games.....	228
Chapter 20 Drawing Out: Non-Static Multi-Street Games.....	243
Chapter 21 A Case Study: Using Game Theory.....	259
<hr/>	
Part IV: Risk	
Chapter 22 Staying in Action: Risk of Ruin.....	274
Chapter 23 Adding Uncertainty: Risk of Ruin with Uncertain Win Rates.....	288
Chapter 24 Growing Bankrolls: The Kelly Criterion and Rational Game Selection.....	297
Chapter 25 Poker Finance: Portfolio Theory and Backing Agreements.....	303
<hr/>	
Part V: Other topics	
Chapter 26 Doubling Up: Tournaments, Part I.....	312
Chapter 27 Chips Aren't Cash: Tournaments, Part II.....	324
Chapter 28 Poker's Still Poker: Tournaments, Part III.....	338
Chapter 29 Three's a Crowd: Multiplayer Games.....	349
Chapter 30 Putting It All Together: Using Math to Improve Play.....	359
Recommended Reading	365
About the Authors	366
About the Publisher	367

Acknowledgments

A book like this is rarely the work of simply the authors; many people have assisted us along the way, both in understanding poker and in the specific task of writing down many of the ideas that we have developed over the last few years. A book called *The Mathematics of Poker* was conceived by Bill, Chuck Weinstock, and Andrew Latto several years ago, before Jerrod and Bill had even met. That book was set to be a fairly formal, textbook-like approach to discussing the mathematical aspects of the game. A few years later, Jerrod and Bill had begun to collaborate on solving some poker games and the idea resurfaced as a book that was less like a mathematics paper and more accessible to readers with a modest mathematical background.

Our deepest thanks to those who read the manuscript and provided valuable feedback. Most notable among those were Andrew Bloch, Andrew Prock, and Andrew Latto, who scoured sections in detail, providing criticisms and suggestions for improvement. Andrew Prock's PokerStove tool (<http://www.pokerstove.com>) was quite valuable in performing many of the equity calculations. Others who read the manuscript and provided useful feedback were Paul R. Pudaite and Michael Maurer.

Jeff Yass at Susquehanna International Group (<http://www.sig.com>) has been such a generous employer in allowing Bill to work on poker, play the World Series, and so on. We thank Doug Costa, Dan Loeb, Jay Siplestien, and Alexei Dvoretzkii at SIG for their helpful comments.

We have learned a great deal from a myriad of conversations with various people in the poker community. Our friend and partner Matt Hawrilenko, and former WSOP champion Chris Ferguson are two individuals especially worth singling out for their insight and knowledge. Both of us participate enthusiastically in the *ARG community, and discussions with members of that community, too, have been enlightening, particularly round table discussions with players including Sabyl Cohen, Jack Mahalingam, JP Massar, Path Beadles, Steve Landrum, and (1999 Tournament of Champions winner) Spencer Sun.

Sarah Jennings, our editor, improved the book significantly in a short period of time by being willing to slog through equations and complain about the altogether too frequent skipped steps. Chuck Weinstock at Conjelco has been amazingly patient with all the delays and extended deadlines that come from working with authors for whom writing is a secondary occupation.

We also appreciate the comments and encouragement from Bill's father Dr. An-Ban Chen who has written his own book and has numerous publications in his field of solid state physics. Jerrod's mother Judy has been a constant source of support for all endeavors, even crazy-sounding ones like playing poker for a living.

Throughout the writing of this book, as in the rest of life, Michelle Lancaster has constantly been wonderful and supportive of Jerrod. This book could not have been completed without her. Patricia Walters has also provided Bill with support and ample encouragement, and it is of no coincidence that the book was completed during the three years she has known Bill.

Foreword

Don't believe a word I say.

It's not that I'm lying when I tell you that this is an important book. I don't even lie at the poker table -- not much, anyway - so why would I lie about a book I didn't even write?

It's just that you can't trust me to be objective. I liked this book before I'd even seen a single page. I liked it when it was just a series of conversations between Bill, myself, and a handful of other math geeks. And if I hadn't made up my mind before I'd read it, I'm pretty sure they'd have won me over with the first sentence.

Don't worry, though. You don't have to trust me. Math doesn't lie. And results don't lie, either. In the 2006 WSOP, the authors finished in the money seven times, including Jerrod's second place finish in Limit Holdem, and Bill's two wins in Limit and Short Handed No Limit Hold'em.

Most poker books get people talking. The best books make some people say, "How could anyone publish our carefully guarded secrets?" Other times, you see stuff that looks fishy enough to make you wonder if the author wasn't deliberately giving out bad advice. I think this book will get people talking, too, but it won't be the usual sort of speculation. No one is going to argue that Bill and Jerrod don't know their math.

The argument will be about whether or not the math is important.

People like to talk about poker as "any man's game." Accountants and lawyers, students and housewives can all compete at the same level - all you need is a buy-in, some basic math and good intuition and you, too, can get to the final table of the World Series of Poker. That notion is especially appealing to lazy people who don't want to have to spend years working at something to achieve success. It's true in the most literal sense that anyone can win, but with some well-invested effort, you can tip the scales considerably in your favor.

The math in here isn't easy. You don't need a PhD in game theory to understand the concepts in this book, but it's not as simple as memorizing starting hands or calculating the likelihood of making your flush on the river. There's some work involved. The people who want to believe intuition is enough aren't going to read this book. But the people who make the effort will be playing with a definite edge. In fact, much of my poker success is the result of using some of the most basic concepts addressed in this book.

Bill and Jerrod have saved you a lot of time. They've saved me a lot of a time, too. I get asked a lot of poker questions, and most are pretty easy to answer. But I've never had a good response when someone asks me to recommend a book for understanding game theory as it relates to poker. I usually end up explaining that there are good poker books and good game theory books, but no book addresses the relationship between the two.

Now I have an answer. And if I ever find myself teaching a poker class for the mathematics department at UCLA, this will be the only book on the syllabus.

Chris "Jesus" Ferguson
Champion, 2000 World Series of Poker
November 2006

Introduction

“If you think the math isn't important, you don't know the right math.”

Chris "Jesus" Ferguson, 2000 World Series of Poker champion

Introduction

In the late 1970s and early 1980s, the bond and option markets were dominated by traders who had learned their craft by experience. They believed that their experience and intuition for trading were a renewable edge; that is, that they could make money just as they always had by continuing to trade as they always had. By the mid-1990s, a revolution in trading had occurred; the old school grizzled traders had been replaced by a new breed of quantitative analysts, applying mathematics to the "art" of trading and making of it a science.

If the latest backgammon programs, based on neural net technology and mathematical analysis had played in a tournament in the late 1970s, their play would have been mocked as overaggressive and weak by the experts of the time. Today, computer analyses are considered to be the final word on backgammon play by the world's strongest players - and the game is fundamentally changed for it.

And for decades, the highest levels of poker have been dominated by players who have learned the game by playing it, "road gamblers" who have cultivated intuition for the game and are adept at reading other players' hands from betting patterns and physical tells. Over the last five to ten years, a whole new breed of player has risen to prominence within the poker community. Applying the tools of computer science and mathematics to poker and sharing information across the Internet, these players have challenged many of the assumptions that underlie traditional approaches to the game. One of the most important features of this new approach to the game is a reliance on quantitative analysis and the application of mathematics to the game. Our intent in this book is to provide an introduction to quantitative techniques as applied to poker and to the application of game theory, a branch of mathematics, to poker.

Any player who plays poker is using some model, no matter what methods he uses to inform it. Even if a player is not consciously using mathematics, a model of the situation is implicit in his decisions; that is, when he calls, raises, or folds, he is making a statement about the relative values of those actions. By preferring one action over another, he articulates his belief that one action is better than another in a particular situation. Mathematics are a particularly appropriate tool for making decisions based on information. Rejecting mathematics as a tool for playing poker puts one's decision-making at the mercy of guesswork.

Common Misconceptions

We frequently encounter players who dismiss a mathematical approach out of hand, often based on their misconceptions about what this approach is all about. We list a few of these here; these are ideas that we have heard spoken, even by fairly knowledgeable players. For each of these, we provide a brief rebuttal here; throughout this book, we will attempt to present additional refutation through our analysis.

1) By analyzing what has happened in the past - our opponents, their tendencies, and so on-we can obtain a permanent and recurring edge.

This misconception is insidious because it seems very reasonable; in fact, we *can* gain an edge over our opponents by knowing their strategies and exploiting them. But this edge can be only temporary; our opponents, even some of the ones we think play poorly, adapt and evolve by reducing the quantity and magnitude of clear errors they make and by attempting to counter-exploit us. We have christened this first misconception the "PlayStation™ theory of poker" - that the poker world is full of players who play the same fixed strategy, and the goal of playing poker is to simply maximize profit against the fixed strategies of our opponents. In fact, our opponents' strategies are dynamic, and so we must be dynamic; no edge that we have is necessarily permanent.

2) Mathematical play is predictable and lacks creativity.

In some sense this is true; that is, if a player were to play the optimal strategy to a game, his strategy would be "predictable" - but there would be nothing at all that could be done with this information. In the latter parts of the book, we will introduce the concept of *balance* this is the idea that each action sequence contains a mixture of hands that prevents the opponent from exploiting the strategy. Optimal play incorporates a precisely calibrated mixture of bluffs, semi-bluffs, and value bets that make it appear entirely unpredictable. "Predictable" connotes "exploitable," but this is not necessarily true. If a player has aces every time he raises, this is predictable and exploitable. However, if a player always raises when he holds aces, this is not necessarily exploitable as long as he also raises with some other hands. The opponent is not able to exploit sequences that contain other actions because it is unknown if the player holds aces.

3) Math is not always applicable; sometimes "the numbers go out the window."

This misconception is related to the idea that for any situation, there is only one mathematically correct play; players assume that even playing exploitively, there is a correct mathematical play - but that they have a "read" which causes them to prefer a different play. But this is simply a narrow definition of "mathematical play" - incorporating new information into our understanding of our opponent's distribution and utilizing that information to play more accurately is the major subject of Part II. In fact, mathematics contains tools (notably Bayes' theorem) that allow us to precisely quantify the degree to which new information impacts our thinking; in fact, playing mathematically is more accurate as far as incorporating "reads" than playing by "feel."

4) Optimal play is an intractable problem for real-life poker games; hence, we should simply play exploitively.

This is an important idea. It is true that we currently lack the computing power to solve heads-up holdem or other games of similar complexity. (We will discuss what it means to "solve" a game in Part III). We have methods that are known to find the answer, but they will not run on modern computers in any reasonable amount of time. "Optimal" play does not even exist for multiplayer games, as we shall see. But this does not prevent us from doing two things: attempting to create strategies which share many of the same properties as optimal strategies and thereby play in a "near-optimal" fashion; and also to evaluate candidate strategies and find out how far away from optimal they are by maximally exploiting them.

5) When playing [online, in a tournament, in high limit games, in low limit games...], you have to change your strategy completely to win.

This misconception is part of a broader misunderstanding of the idea of a "strategy" - it is in fact true that in some of these situations, you must take different actions, particularly exploitively, in order to have success. But this is not because the games are fundamentally different; it is because the other players play differently and so your responses to their play take different forms. Consider for a moment a simple example. Suppose you are dealt A9s on the button in a full ring holdem game. In a small-stakes limit holdem game, six players might limp to you, and you should raise. In a high limit game, it might be raised from middle position, and you would fold. In a tournament, it might be folded to you, and you would raise. These are entirely different actions, but the broader strategy is the same in all - choose the most profitable action.

Throughout this book, we will discuss a wide variety of poker topics, but overall, our ideas could be distilled to one simple piece of play advice: *Maximize average profit*. This idea is at the heart of all our strategies, and this is the one thing that doesn't change from game condition to game condition.

Psychological Aspects

Poker authors, when faced with a difficult question, are fond of falling back on the old standby, "It depends." - on the opponents, on one's 'read', and so on. And it is surely true that the most profitable action in many poker situations does in fact depend on one's sense, whether intuitive or mathematical, of what the opponent holds (or what he can hold). But one thing that is often missing from the qualitative reasoning that accompanies "It depends," is a real answer or a methodology for arriving at an action. In reality, the answer does in fact depend on our assumptions, and the tendencies and tells of our opponents are certainly something about which reasonable people can disagree. But once we have characterized their play into assumptions, the methods of mathematics take over and intuition fails as a guide to proper play.

Some may take our assertion that quantitative reasoning surpasses intuition as a guide to play as a claim that the psychological aspects of poker are without value. But we do not hold this view. The psychology of poker can be an absolutely invaluable tool for exploitive play, and the assumptions that drive the answers that our mathematical models can generate are often strongly psychological in nature. The methods by which we utilize the information that our intuition or people-reading skills give us is our concern here. In addition, we devote time to the question of what we ought to do when we are unable to obtain such information, and also in exposing some of the poor assumptions that often undermine the information-gathering efforts of intuition. With that said, we will generally, excepting a few specific sections, ignore physical tells and opponent profiling as being beyond the scope of this book and more adequately covered by other writers, particularly in the work of Mike Garo.

About This Book

We are practical people - we generally do not study poker for the intellectual challenge, although it turns out that there is a substantial amount of complexity and interest to the game. We study poker with mathematics because by doing so, we make more money. As a result, we are very focused on the practical application of our work, rather than on generating proofs or covering esoteric, improbable cases. This is not a mathematics textbook, but a primer on the application of mathematical techniques to poker and in how to turn the insights gained into increased profit at the table.

Certainly, there are mathematical techniques that can be applied to poker that are difficult and complex. But we believe that most of the mathematics of poker is really not terribly difficult, and we have sought to make some topics that may seem difficult accessible to players without a very strong mathematical background. But on the other hand, it is math, and we fear that if you are afraid of equations and mathematical terminology, it will be somewhat difficult to follow some sections. But the vast majority of the book should be understandable to anyone who has completed high school algebra. We will occasionally refer to results or conclusions from more advanced math. In these cases, it is not of prime importance that you understand exactly the mathematical technique that was employed. The important element is the concept - it is very reasonable to just "take our word for it" in some cases.

To help facilitate this, we have marked off the start and end of some portions of the text so that our less mathematical readers can skip more complex derivations. Just look for this icon for guidance, indicating these cases.

In addition,

Solution:

Solutions to example problems are shown in shaded boxes.

As we said, this book is not a mathematical textbook or a mathematical paper to be submitted to

a journal. The material here is not presented in the manner of formal proof, nor do we intend it to be taken as such. We justify our conclusions with mathematical arguments where necessary and with intuitive supplemental arguments where possible in order to attempt to make the principles of the mathematics of poker accessible to readers without a formal mathematical background, and we try not to be boring. The primary goal of our work here is not to solve game theory problems for the pure joy of doing so; it is to enhance our ability to win money at poker.

This book is aimed at a wide range of players, from players with only a modest amount of experience to world-class players. If you have never played poker before, the best course of action is to put this book down, read some of the other books in print aimed at beginners, play some poker, learn some more, and then return after gaining additional experience. If you are a computer scientist or options trader who has recently taken up the game, then welcome. This book is for you. If you are one of a growing class of players who has read a few books, played for some time, and believe you are a solid, winning player, are interested in making the next steps but feel like the existing literature lacks insight that will help you to raise your game, then welcome. This book is also for you. If you are the holder of multiple World Series of Poker bracelets who plays regularly in the big game at the Bellagio, you too are welcome. There is likely a fair amount of material here that can help you as well.

Organization

The book is organized as follows:

Part I: Basics, is an introduction to a number of general concepts that apply to all forms of gambling and other situations that include decision making under risk. We begin by introducing probability, a core concept that underlies all of poker. We then introduce the concept of a probability distribution, an important abstraction that allows us to effectively analyze situations with a large number of possible outcomes, each with unique and variable probabilities. Once we have a probability distribution, we can define expected value, which is the metric that we seek to maximize in poker. Additionally, we introduce a number of concepts from statistics that have specific, common, and useful applications in the field of poker, including one of the most powerful concepts in statistics, Bayes' theorem.

Part II: Exploitive Play, is the beginning of our analysis of poker. We introduce the concept of a toy game, which is a smaller, simpler game that we can solve in order to gain insight about analogous, more complicated games. We then consider examples of toy games in a number of situations. First we look at playing poker with the cards exposed and find that the play in many situations is quite obvious; at the same time, we find interesting situations with some counter-intuitive properties that are helpful in understanding full games. Then we consider what many authors treat as the heart of poker, the situation where we play our single hand against a distribution of the opponent's hands and attempt to exploit his strategy, or maximize our win against his play. This is the subject of the overwhelming majority of the poker literature. But we go further, to the (in our view) much more important case, where we are not only playing a single hand against the opponent, but playing an entire distribution of hands against his distribution of hands. It is this view of poker, we claim, that leads to truly strong play.

Part III: Optimal Play, is the largest and most important part of this book. In this part, we introduce the branch of mathematics called game theory. Game theory allows us to find optimal strategies for simple games and to infer characteristics of optimal strategies for more complicated games even if we cannot solve them directly. We do work on many variations of the AKQ game, a simple toy game originally introduced to us in Card Player magazine by Mike Caro. We then spend a substantial amount of time introducing and solving $[0,1]$ poker games, of the type introduced by John von Neumann and Oskar Morganstem in their seminal text on game

theory "Theory of Games and Economic Behavior (1944), but with substantially more complexity and relevance to real-life poker. We also explain and provide the optimal play solution to short-stack headsup no-limit holdem.

Part IV: Bankroll and Risk includes material of interest on a very important topic to anyone who approaches poker seriously. We present the risk of ruin model, a method for estimating the chance of losing a fixed amount playing in a game with positive expectation but some variance. We then extend the risk of ruin model in a novel way to include the uncertainty surrounding any observation of win rate. We also address topics such as the Kelly criterion, choosing an appropriate game level, and the application of portfolio theory to the poker metagame.

Part V: Other Topics includes material on other important topics, tournaments are the fastest-growing and most visible form of poker today; we provide an explanation of concepts and models for calculating equity and making accurate decisions in the tournament environment. We consider the game theory of multiplayer games, an important and very complex branch of game theory, and show some reasons why the analysis of such games is so difficult. In this section we also articulate and explain our strategic philosophy of play, including our attempts to play optimally or at least pseudo-optimally as well as the situations in which we play exploitively.

How This Book Is Different

This book differs from other poker books in a number of ways. One of the most prominent is in its emphasis on quantitative methods and modeling. We believe that intuition is often a valuable tool for understanding what is happening. But at the same time, we eschew its use as a guide to what action to take. We also look for ways to identify situations where our intuition is often wrong, and attempt to retrain it in such situations in order to improve the quality of our reads and our overall play. For example, psychologists have identified that the human brain is quite poor at estimating probabilities, especially for situations that occur with low frequency. By creating alternate methods for estimating these probabilities, we can gain an advantage over our opponents.

It is reasonable to look at each poker decision as a two-part process of gathering information and then synthesizing that information and choosing the right action. It is our contention that intuition has no place in the latter. Once we have a set of assumptions about the situation - how our opponent plays, what our cards are, the pot size, etc., then finding the right action is a simple matter of calculating expectation for the various options and choosing the option that maximizes this.

The second major way in which this book differs from other poker books is in its emphasis on *strategy*, contrasted to an emphasis on *decisions*. Many poker books divide the hand into sections, such as "preflop play," "flop play," "turn play" etc. By doing this, however, they make it difficult to capture the way in which a player's preflop, flop, turn, and river play are all intimately connected, and ultimately part of the same strategy. We try to look at hands and games in a much more organic fashion, where, as much as possible, the evaluation of expectation occurs not at each decision point but at the beginning of the hand, where a full strategy for the game is chosen. Unfortunately, holdem and other popular poker games are extraordinarily complex in this sense, and so we must sacrifice this sometimes due to computational infeasibility. But the idea of carrying a strategy forward through different betting rounds and being constantly aware of the potential hands *we* could hold at this point, which our fellow poker theorists Chris Ferguson and Paul R. Pudaite call "reading your own hand," is essential to our view of poker.

A third way in which this book differs from much of the existing literature is that it is not a book

about how to play poker. It is a book about how to think about poker. We offer very little in terms of specific recommendations about how to play various games; instead this book is devoted to examining the issues that are of importance in determining a strategy. Instead of a roadmap to how to play poker optimally, we instead try to offer a roadmap to how to think about optimal poker.

Our approach to studying poker, too, diverges from much of the existing literature. We often work on toy games, small, solvable games from which we hope to gain insight into larger, more complex games. In a sense, we look at toy games to examine dimensions of poker, and how they affect our strategy. How does the game change when we move from cards exposed to cards concealed? From games where players cannot fold to games where they can? From games where the first player always checks to games where both players can bet? From games with one street to games with two? We examine these situations by way of toy games - because toy games, unlike real poker, are solvable in practice - and attempt to gain insight into how we should approach the larger game.

Our Goals

It is our hope that our presentation of this material will provide at least two things; that it will aid you to play more strongly in your own poker endeavors and to think about situations in poker in a new light, and that it will serve as a jumping-off point toward the incredible amount of serious work that remains to be done in this field. Poker is in a critical stage of growth at this writing; the universe of poker players and the mainstream credibility of the game have never been larger. Yet it is still largely believed that intuition and experience are determining factors of the quality of play - just as in the bond and options markets in the early 1980s, trading was dominated by old-time veterans who had both qualities in abundance. A decade later, the quantitative analysts had grasped control of the market, and style and intuition were on the decline. In the same way, even those poker players regarded as the strongest in the world make serious errors and deviations from optimal strategies. This is not an indictment of their play but a reminder that the distance between the play of the best players in the world and the best play possible is still large, and that therefore there is a large amount of profit available to those who can bridge that gap.

Part I: Basics

“As for as the laws of mathematics refer to reality, they are not certain; as for as they are certain, they do not refer to reality.”

Albert Einstein

Chapter 1

Decisions Under Risk: Probability and Expectation

There are as many different reasons for playing poker as there are players who play the game. Some play for social reasons, to feel part of a group or "one of the guys," some play for recreation, just to enjoy themselves. Many play for the enjoyment of competition. Still others - lay to satisfy gambling addictions or to cover up other pain in their lives. One of the difficulties of taking a mathematical approach to these reasons is that it's difficult to quantify the value of having fun or of a sense of belonging.

In addition to some of the more nebulous and difficult to quantify reasons for playing poker, there may also be additional financial incentives not captured within the game itself. For example, the winner of the championship event of the World Series of Poker is virtually guaranteed to reap a windfall from endorsements, appearances, and so on, over and above the large first prize.

There are other considerations for players at the poker table as well; perhaps losing an additional hand would be a significant psychological blow. While we may criticize this view as irrational, it must still factor into any exhaustive examination of the incentives to play poker. Even if we restrict our inquiry to monetary rewards, we find that preference for money is non-linear. For most people, winning five million dollars is worth much more (or has much more *utility*) than a 50% chance of winning ten million; five million dollars is life-changing money for most, and the *marginal value* of the additional five million is much smaller.

In a broader sense, all of these issues are included in the *utility theory* branch of economics. Utility theorists seek to quantify the preferences of individuals and create a framework under which financial and non-financial incentives can be directly compared. In reality, it is utility that we seek to maximize when playing poker (or in fact, when doing any thing). However, the use of utility theory as a basis for analysis presents a difficulty; each individual has his own utility curves and so general analysis becomes extremely difficult.

In this book, we will therefore refrain from considering utility and instead use money won inside the game as a proxy for utility. In the bankroll theory section in Part IV, we will take an in-depth look at certain meta-game considerations, introduce such concepts as *risk of ruin*, *the Kelly criterion*, and *certainty equivalent*. All of these are measures of risk that have primarily to do with factors outside the game. Except when expressly stated, however, we will take as a premise that players are adequately bankrolled for the games they are playing in, and that their sole purpose is to maximize the money they will win by making the best decisions at every point.

Maximizing total money won in poker requires that a player maximize the *expected value* of his decisions. However, before we can reasonably introduce this cornerstone concept, we must first spend some time discussing the concepts of probability that underlie it. The following material owes a great debt to Richard Epstein's text *The Theory of Gambling and Statistical Logic* (1967), a valuable primer on probability and gambling.

Probability

Most of the decisions in poker take place under conditions where the outcome has not yet been determined. When the dealer deals out the hand at the outset, the players' cards are unknown, at least until they are observed. Yet we still have some information about the contents of the other players' hands. The game's rules constrain the contents of their hands-while a player may hold the jack-ten of hearts, he cannot hold the ace-prince of billiard tables, for example. The composition of the deck of cards is set before starting and gives us information about the hands.

Consider a holdem hand. What is the chance that the hand contains two aces? You may know the answer already, but consider what the answer means. What if we dealt a million hands just like this? How many pairs of aces would there be? What if we dealt ten million? Over time and many trials, the ratio of pairs of aces to total hands dealt will converge on a particular number. We define **probability** as this number. Probability is the key to decision-making in poker as it provides a mathematical framework by which we can evaluate the likelihood of uncertain events.

If n trials of an experiment (such as dealing out a holdem hand) produce *no* occurrences of an event x , we define the probability p of x occurring $p(x)$ as follows:

$$p(x) = \lim_{n \rightarrow \infty} \frac{n_0}{n} \tag{1.1}$$

Now it happens to be the case that the likelihood of a single holdem hand being a pair of aces is $1/221$. We could, of course, determine this by dealing out ten billion hands and observing the ratio of pairs of aces, to total hands dealt. This, however, would be a lengthy and difficult process, and we can do better by breaking the problem up into components. First we consider just one card. What is the probability that a single card is an ace? Even this problem can be broken down further - what is the probability that a single card is the ace of spades?

This final question can be answered rather directly. We make the following assumptions:

- There are fifty-two cards in a standard deck.
- Each possible card is equally likely.

Then the probability of any particular card being the one chosen is $1/52$. If the chance of the card being the ace of spades is $1/52$, what is the chance of the card being any ace? This is equivalent to the chance that the card is the ace of spades OR that it is the ace of hearts OR that it is the ace of diamonds OR that it is the ace of clubs. There are four aces in the deck, each with a $1/52$ chance of being the card, and summing these probabilities, we have:

$$p(A) = (4) \left(\frac{1}{52} \right)$$

$$p(A) = \left(\frac{1}{13} \right)$$

We can sum these probabilities directly because they are **mutually exclusive**; that is, no card can simultaneously be both the ace of spades and the ace of hearts. Note that the probability $1/13$ is exactly equal to the ratio (number of aces in the deck)/(number of cards total). This relationship holds just as well as the summing of the individual probabilities.

Independent Events

Some events, however, are not mutually exclusive. Consider for example, these two events:

1. The card is a heart
2. The card is an ace.

If we try to figure out the probability that a single card is a heart OR that it is an ace, we find there are thirteen hearts in the deck out of fifty-cards, so the chance that the card is a heart is $1/4$. The chance that the card is an ace is, as before, $1/13$. However, we cannot simply add these probabilities as before, because it is possible for a card to both be an ace and a heart.

There are two types of relationships between events. The first type are events that have no effect

on each other. For example, the closing value of the NASDAQ stock index and the value of the dice on a particular roll at the craps table in a casino in Monaco that evening are basically unrelated events; neither one should impact the other in any way that is not negligible. If the probability of both events occurring equals the product of the individual probabilities, then the events are said to be *independent*. The probability that both A and B occur is called the *joint probability* of A and B.

In this case, the joint probability of a card being both a heart and an ace is $(\frac{1}{13})(\frac{1}{4})$, or $\frac{1}{52}$. This is because the fact that the card is a heart does not affect the chance that it is an ace - all four suits have the same set of cards.

Independent events are not mutually exclusive except when one of the events has probability zero. In this example, the total number of hearts in the deck is thirteen, and the total of aces in the deck is four. However, by adding these together, we are double-counting one single card (the ace of hearts). There are actually thirteen hearts and three other aces, or if you prefer, four aces, and twelve other hearts. It turns out that the general application of this concept is that the probability that at least one of two mutually non-exclusive events A and B will occur is the sum of the probabilities of A and B minus the joint probability of A and B. So the probability of the card being a heart *or* an ace is equal to the chance of it being a heart ($\frac{1}{4}$) plus the chance of it being an ace minus the chance of it being both ($\frac{1}{52}$), or $\frac{4}{13}$. This is true for all events, independent or dependent.

Dependent Events

Some events, by contrast, do have impacts on each other. For example, before a baseball game, a certain talented pitcher might have a 3% chance of pitching nine innings and allowing no runs, while his team might have a 60% chance of winning the game. However, the chance of the pitcher's team winning the game and him also pitching a shutout is obviously not 60% times 3%. Instead, it is very close to 3% itself, because the pitcher's team will virtually always win the game when he accomplishes this. These events are called *dependent*. We can also consider the *conditional probability* of A given B, which is the chance that *if* B happens, A will also happen. The probability of A and B both occurring for dependent events is equal to the probability of A multiplied by the conditional probability of B given A. Events are independent if the conditional probability of A given B is equal to the probability of A alone.

Summarizing these topics more formally, if we use the following notation:

$p(A \cup B)$ = Probability of A or B occurring.

$p(A \cap B)$ = Probability of A and B occurring.

$p(A|B)$ = Conditional probability of A occurring given B has already occurred.

The \cup and \cap notations are from set theory and formally represent "union" and "intersection." We prefer the more mundane terms "or" and "and." Likewise, $|$ is the symbol for "given," so we pronounce these expressions as follows:

$p(A \cup B)$ = "p of A or B"

$p(A \cap B)$ = "p of A and B"

$p(A|B)$ = "p of A given B"

Then for mutually exclusive events:

$$p(A \cup B) = p(A) + p(B) \tag{1.2}$$

For independent events:

$$p(A \cap B) = p(A)p(B) \quad (1.3)$$

For all events:

$$p(A \cup B) = p(A) + p(B) - p(A \cap B) \quad (1.4)$$

For dependent events:

$$p(A \cap B) = p(A)p(B|A) \quad (1.5)$$

Equation 1.2 is simply a special case of Equation 1.4 for mutually exclusive events, $p(A \cap B) = 0$. Likewise, Equation 1.3 is a special case of Equation 1.5, as for independent events, $p(B|A) = p(B)$. Additionally, if $p(B|A) = p(B)$, then $p(A|B) = p(A)$.

We can now return to the question at hand. How frequently will a single holdem hand dealt from a full deck contain two aces? There are two events here:

- A: The first card is an ace.
- B: The second card is an ace.

$p(A) = 1/13$, and $p(B) = 1/13$ as well. However, these two events are dependent, if A occurs (the first card is an ace), then it is less likely that B will occur, as the cards are dealt without replacement. So $p(B|A)$ is the chance that the second card is an ace given that the first card is an ace. There are three aces remaining, and fifty-one possible cards, so $p(B|A) = 3/51$, or $1/17$.

$$\begin{aligned} p(A \cap B) &= p(A)p(B|A) \\ p(A \cap B) &= (1/13)(1/17) \\ p(A \cap B) &= 1/221 \end{aligned}$$

There are a number of other simple properties that we can mention about probabilities. First, the probability of any event is at least zero and no greater than one. Referring back to the definition of probability, n trials will never result in more than n occurrences of the event, and never less than zero occurrences. The probability of an event that is certain to occur is one. The probability of an event that never occurs is zero. The probability of an event's **complement** -that is, the chance that an event does not occur, is simply one minus the event's probability.

Summarizing, if we use the following notation:

$$p(\bar{A}) = \text{Probability that } A \text{ does not occur.}$$

C = a certain event

I = an impossible event

Then we have:

$$0 \leq p(A) \leq 1 \text{ for any } A \quad (1.6)$$

$$p(C) = 1 \quad (1.7)$$

$$p(I) = 0 \quad (1.8)$$

$$p(A) + p(\bar{A}) = 1 \quad (1.9)$$

Equation 1.9 can also be restated as:

$$p(A) = 1 - p(\bar{A}) \quad (1.10)$$

We can solve many probability problems using these rules

Some common questions of probability are simple, such as the chance of rolling double sixes on two dice, in terms of probability, this can be stated using equation 1.3, since the die rolls are independent. Let $p(A)$ be the probability of rolling a six on the first die and $p(B)$ be the probability of rolling a six on the second die. Then:

$$\begin{aligned} p(A \cap B) &= p(A)p(B) \\ p(A \cap B) &= (1/6)(1/6) \\ p(A \cap B) &= 1/36 \end{aligned}$$

Likewise, using equation 1.2, the chance of a single player holding aces, kings, or queens becomes:

$$\begin{aligned} p(AA) &= 1/221 \\ p(KK) &= 1/221 \\ p(QQ) &= 1/221 \\ p(\{AA, KK, QQ\}) &= p(AA) + p(KK) + p(QQ) = 3/221 \end{aligned}$$

Additionally we can solve more complex questions, such as:

How likely is it that a suited hand will flop a flush?

We hold two of the flush suit, leaving eleven in the deck. All three of the cards must be of the flush suit, meaning that we have A — the first card being a flush card, B — the second card being a flush card given that the first card is a flush card, and C = the third card being a flush card given that both of the first two are flush cards.

$$\begin{aligned} p(A) &= 11/50 && \text{(two cards removed from the deck in the player's hand)} \\ p(B|A) &= 10/49 && \text{(one flush card and three total cards removed)} \\ p(C|(A \cap B)) &= 9/48 && \text{(two flush cards and four total cards removed)} \end{aligned}$$

Applying equation 1.5, we get:

$$\begin{aligned} p(A \cap B) &= p(A)p(B|A) \\ p(A \cap B) &= (1/50)(10/49) \\ p(A \cap B) &= 11/245 \end{aligned}$$

Letting $D = (A \cap B)$, we can use equation 1.5 again:

$$\begin{aligned} p(D \cap C) &= p(D)p(C|D) \\ p(A \cap B \cap C) &= p(A \cap B)p(C|(A \cap B)) \\ p(A \cap B \cap C) &= (11/245)(9/48) \\ p(A \cap B \cap C) &= 33/3920, \text{ or a little less than 1\%.} \end{aligned}$$

We can apply these rules to -virtually any situation, and throughout the text we will use these properties and rules to calculate probabilities for single events.

Probability Distributions

Though single event probabilities are important, it is often the case that they are inadequate to fully analyze a situation. Instead, it is frequently important to consider many different probabilities at the same time. We can characterize the possible outcomes and their probabilities from an event as a probability distribution.

Consider a fair coin flip. The coin flip has just two possible outcomes - each outcome is mutually exclusive and has a probability of $1/2$. We can create a probability distribution for the coin flip by taking each outcome and pairing it with its probability. So we have two pairs: (heads, $1/2$) and (tails, $1/2$).

If C is the probability distribution of the result of a coin flip, then we can write this as:

$$C = \{(\text{heads}, 1/2), (\text{tails}, 1/2)\}$$

Likewise, the probability distribution of the result of a fair six-sided die roll is:

$$D = \{(1, 1/6), (2, 1/6), (3, 1/6), (4, 1/6), (5, 1/6), (6, 1/6)\}$$

We can construct a discrete probability distribution for any event by enumerating an exhaustive and mutually exclusive list of possible outcomes and pairing these outcomes with their corresponding probabilities.

We can therefore create different probability distributions from the same physical event. From our die roll we could also create a second probability distribution, this one the distribution of the odd-or-evenness of the roll:

$$D' = \{(\text{odd}, 1/2), (\text{even}, 1/2)\}$$

In poker, we are almost always very concerned with the contents of our opponents' hands. But it is seldom possible to narrow down our estimate of these contents to a single pair of cards. Instead, we use a probability distribution to represent the hands he could possibly hold and the corresponding probabilities that he holds them. At the beginning of the hand, before anyone has looked at their cards, each player's probability distribution of hands is identical. As the hand progresses, however, we can incorporate new information we gain through the play of the hand, the cards in our own hand, the cards on the board, and so on, to continually refine the probability estimates we have for each possible hand.

Sometimes we can associate a numerical value with each element of a probability distribution. For example suppose that a friend offers to flip a fair coin with you. The winner will collect \$10 from the loser. Now the results of the coin flip follow the probability distribution we identified earlier:

$$C = \{(\text{heads}, 1/2), (\text{tails}, 1/2)\}$$

Since we know the coin is fair, it doesn't matter who calls the coin or what they call, so we can identify a second probability distribution that is the result of the bet:

$$C' = \{(\text{win}, 1/2), (\text{lose}, 1/2)\}$$

We can then go further, and associate a numerical value with each result. If we win the flip, our friend pays us \$10. If we lose the flip, then we pay him \$10. So we have the following:

$$B = \{(+\$10, 1/2), (-\$10, 1/2)\}$$

When a probability distribution has numerical values associated with each of the possible outcomes, we can find the *expected value (EV)* of that distribution, which is the value of each outcome multiplied by its probability, all summed together. Throughout the text, we will use the notation $\langle X \rangle$ to denote "the expected value of X." For this example, we have:

$$\begin{aligned}\langle B \rangle &= (1/2)(+\$10) + (1/2)(-\$10) \\ \langle B \rangle &= \$5 + (-\$5) \\ \langle B \rangle &= 0\end{aligned}$$

Hopefully this is intuitively obvious - if you flip a fair coin for some amount, half the time you win and half the time you lose. The amounts are the same, so you break even on average. Also, the EV of declining your friend's offer by not flipping at all is also zero, because no money changes hands.

For a probability distribution P , where each of the n outcomes has a value x_i and a probability then p_i then P 's expected value $\langle P \rangle$ is:

$$\langle P \rangle = \sum_{i=1}^n p_i x_i \tag{1.11}$$

At the core of winning at poker or at any type of gambling is the idea of maximizing expected value. In this example, your friend has offered you a fair bet. On average, you are no better or worse off by flipping with him than you are by declining to flip.

Now suppose your friend offers you a different, better deal. He'll flip with you again, but when you win, he'll pay you \$11, while if he wins, you'll only pay him \$10. Again, the EV of not flipping is 0, but the EV of flipping is not zero any more. You'll win \$11 when you win but lose \$10 when you lose. Your expected value of this new bet B_n is:

$$\begin{aligned}\langle B_n \rangle &= (1/2)(+\$11) + (1/2)(-\$11) \\ \langle B_n \rangle &= \$0.50\end{aligned}$$

On average here, then, you will win fifty cents per flip. Of course, this is not a guaranteed win; in fact, it's impossible for you to win 50 cents on any particular flip. It's only in the aggregate that this expected value number exists. However, by doing this, you will average fifty cents better than declining.

As another example, let's say your same friend offers you the following deal. You'll roll a pair of dice once, and if the dice come up double sixes, he'll pay you \$30, while if they come up any other number, you'll pay him \$1. Again, we can calculate the EV of this proposition.

$$\begin{aligned}\langle B_d \rangle &= (+\$30)(1/36) + (-\$1)(35/36) \\ \langle B_d \rangle &= \$ 30/36 - \$ 35/36 \\ \langle B_d \rangle &= -\$ 5/36 \quad \text{or about 14 cents.}\end{aligned}$$

The value of this bet to you is about negative 14 cents. The EV of not playing is zero, so this is a bad bet and you shouldn't take it. Tell your friend to go back to offering you 11-10 on coin flips. Notice that this exact bet is offered on craps layouts around the world.

A very important property of expected value is that it is *additive*. That is, the EV of six different bets in a row is the sum of the individual EVs of each bet individually. Most gambling games - most things in life, in fact, are just like this. We are continually offered little coin flips or dice rolls - some with positive expected value, others with negative expected value. Sometimes the event in question isn't a die roll or a coin flip, but an insurance policy or a bond fund. The free drinks and neon lights of Las Vegas are financed by the summation of millions of little coin flips, on each of which the house has a tiny edge. A skillful poker player takes advantage of this additive property of expected value by constantly taking advantage of favorable EV situations.

In using probability distributions to discuss poker, we often omit specific probabilities for each hand. When we do this, it means that the relative probabilities of those hands are unchanged from their probabilities at the beginning of the hand. Supposing that we have observed a very tight player raise and we know from our experience that he raises if and only if he holds aces, kings, queens, or ace-king, we might represent his distribution of hands as:

$$H = \{AA, KK, QQ, AKs, AKo\}$$

The omission of probabilities here simply implies that the relative probabilities of these hands are as they were when the cards were dealt. We can also use the $\langle X \rangle$ notation for situations where we have more than one distribution under examination. Suppose we are discussing a poker situation where two players A and B have hands taken from the following distributions:

$$A = \{AA, KK, QQ, JJ, AKo, AKs\}$$

$$B = \{AA, KK, QQ\}$$

We have the following, then:

$\langle A, B \rangle$: the expectation for playing the distribution A against the distribution B.

$\langle A, AA|B \rangle$: the expectation for playing the distribution A against the hand AA taken from the distribution B.

$\langle AA|A, AA|B \rangle$: the expectation for playing AA from A against AA from B.

$\langle A, B \rangle = p(AA) \langle A, AA|B \rangle + p(KK) \langle A, KK|B \rangle + p(QQ) \langle A, QQ|B \rangle$...and so on.

Additionally, we can perform some basic arithmetic operations on the elements of a distribution. For example, if we multiply all the values of the outcomes of a distribution by a real constant, the expected value of the resulting distribution is equal to the expected value of the original distribution multiplied by the constant. Likewise, if we add a constant to each of the values of the outcomes of a distribution, the expected value of the resulting distribution is equal to the expected value of the original distribution plus the constant.

We should also take a moment to describe a common method of expressing probabilities, *odds*. Odds are defined as the ratio of the probability of the event not happening to the probability of the event happening. These odds may be scaled to any convenient base and are commonly expressed as "7 to 5," "3 to 2," etc. *Shorter* odds are those where the event is more likely: *longer* odds are those where the event is less likely. Often, relative hand values might be expressed this way: "That hand is a 7 to 3 favorite over the other one," meaning that it has a 70% of winning, and so on.

Odds are usually more awkward to use than probabilities in mathematical calculations because they cannot be easily multiplied by outcomes to yield expectation. True "gamblers" often use odds, because odds correspond to the ways in which they are paid out on their bets. Probability is more of a mathematical concept. Gamblers who utilize mathematics may use either, but often

prefer probabilities because of the ease of converting probabilities to expected value.

Key Concepts

- The probability of an outcome of an event is the ratio of that outcome's occurrence over an arbitrarily large number of trials of that event.
- A probability distribution is a pairing of a list of complete and mutually exclusive outcomes of an event with their corresponding probabilities.
- The expected value of a valued probability distribution is the sum of the probabilities of the outcomes times their probabilities.
- Expected value is additive.
- If each outcome of a probability distribution is mapped to numerical values, the expected value of the distribution is the summation of the products of probabilities and outcomes.
- A mathematical approach to poker is concerned primarily with the maximization of expected value.

Chapter 2

Predicting the Future: Variance and Sample Outcomes

Probability distributions that have values associated with the elements have two characteristics which, taken together, describe most of the behavior of the distribution for repeated trials. The first, described in the previous chapter, is expected value. The second is **variance**, a measure of the dispersion of the outcomes from the expectation. To characterize these two terms loosely, expected value measures how much you will win on average; variance measures how far your specific results may be from the expected value.

When we consider variance, we are attempting to capture the range of outcomes that can be expected from a number of trials. In many fields, the range of outcomes is of particular concern on both sides of the mean. For example, in many manufacturing environments there is a band of acceptability and outcomes on either side of this band are undesirable. In poker, there is a tendency to characterize variance as a one-sided phenomenon, because most players are unconcerned with outcomes that result in winning much more than expectation. In fact, "variance" is often used as shorthand for negative swings.

This view is somewhat practical, especially for professional players, but creates a tendency to ignore positive results and to therefore assume that these positive outcomes are more representative of the underlying distribution than they really are. One of the important goals of statistics is to find the probability of a certain measured outcome given a set of initial conditions, and also the inverse of this - inferring the initial conditions from the measured outcome. In poker, both of these are of considerable use. We refer to the underlying distribution of outcomes from a set of initial conditions as the **population** and the observed outcomes as the **sample**. In poker, we often cannot measure all the elements of the population, but must content ourselves with observing samples.

Most statistics courses and texts provide material on probability as well as a whole slew of sampling methodologies, hypothesis tests, correlation coefficients, and so on. In analyzing poker we make heavy use of probability concepts and occasional use of other statistical methods. What follows is a quick-and-dirty overview of some statistical concepts that are useful in analyzing poker, especially in analyzing observed results. Much information deemed to be irrelevant is omitted from the following and we encourage you to consult statistics textbooks for more information on these topics.

A commonly asked question *in* poker is "How often should I expect to have a winning session?" Rephrased, this question is "what is the chance that a particular sample taken from a population that consists of my sessions in a game will have an outcome > 0 ?" The most straightforward method of answering this question would be to examine the probability distribution of your sessions in that game and sum the probabilities of all those outcomes that are greater than zero.

Unfortunately, we do not have access to that distribution - no matter how much data you have collected about your play in that game from the past, all you have is a sample. However, suppose that we know somehow your per-hand expectation and variance in the game, and we know how long the session you are concerned with is. Then we can use statistical methods to estimate the probability that you will have a winning session. The first of these items, expected value (which we can also call the **mean** of the distribution) is familiar by now; we discussed it in Chapter 1.

Variance

The second of these measures, **variance**, is a measure of the deviation of outcomes from the expectation of a distribution. Consider two bets, one where you are paid even money on a coin

flip, and one where you are paid 5 to 1 on a die roll, winning only when the die comes up 6. Both of these distributions have an EV of 0, but the die roll has significantly higher variance. $1/16$ of the time, you get a payout that is 5 units away from the expectation, while $5/6$ of the time you get a payout that is only 1 unit away from the expectation. To calculate variance, we first square the distances from the expected value, multiply them by the probability they occur, and sum the values.

For a probability distribution P , where each of the n outcomes has a value x_i and a probability p_i , then the variance of P , V_p is:

$$V_p = \sum_{i=1}^n p_i (x_i - \langle P \rangle)^2 \quad (2.1)$$

Notice that because each term is squared and therefore positive, variance is always positive. Reconsidering our examples, the variance of the coinflip is:

$$V_C = (1/2)(1-0)^2 + (1/2)(-1-0)^2$$

$$V_C = 1$$

While the variance of the die roll is:

$$V_D = (5/6)(-1-0)^2 + (1/6)(5-0)^2$$

$$V_D = 5$$

In poker, a loose-wild game will have much higher variance than a tight-passive game, because the outcomes will be further from the mean (pots you win will be larger, but the money lost in pots you lose will be greater). Style of play will also affect your variance; thin value bets and semi-bluff raises are examples of higher-variance plays that might increase variance, expectation, or both. On the other hand, loose-maniacal players may make plays that increase their variance while decreasing expectation. And playing too tightly may reduce both quantities. In Part IV, we will examine bankroll considerations and risk considerations and consider a framework by which variance can affect our utility value of money. Except for that part of the book, we will ignore variance as a decision-making criterion for poker decisions. In this way variance is for us only a descriptive statistic, not a prescriptive one (as expected value is).

Variance gives us information about the expected distance from the mean of a distribution. The most important property of variance is that it is directly additive across trials, just as expectation is. So if you take the preceding dice bet twice, the variance of the two bets combined is twice as large, or 10.

Expected value is measured in units of expectation per event; by contrast, variance is measured in units of expectation squared per event squared. Because of this, it is not easy to compare variance to expected value directly. If we are to compare these two quantities, we must take the square root of the variance, which is called the **standard deviation**. For our dice example, the standard deviation of one roll is $\sqrt{5} \approx 2.23$. We often use the Greek letter σ (sigma) to represent standard deviation, and by extension σ^2 is often used to represent variance in addition to the previously utilized V .

$$\sigma = \sqrt{V} \quad (2.2)$$

$$\sigma^2 = V \quad (2.3)$$

The Normal Distribution

When we take a single random result from a distribution, it has some value that is one of the possible outcomes of the underlying distribution. We call this a **random variable**. Suppose we flip a coin. The flip itself is a random variable. Suppose that we label the two outcomes 1 (heads) and 0 (tails). The result of the flip will then either be 1 (half the time) or 0 (half the time). If we take multiple coin flips and sum them up, we get a value that is the summation of the outcomes of the random variable (for example, heads), which we call a **sample**. The sample value, then, will be the number of heads we flip in whatever the size of the sample.

For example, suppose we recorded the results of 100 coinflips as a single number - the total number of heads. The expected value of the sample will be 50, as each flip has an expected value of 0.5.

The variance and standard deviation of a single flip are:

$$\begin{aligned}\sigma^2 &= (1/2)(1 - 1/2)^2 + (1/2)(0 - 1/2)^2 \\ \sigma^2 &= 1/4 \\ \sigma &= 1/2\end{aligned}$$

From the previous section, we know also that the variance of the flips is additive.

So the variance of 100 flips is 25.

Just as an individual flip has a standard deviation, a sample has a standard deviation as well. However, unlike variance, standard deviations are not additive. But there is a relationship between the two.

For N trials, the variance will be:

$$\begin{aligned}\sigma^2_N &= N\sigma^2 \\ \sigma_N &= \sigma\sqrt{N}\end{aligned}\tag{2.4}$$

The square root relationship of trials to standard deviation is an important result, because it shows us how standard deviation scales over multiple trials. If we flip a coin once, it has a standard deviation of $1/2$. If we flip it 100 times, the standard deviation of a sample containing 100 trials is not 50, but 5, the square root of 100 times the standard deviation of one flip. We can see, of course, that since the variance of 100 flips is 25, the standard deviation of 100 flips is simply the square root, 5.

The distribution of outcomes of a sample is itself a probability distribution, and is called the **sampling distribution**. An important result from statistics, the **Central Limit Theorem**, describes the relationship between the sampling distribution and the underlying distribution. What the Central Limit Theorem says is that as the size of the sample increases, the distribution of the aggregated values of the samples converges on a special distribution called the **normal distribution**.

The normal distribution is a bell-shaped curve where the peak, of the curve is at the population mean, and the tails asymptotically approach zero as the x-values go to negative or positive infinity. The curve is also scaled by the standard deviation of the population. The total area under the curve of the normal distribution (as with all probability distributions) is equal to 1, and the area under the curve on the interval $[x_1, x_2]$ is equal to the probability that a particular result will

fall between x_1 and x_2 . This area is marked region A in Figure 2.1.




Figure 2.1. Std. Normal Dist, A = p(event between x_1 and x_2)

A little less formally, the Central Limit Theorem says that if you have some population and take a lot of big enough samples (how big depends on the type of data you're looking at), the outcomes of the samples will follow a bell-shaped curve around the mean of the population with a variance that's related to the variance of the underlying population.

The equation of the normal distribution function of a distribution with mean μ and standard deviation σ is:

$$N(x, \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \quad (2.5)$$

Finding the area between two points under the normal distribution curve gives us the probability that the value of a sample with the corresponding mean and variance will fall between those two points. The normal distribution is symmetric about the mean, so $1/2$ of the total area is to the right of the mean, and $1/2$ is to the left. A usual method of calculating areas under the normal curve involves creating for each point something called a **z-score**, where $z = (x - \mu)/\sigma$. This z-score represents the number of standard deviations that the outcome x is away from the mean.

$$z = (x - \mu)/\sigma \quad (2.6)$$

We can then find something called the **cumulative normal distribution** for a z-score z , which is the area to the left of z under the curve (where the mean is zero and the standard deviation is 1). We call this function $\Phi(z)$. See Figure 2.2

If z is a normalized z-score value, then the cumulative normal distribution function for z is:

$$\Phi(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^z \exp\left(-\frac{x^2}{2}\right) \quad (2.7)$$




Figure 2.2, Cumulative normal distribution

Finding the area between two values x_1 and x_2 is done by calculating the z-scores z_1 and z_2 for x_1 and x_2 , finding the cumulative normal distribution values $\Phi(z_1)$ and $\Phi(z_2)$ and subtracting them.

If $\Phi(z)$ is the cumulative normal distribution function for a z-score of z , then the probability that a sample taken from a normal distribution function with mean μ and standard deviation σ will fall between two z-scores x_1 and x_2 is:

$$p = \Phi\left(\frac{x_1 - \mu}{\sigma}\right) - \Phi\left(\frac{x_2 - \mu}{\sigma}\right) \quad (2.8)$$

Statisticians have created tables of $\Phi(z)$ values and spreadsheet programs can derive these values as well. Some key values are well known to students of statistics. The area between $z = -1$ and $z = +1$ is approximately 0.68; the area between $z = -2$ and $z = +2$ is approximately 0.955, and the area between $z = -3$ and $z = +3$ is approximately 0.997.

These values mean that the value of a sample from a normal distribution will fall:

Between $(\mu - \sigma)$ and $(\mu + \sigma)$ of the mean 68% of the time.

Between $(\mu - 2\sigma)$ and $(\mu + 2\sigma)$ of the mean 95.5% of the time.

Between $(\mu - 3\sigma)$ and $(\mu + 3\sigma)$ of the mean 99.7% of the time.

An example may help to make this clear. Let us modify the parameters of the die roll game we discussed before. A die is thrown, and the player loses 1 unit on 1-5, but wins 6 units on a roll of

6. We'll call this new game D_2 .

The expectation of the game is:

$$\begin{aligned} \langle D_2 \rangle &= \left(\frac{5}{6}\right)(-1) + \left(\frac{1}{6}\right)(6) \\ \langle D_2 \rangle &= \frac{1}{6} \text{ units /trial} \end{aligned}$$

When the player wins, he gains six units. Subtracting the mean value of $\frac{1}{6}$ from this outcome, we obtain:

$$\begin{aligned} V_{win} &= \left(6 - \frac{1}{6}\right)^2 \\ V_{win} &= \left(\frac{35}{6}\right)^2 \end{aligned}$$

Likewise, when he loses he loses one unit. Subtracting the mean value of $\frac{1}{6}$ from this, we have:

$$\begin{aligned} V_{lose} &= \left(-1 - \frac{1}{6}\right)^2 \\ V_{lose} &= \left(-\frac{7}{6}\right)^2 \end{aligned}$$

The variance of the game is:

$$\begin{aligned} V_{D_2} &= p(\text{win})(V_{win}) + p(\text{lose})(V_{lose}) \\ V_{D_2} &= \left(\frac{1}{6}\right)\left(\frac{35}{6}\right)^2 + \left(\frac{5}{6}\right)\left(-\frac{7}{6}\right)^2 \\ V_{D_2} &\approx 6.806 \text{ units}^2/\text{trial}^2 \end{aligned}$$

Suppose we toss the die 200 times. What is the chance that the player will win overall in 200 tosses? Win 40 units or more? Win 100 units or more?

We can solve this problem using the techniques we have just summarized. We first calculate the standard deviation, sometimes called the standard error, of a sample of 200 trials. This will be:

$$\sigma = \sqrt{V} = \sqrt{6.806} = 2.61 \text{ units/trial}$$

Applying equation 2.4 we get:

$$\begin{aligned} \sigma_N &= \sigma\sqrt{N} \\ \sigma_{200} &= 2.61\sqrt{200} = 36.89 \text{ units/200 trials} \end{aligned}$$

For 200 trials, the expected value, or mean (μ) of this distribution is $\frac{1}{6}$ units/trial times 200 trials or 33.33 units. Using Equation 2.6, we find the z-score of the point 0 as:

$$\begin{aligned} z_x &= (x-\mu)/\sigma && \text{where } x=0 \\ z_0 &= (0-33.33)/36.89 \\ z_0 &= -33.33/36.89 \\ z_0 &= -0.9035 \end{aligned}$$

Consulting a probability table in a statistics textbook, we find that the probability that an observed outcome will lie to the left of this z-score, $\Phi(-0.9035)$ is 0.1831, Hence, there is an 18.31% chance that the player will be behind after 200 trials.

To find the chance of being 40 units ahead, we find that point's z-score:

$$z_{40} = (40-33.33)/(36.89) = 0.1807$$

$$\Phi(0,1807) = 0,5717$$

But $\Phi(0.1807)$ is the probability that the observed outcome lies to the left of 40 units, or that we lose at least 40 units. To find the probability that we are to the right of this value, or are ahead 40 units, we must actually find $1 - \Phi(0.1807)$.

$$1 - \Phi(0.1807) = 1 - 0.5717$$

$$= 0.4283$$

So there is a 42.83% chance of being ahead at least 40 units after 200 tosses.

And similarly for 100 units ahead:

$$z_{100} = (100-33.33)/(36.89) = 1.8070$$

From a probability table we find that $\Phi(1.8070) = 0.9646$. Thus, for:

$$p = 1 - \Phi(1.8070)$$

$$p = 0.0354$$

The probability of being 100 units ahead after 200 tosses is 3.54%.

These values, however, are only approximations; the distribution of 200-roll samples is not quite normal. We can actually calculate these values directly with the aid of a computer. Doing this yields:

	Direct Calculation	Normal Approx
Chance of being ahead after 200 trials:	81.96%	81.69%
Chance of being ahead at least 40 units:	40.46%	42.83%
Chance of being ahead at least 100 units:	4.44%	3,54%

As you can see, these values have slight differences. Much of this is caused by the fact that the direct calculations have only a discrete amount of values. 200 trials of this game can result in outcomes of +38 and +45, but not +39 or +42, because there is only one outcome possible from winning, a +6. The normal approximation assumes that all values are possible.

Using this method, we return to the question posed at the beginning of the chapter: "How often will I have a winning session?" To solve this problem, we need to know the player's expectation per hand, his variance per hand and the size of his sessions. Consider a player whose win rate in some particular game is approximately 0.015 bets per hand, and whose standard deviation over the same interval is 2 bets per hand. Assume this player plans to play sessions of 300 hands (these would be roughly nine-hour sessions in live poker; perhaps only two to three hours online). How often should this player expect to win (have a result greater than 0) in a session?

First we can calculate the expected value of a sample, μ_N :

$$\mu_N = N\mu$$

$$\mu_N = (300)(0.015)$$

$$\mu_N = 4.5$$

Second, we calculate his standard deviation for a 300-hand session.

$$\begin{aligned}\sigma_N &= \sigma\sqrt{N} \\ \sigma_{300} &= (2)(\sqrt{300}) \\ \sigma_{300} &= 34.6\end{aligned}$$

Next, we find the z-score and $\Phi(z)$ for this value:

$$\begin{aligned}z_x &= (x - \mu_N) / \sigma \\ z_0 &= (0 - 4.5) / 34.6 \\ z_0 &= -0.1299\end{aligned}$$

From a probability table we find:

$$\begin{aligned}\Phi(-0.1299) &= 44.83\% \\ p &= 1 - \Phi(-0.1299) \\ p &= 1 - 0.4483 \\ p &= 0.55171\end{aligned}$$

This indicates that the player has a result of 0 or less 44.83% of the time - or correspondingly has a winning session 55.17% of the time. In reality, players may change their playing habits in order to produce more winning sessions for psychological reasons, and in reality "win rate" fluctuates significantly based on game conditions and for other reasons. But a player who simply plays 300 hands at a time with the above performance metrics should expect to be above zero just over 55% of the time.

A more extreme example: A player once told one of the authors that he had kept a spreadsheet where he marked down every AQ vs. AK all-in preflop confrontation that he saw online over a period of months and that after 2,000 instances, they were running about 50-50.

How likely is this, assuming the site's cards were fair and we have no additional information based on the hands that were folded?

First, let's consider the variance or standard deviation of a single confrontation.

AK vs. AQ is about 73.5% all-in preflop (including all suitedness combinations). Let's assign a result of 1 to AK winning and 0 to AQ winning. Then our mean is 0.735. Calculating the variance of a single confrontation:

$$\begin{aligned}V &= (0.735)(1 - 0.735)^2 + (0.265)(0 - 0.735)^2 \\ V &= 0.1948 \\ \sigma &= 0.4413\end{aligned}$$

The mean of 2,000 hands is:

$$\begin{aligned}\mu_N &= N\mu \\ \mu_N &= (2000)(0.735) \\ \mu_N &= 1470\end{aligned}$$

For a sample of 2,000 hands, using Equation 2.4, the standard deviation is:

$$\sigma_N = \sigma\sqrt{N}$$

$$\sigma_{2000} = (0.4413)(\sqrt{2000})$$

$$\sigma_{2000} = 19.737$$

The result reported here was approximately 50% of 2000, or 1000, while the sample mean would be about 1470 out of 2000. We can find the z-score using Equation 2.6:

$$z_x = (x - \mu_N) / \sigma$$

$$z_{1000} = (1000 - 1470) / (19.737)$$

$$z_{1000} = -23.815$$

The result reported in this case was 1000 out of 2000, while the expected population mean would be 1470 out of 2000. This result is then about 23.815 standard deviations away from the mean. Values this small are not very easy to evaluate because they are so small - in fact, my spreadsheet program calculates $\Phi(-23.815)$ to equal exactly zero. Suffice it to say that this is a staggeringly low probability.

What's likely is that in fact, this player was either exaggerating, outright lying, or perhaps made the common mistake of forgetting to notice when AK beat AQ, because that is the "expected" result. This is related to a psychological effect known as "perception bias" - we tend to notice things that are out of the ordinary while failing to notice things that are expected. Or perhaps the online site was in fact "rigged." When this example was posed to a mailing list, the reaction of some participants was to point to the high variance of poker and assert that a wide range of outcomes is possible in 2000 hands. However, this confuses the altogether different issue of win/loss outcomes in 2000 poker hands (which does have quite high variance) with the outcome of a single poker hand (which has far less). The variance of a single hand in terms of big bets in most forms of poker is much higher than the variance of the winner of an all-in preflop confrontation. One lesson of this example is not to confuse the high variance of poker hand dollar outcomes with the comparatively low variance of other types of distributions.

When we play poker, many random events happen. We and the other players at the table are dealt random cards taken from a distribution that includes all two-card combinations. There is some betting, and often some opportunity to either change our cards, or to have the value of our hand change by the dealing of additional cards, and so on. Each hand results in some outcome for us, whether it is whining a big pot, stealing the blinds, losing a bet or two, or losing a large pot. This outcome is subject to all the smaller random events that occur within the hand, as well as events that are not so random - perhaps we get a tell on an opponent that enables us to win a pot we would not otherwise have, or save a bet when we are beaten. Nevertheless, the power of the Central Limit Theorem is that outcomes of individual hands function approximately as though they were random variables selected from our "hand outcomes" distribution. And likewise, outcomes of sessions, and weeks, and months, and our whole poker career, behave as though they were appropriately-sized samples taken from this distribution.

The square root relationship of trials to standard deviation makes this particularly useful, because as the number of trials increases, the probability that our results will be far away from our expected value in relative terms decreases.

Assume we have a player whose distribution of hand outcomes at a particular limit has a mean of \$75 per 100 hands, with a variance of \$6,400 per hand. If we sample different numbers of hands from this player's distribution, we can see how the size of the sample impacts the dispersion of the results. We know that the probability that this player's results for a given sample will be between the mean minus two standard deviations and the mean plus two standard deviations is 95.5%. We will identify for each sample size:

- The mean μ_N
- The standard deviation σ
- The two endpoints of the 95.5% probability interval.

Hands	μ_N	σ	Lower endpoint	Higher endpoint
100	\$75	\$800.00	(\$1,525.00)	\$1,675.00
500	\$375.00	\$1,788.85	(\$3,202.71)	\$3,952.71
1,000	\$750.00	\$2,529.82	(\$4,309.64)	\$5,809.64
5,000	\$3,750.00	\$5,656.85	(\$7,563.71)	\$15,063.71
25,000	\$18,750.00	\$12,649.11	(\$6,548.22)	\$44,048.22
50,000	\$37500.00	\$17888.54	\$1,722.91	\$73,277.09
100,000	\$75,000.00	\$25,298.22	\$24,403.56	\$125,596.44
1,000,000	\$750,000.00	\$80,000.00	\$590,000.00	\$910,000.00

As you can see, for smaller numbers of hands, outcomes vary widely between losses and wins. However, as the size of the sample grows, the relative closeness of the sample result becomes larger and larger - although its absolute magnitude continues to grow. Comparing the standard deviation of one million hands to the standard deviation for one hundred hands, the size of a standard deviation is a hundred times as large in absolute terms, but more than a hundred times smaller relative to the number of hands. This is the *law of large numbers* at work; the larger the sample, the closer on a relative basis the outcomes of the sample will be.

Key Concepts

- Variance, the weighted sum of the squares of the distance from the mean of the outcomes of a distribution, is a valuable measurement of dispersion from the mean.
- To compare variance to expected value, we often use its square root, standard deviation.
- Variance is additive across trials. This gives rise to the square root relationship between the standard deviation of a sample of multiple trials and the standard deviation of one trial.
- The Central Limit Theorem tells us that as the size of a sample increases, the distribution of the samples behaves more and more like a normal distribution. This allows us to approximate complex distributions by using the normal distribution and to understand the behavior of multiple-trial samples taken from a known distribution.

Chapter 3

Using All the Information: Estimating Parameters and Bayes' Theorem

In the last chapter, we described some statistical properties of valued probability distributions, as well as the relationship of samples taken from those distributions to the normal distribution. However, throughout the chapter, we simply assumed in examining the sampling distribution that we knew the parameters of the underlying distribution. But in real life, we often don't know these things. Even for simple things such as coin flips, the "true" distribution of outcomes is that the coin is very slightly biased in one direction or the other. A die has tiny imperfections in its surface or composition that make it more likely to land on one side or another. However, these effects are usually small, and so using the "theoretical" coin which is truly 50-50 is a reasonable approximation for our purposes. Likewise, we generally assume for the purposes of poker analysis that the deck is fair and each card is random until it is observed.

We can mitigate real-world difficulties with distributions that we can reasonably approximate (such as coin flips or die rolls). Other types of distributions, however, pose much more difficult problems. In analyzing poker results, we are often interested in a distribution we discussed in the last chapter - per hand won/loss amounts. When playing in a casino, it would be quite difficult and time-consuming to record the results of every hand - not to mention it might draw unwanted attention. The process is somewhat easier online, as downloadable hand histories and software tools can automate the process. But even if we have all this data, it's just a sample. For poker hands, the probabilities of the underlying distribution won't be reflected in the observed data unless you have a really large amount of data.

We can get around this to some extent by using the subject matter of the last chapter. Suppose that we could get the mean and variance of our hand outcome distribution. Then we could find the sampling distribution and predict the aggregate outcomes from playing different numbers of hands. We can't predict the actual outcome of a particular future sample, but we can predict the distribution of outcomes that will occur.

Now the problem is to try to infer the population mean and variance from the sample mean and variance. We will examine two approaches to this process. The first is the approach of classical statistics, and the second is an approach that utilizes the primary topic of this chapter, Bayes' theorem. The first approach takes as one of its assumptions that we have no information other than the sample about the likelihood of any particular win rate. The second approach postulates a distribution of win rates that exists outside of our particular sample that can be used to refine our estimates of mean and variance for the population distribution.

Estimating Parameters: Classical Statistics

Suppose that we have a player who has played a total of 16,900 hands of limit poker. Normalizing his results to big bets (BB) in order to account for different limits he has played, he has achieved a win rate of $\bar{x} = 1.15$ BB/100 hands with a standard deviation of $s = 2.1$ BB/hand. Here instead of μ and σ , which represent population parameters, we use \bar{x} and s , which are sample parameters. Assuming that he plans to continue to play in a similar mix of games with similar lineups of players, what can we say about his "true" win rate μ in the games he has played? We assume in this section that we have no other information about the likelihood of various win rates that might be possible; all win rates from -1 BB/hand to +1BB/hand are deemed to be equally probable.

First of all, it's important to note that we only have a sample to work with. As a result, there will

be uncertainty surrounding the value of his win rate. However, we know that the sampling distribution of 16,900 hands of limit poker will be approximately normal because of the Central Limit Theorem. The observed standard deviation $s = 2.1$ BB/h is a reasonable estimate of the standard deviation of the population, particularly because the sample is relatively large. We can use these facts to obtain the *maximum likelihood estimate* of the population mean.

Consider all possible win rates. For each of these win rates, there will be a corresponding sampling distribution, a normal curve with a mean of μ and a standard deviation σ_N . The peak of each of these normal curves will be at $x = \mu$, and all the other points will be lower. Now suppose that we assume for a moment that the population mean is in fact some particular μ . The height of the curve at $x = \bar{x}$ will be associated with the probability that the observed sample mean would have been the result of the sample. We can find this value for all possible values of μ . Since all these normal curves have the same standard deviation σ_N , they will all be identical, but shifted along the X-axis, as in Figure 3.1.




Figure 3.1, Shifted normal distributions (labeled points at $x=1.15$)

Since the peak of the curve is the highest point, and the observed value \bar{x} is the peak when $\mu = \bar{x}$, this means that $\bar{x} = 1.15$ BB/100h is the maximum likelihood estimate of the mean of the distribution. This may seem intuitive, but we will see when we consider a different approach to this problem that the maximum likelihood estimate does not always have to equal the sample mean, if we incorporate additional information.

Knowing that the single win rate that is most likely given the sample is the sample mean is a useful piece of data, but it doesn't help much with the uncertainty. After all, our hypothetical player might have gotten lucky or unlucky. We can calculate the standard deviation of a sample of 16,900 hands and we can do some what-if analysis about possible win rates.

Suppose we have a sample N that consists of 16,900 hands taken from an underlying distribution with a mean, or win rate, of 1.15 BB/100h and a standard deviation of 2.1 BB/h.

Then, using equation 2.4:

$$\begin{aligned}\sigma_N &= \sigma\sqrt{N} \\ \sigma_{16,900} &= (2.1 \text{ BB/h})(\sqrt{16,900} \text{ hands}) \\ \sigma_{16,900} &= 273 \text{ BB, so} \\ \sigma_N/100h &= 273/169 \approx 1.61\end{aligned}$$

The standard deviation of a sample of this size is greater than the win rate itself. Suppose that we knew that the parameters of the underlying distribution were the same as the observed ones. If we took another sample of 16,900 hands, 32% of the time, the observed outcome of the 16,900 hand sample would be lower than -0.46 BB/100 or higher than 2.76 BB/100.

This is a little troubling. How can we be confident in the idea that the sample represents the true population mean, when even if that were the case, another sample would be outside of even those fairly wide bounds 32% of the time? And what if the true population mean were actually, say, zero? Then 1.15 would fall nicely into the one-sigma Interval. In fact, it seems like we can't tell the difference very clearly based on this sample between a win rate of zero and a win rate of 1.15 BB/100.

What we can do to help to capture this uncertainty is create a *confidence interval*. To create a confidence interval, we must first decide on a level of tolerance. Because we're dealing with statistical processes, we can't simply say that the probability that the population mean has some certain value is zero - we might have gotten extremely lucky or extremely unlucky. However, we can choose what is called a *significance level*. This is a probability value that represents our tolerance for error, Then the confidence interval is the answer to the question, "What are all the population mean values such that the probability of the observed outcome occurring is less than the chosen significance level?"

Suppose that for our observed player, we choose a significance level of 95%. Then we can find a confidence level for our player. If our population mean is μ , then a sample of this size taken from this population will be between $(\mu - 2\sigma)$ and $(\mu + 2\sigma)$ 95% of the time. So we can find all the values of μ such that the observed value $\bar{x} = 1.15$ is between these two boundaries.

As we calculated above, the standard deviation of a sample of 16,900 hands is 1.61 units/100 hands:

$$\begin{aligned}\sigma_N &= \sigma\sqrt{N} \\ \sigma &= (2.1 \text{ BB/h})(\sqrt{16,900}) \\ \sigma &= 273 \text{ BB per 16,900 hands} \\ \sigma/100h &= 273 \text{ BB}/169 = 1.61\end{aligned}$$

So as long as the population mean satisfies the following two equations, it will be within the confidence interval:

$$\begin{aligned}(\mu - 2\sigma) &< 1.15 \\ (\mu + 2\sigma) &> 1.15 \\ (\mu - 2\sigma) &< 1.15 \\ \mu - (2)(1.61) &< 1.15 \\ \mu &< 4.37 \\ (\mu + 2\sigma) &> 1.15 \\ \mu + (2)(1.61) &> 1.15\end{aligned}$$

$$\mu > -2.07$$

So a 95% confidence interval for this player's win rate (based on the 16,900 hand sample he has collected) is [-2.07 BB/100, 4.37 BB/100].

This does not mean that his true rate is 95% likely to lie on this interval. This is a common misunderstanding of the definition of confidence intervals. The confidence interval is all values that, if they were the true rate, then the observed rate would be inside the range of values that would occur 95% of the time. Classical statistics doesn't make probability estimates of parameter values - in fact, the classical view is that the true win rate is either in the interval or it isn't, because it is not the result of a random event. No amount of sampling can make us sure or unsure as to what the parameter value is. Instead, we can only make claims about the likelihood or unlikelihood that we would have observed particular outcomes if a parameter had a particular value.

The maximum likelihood estimate and a confidence interval are useful tools for evaluating what information we can gain from a sample. In this case, even though a rate of 1.15 BB/100 might look reasonable, concluding that this rate is close to the true one is a bit premature. The confidence interval can give us an idea of how wide the range of possible population rates might be. However, if pressed, the best single estimate of the overall win rate is the maximum likelihood estimate, which is the sample mean of 1.15 BB/100 in this case.

To this point, we have assumed that we had no information about the likelihood of different win rates - that is, that our experimental evidence was the only source of data about what win rate a player might have. But in truth, some win rates are likelier than others, even before we take a measurement. Suppose that you played, say, 5,000 hands of casino poker and in those hands you won 500 big bets, a rate of 10 big bets per 100 hands. In this case, the maximum likelihood estimate from the last section would be that your win rate was exactly that - 10 big bets per 100 hands,

But we do have other information. We have a fairly large amount of evidence, both anecdotal and culled from hand history databases and the like that indicates that among players who play a statistically significant number of hands, the highest win rates are near 3-4 BB/100. Even the few outliers who have win rates higher than this do not approach a rate of 10 BB/100. Since this is information that we have before we even start measuring anything, we call it a priori information.

In fact, if we just consider any particular player we measure to be randomly selected from the universe of all poker players, there is a probability distribution associated with his win rate. We don't know the precise shape of this distribution, of course, because we lack observed evidence of the entire universe of poker players. However, if we can make correct assumptions about the underlying *a priori* distribution of win rates, we can make better estimates of the parameters based on the observed evidence by combining the two sets of evidence.

Bayes' theorem

In Chapter 2, we stated the basic principle of probability (equation 1.5).

$$p(A \cap B) = p(A)p(B|A)$$

In this form, this equation allows us to calculate the joint probability of A and B from the probability of A and the conditional probability of B given A. However, in poker, we are often most concerned with calculating the conditional probability of B given that A has already occurred - for example, we know the cards in our own hand (A), and now we want to know how

this information affects the cards in our opponents hand (B). What we are looking for is the conditional probability of B given A .

So we can reorganize equation 1.5 to create a formula for conditional probability.

This is the equation we will refer to as **Bayes' theorem**:

$$p(B|A) = \frac{p(A \cap B)}{p(A)} \quad (3.1)$$

Recall that we defined \bar{B} as the complement of B in Chapter 1; that is:

$$\begin{aligned} p(\bar{B}) &= 1 - p(B) \\ p(B) + p(\bar{B}) &= 1 \end{aligned}$$

We already have the definitions:

$$p(A \cap B) = p(A)p(B|A)$$

Since we know that B and \bar{B} sum to 1, $p(A)$ can be expressed as the probability of A given B when B occurs, plus the probability of A given B when \bar{B} occurs.

So we can restate equation 3.1 as:

$$p(B|A) = \frac{p(A|B)p(B)}{p(A|B)p(B) + p(A|\bar{B})p(\bar{B})} \quad (3.2)$$

In poker, Bayes' theorem allows us to refine our judgments about probabilities based on new information that we observe. In fact, strong players use Bayes' theorem constantly as new information appears to continually refine their probability assessments; the process of Bayesian inference is at the heart of reading hands and exploitive play, as we shall see in Part II.

A classic example of Bayes' theorem comes from the medical profession. Suppose that we have a screening process for a particular condition. If an individual with the condition is screened, the screening process properly identifies the condition 80% of the time. If an individual without the condition is screened, the screening process improperly identifies him as having the condition 10% of the time. 5% of the population (on average) has the condition.

Suppose, then, that a person selected randomly from the population is screened, and the screening returns a positive result. What is the probability that this person has the condition (absent further screening or diagnosis)?

If you answered "about or a little less than 80%," you were wrong, but not alone. Studies of doctors' responses to questions such as these have shown a perhaps frightening lack of understanding of Bayesian concepts.

We can use Bayes' theorem to find the answer to this problem as follows:

A = the screening process returns a positive result.

B = the patient has the condition.

Then we are looking for the probability $p(B|A)$ and the following are true:

$$p(B|A) = 0.8$$

(if the patient has the condition, the screening will be positive 80% of the time)

$$p(A|\overline{B}) = 0.1$$

(if the patient doesn't have the condition, the result will be positive 10% of the time)

$$p(B) = 0.05 \quad (5\% \text{ of all people have the condition})$$

$$p(\overline{B}) = 0.95 \quad (95\% \text{ of all people don't have the condition})$$

And using Equation 3.2, we have:

$$p(B|A) = \frac{p(A|B)p(B)}{p(A|B)p(B) + p(A|\overline{B})p(\overline{B})}$$

$$p(B|A) = \frac{(0.8)(0.05)}{(0.8)(0.05) + (0.1)(0.95)}$$

$$p(B|A) \approx 29.63\%$$

As you can see, the probability that a patient with a positive screening result actually has the condition is much lower than the 80% "accuracy" the test has in identifying the condition in someone who has it. Testing with this screening process would need to be followed up by additional testing to confirm the presence of the condition.

But rather than just plugging into an equation, let's discuss what's happening here by considering a population of 100,000 people, all of whom are tested using this screening process.

Of these 100,000 people:

5,000 actually have the condition. (5% of the population)

95,000 actually do not have the condition. (95% of the population)

Of the 5,000 who have the condition:

4,000 will test positive. (80% of those with the condition)

1,000 will test negative. (20% false negatives)

Of the 95,000 who do not have the condition:

9,500 will test positive. (10% false positives)

85,500 will test negative. (90% of those without the condition)

Now our question was: given that a person has tested positive, how likely are they to have the condition? Out of our 100,000 tested individuals, a total of 13,500 tested positive. Of those, only 4,000 actually have the condition.

$$p(B|A) = 4,000/13,500$$

$$p(B|A) \approx 29.6\%$$

Additionally, we can see that by increasing the accuracy of the test, either by making it more accurate in identifying the condition in those who have it, or in producing fewer false positive identifications of the condition in those who do not, we can increase the conditional probability that someone who tests positive has the condition. Note that this does not increase the chance that somebody actually has the condition - we would not want to increase that! - but rather decreases the number of people who incorrectly test positive, and potentially have to incur the stress of incorrectly believing they have a condition.

Suppose that we increase the effectiveness of the test to always identify those who have the condition, while the false positive rate remains the same. Then, using Bayes' theorem:

$$p(B|A) = \frac{p(A|B)p(B)}{p(A|B)p(B) + p(A|\bar{B})p(\bar{B})}$$

$$p(B|A) = \frac{(1)(0.05)}{(1)(0.05) + (0.1)(0.95)}$$

$$p(B|A) \approx 34.5\%$$

Likewise, if we hold the 80% identification rate constant and drop the rate of false positives from 10% to 6%, we obtain:

$$p(B|A) = \frac{(0.8)(0.05)}{(0.8)(0.05) + (0.06)(0.95)}$$

$$p(B|A) \approx 41.2\%$$

The key to the application of Bayes' theorem is the existence of a prior probability and obtaining new information. In the above example, we began with a prior probability that our patient had the condition of 5%. After the screening, the positive result was new information that allowed us to revise the probability for the patient - in this case upward. This process is called **Bayesian inference** and is critical to successful poker play.

There are countless ways in which Bayesian inference can be used; in fact, many players employ this process unconsciously all the time. Consider a situation in a tournament when a player is at a table that breaks very early on. He is moved to a new table where the player on his right has already amassed a stack seven times as large as the starting stacks, while most of *the* other players in the tournament are still near their starting stacks. Most players would conclude that two scenarios are likely: the player with the huge stack is either an extremely loose and aggressive player or he happens to simply be an ordinary player who got extremely lucky. The natural inclination to believe that he is more likely to be loose and aggressive than very lucky is a direct consequence of Bayes' theorem, whether the player from the broken table is aware of this concept or not.

We can quantify this effect using the formal version of Bayes' theorem, and this can lead us to making stronger plays even without solid information. Consider the following common situation:

A new player sits down in the game. Using all the observational information available to us, which might include stereotyping the player's ethnicity, gender, manner, wardrobe, personal appearance, and so on, we conclude that he is 10% likely to be a "maniac" who will raise 80% of his hands from the cutoff and 90% likely to be a tight player who will raise 10% of his hands from that seat. On the first hand he plays, he raises from the cutoff. (Assume there is no posting.) Now what is the probability- that he is a maniac?

We can use Bayes' theorem to solve this problem, but we invite you to estimate this probability before continuing. We believe that testing intuition is one of the best ways to refine our ability to estimate accurately, a skill that is invaluable not only in poker but in life.

A = The opponent will raise the first hand he plays from the cutoff.

B = The opponent is a maniac.

$p(A B) = 0.8$	(if the player is a maniac, he will raise 80% of the time)
$p(A \bar{B}) = 0.1$	(if the player is not a maniac, he will raise 10% of the time)
$p(B) = 0.1$	(10% of the time, he is a maniac <i>a priori</i>)
$p(\bar{B}) = 0.9$	(90% of the time, he is not a maniac <i>a priori</i>)

Applying Bayes' theorem again:

$$p(B|A) = \frac{p(A|B)p(B)}{p(A|B)p(B) + p(A|\bar{B})p(\bar{B})}$$

$$p(B|A) = \frac{(0.8)(0.1)}{(0.8)(0.1) + (0.1)(0.9)}$$

$$p(B|A) \approx 47.1\%$$

So simply by observing this player raising the first hand, we can adjust the probability of this player being a maniac from just 10% to 47% immediately. If the player raises the first two hands (assuming the same inference for the seat next to the cutoff), this probability moves to nearly 87%! Of course, these probabilities are subject to the accuracy of our original assumptions - in reality, there are not just two types of players, and our probability estimates are probably not so crisp about what type of player he is.

One tendency among players is to delay characterizing and adjusting to a player's play until gaining a little more information, by observing some hands or the like. But this view is overly passive in our view; maximizing EV means taking advantage of all the information we have at our disposal and not necessarily waiting for confirmation that the information is reliable before trying to take advantage of it. The error that these players are making is that they- do not realize the power of the information they have gained. It is worth noting that many players, even players who do not play well often make this adjustment, or a similar one, intuitively. But beware! This adjustment is open to exploitation by players who will sit down in a game and play very differently from their usual style in an attempt to induce significant adjustments by players in the game.

Strong players use Bayes' theorem constantly as new information appears to continually refine their probability assessments; the process of Bayesian inference is at the heart of reading hands and exploitive play, as we shall see in Part II. But even away from the table, Bayes' theorem can allow us to make more informed conclusions about data. To see an example of this, we return to the topic of win rates.

Estimating Parameters: Bayesian Statistics

Recall earlier in this chapter we discussed a player who had played a total of 16,900 hands of limit poker with the following observed statistics:

Win rate of $\bar{x} = 1.15$ BB/100 hands

Standard deviation of $s = 2.1$ BB/hand.

We were concerned with some of the statements that we could make about his "true" win rate based on these observations. Using methods from classical statistics, we found that his maximum likelihood estimator for win rate was 1.15 BB/100 hands and his 95% confidence interval was [-2.07 BB/100, 4.37 BB/100]. These statements relied on an assumption that we had no additional

information about the distribution of this player's win rates.

However, suppose that we guess at a distribution of win rates and try to apply Bayes' theorem to the distribution of win rates and this player's results in order to produce a more accurate estimate. To do this, we must first hypothesize an overall distribution of win rates for this player. Let's assume that he is pulled randomly from the population of regular poker players. What is the shape of this distribution, which is called a *prior probability distribution*?

It's pretty unclear at a glance what this distribution looks like - after all, we do not have access to the results and records of the broad population. But we can simplify and estimate, hoping that our distribution of win rates will be close enough to the truth that we can gain from incorporating it into our estimates. Assuming that our player plays some mixture of lower mid-limit games such as \$10-\$20 to \$30- \$60, we can estimate the total rake paid by the game as about \$3-\$4 per hand, or perhaps 0.1 BB/h. Dividing this amongst all the players approximately equally, we obtain a net rake effect on everyone's win rate of about 0.01 BB/h, or 1 BB/100.

The mean of the distribution of all players' win rates, then, is equal to this value, as this is the net flow of money out of the game. Suppose that we have a roughly normal distribution of win rates, let's just estimate a standard deviation (of win rates) of about 0.015 BB/h. This would lead to a population where 68% of players have a rate between -2.5 BB/100 and +0.5 BB/100 and where 95% of players would have a rate between -4 bb/100 and +2 BB/100. This might square with your intuition - if not, these numbers can be tweaked to reflect different assumptions without changing the underlying process.

To simplify the computation, instead of using the continuous normal distribution, we will create a discrete distribution that roughly mirrors our assumptions. We assume that the underlying distribution of all poker players is as follows:

Win Rate	% of players with this win rate
-5 BB/100	0.25%
-4 BB/100	2%
-3 BB/100	8%
-2 BB/100	20%
-1 BB/100	39.5%
0 BB/100	20%
+1 BB/100	8%
+2 BB/100	2%
+3 BB/100	0.25%

Now that we have an *a priori* distribution of win rates, we can apply Bayes' theorem to this problem. For each win rate, we calculate:

A = the chance of a win rate of 1.15 being observed.

B = the chance that this particular win rate is the true one (*a priori*).

We cannot directly find the probability of a particular win rate being observed (because the normal is a continuous distribution). We will instead substitute the probability of a win rate between 1.14 and 1.16 being observed as a proxy for this value. Recall that the standard

deviation of a sample of this size was 1.61 bets. $p(A|\bar{B})$ is calculated by simply calculating the weighted mean of $p(A|B)$ excluding the current row. From a probability chart:

Win Rate	$p(\bar{B})$	$p(A B)$	$p(A \bar{B})$	$p(B)$
- 5 BB/100	0.25%	0.000004	0.00222555	99.75%
- 4 BB/100	2%	0.000031	0.00226468	98%
- 3 BB/100	8%	0.000182	0.00239720	92%
- 2 BB/100	20%	0.000738	0.00259053	80%
- 1 BB/100	39.5%	0.002037	0.00233943	60.5%
0 BB/100	20%	0.003834	0.00181659	80%
- 1 BB/100	8%	0.004918	0.00198539	92%
- 2 BB/100	2%	0.004301	0.00217753	98%
- 3 BB/100	0.25%	0.002564	0.00221914	99.75%

Applying Bayes' theorem to each of these rows:

$$p(B|A) = \frac{p(A|B)p(B)}{p(A|B)p(B) + p(A|\bar{B})p(\bar{B})}$$

Win Rate	$p(A B)$
- 5 BB/100	0.00%
- 4 BB/100	0.03%
- 3 BB/100	0.66%
- 2 BB/100	6.65%
- 1 BB/100	36.25%
0 BB/100	34.54%
- 1 BB/100	17.72%
- 2 BB/100	3.87%
- 3 BB/100	0.29%
Total	100%

When we looked at the classical approach, we were able to generate a maximum likelihood estimator. In the same way, we can identify -1 BB/100 as the maximum likelihood estimate given these assumptions. Of course, these assumptions aren't really the truth - possible win rates for players are close to continuous. Nevertheless, if we used a continuous distribution and did the more complex math that arises, we would find a distribution similar to this one. And the key implication of this approach is that because of the relative scarcity of winning players, our hypothetical hero is nearly as likely to be a losing player who has gotten lucky as he is to have a positive win rate.

We could see this even more starkly if we considered a player with a much higher observed win rate, perhaps 5 BB/100h. The classical approach would still assign a maximum likelihood estimate of 5 BB/100h, because it considers all win rates to be equally likely (because in the classical method we have no additional information about the distribution of win rates). However, recalculating the above analysis with an observed win rate of 5 BB. we find:

Win Rate	$p(A B)$
- 5 BB/100	0.00%
- 4 BB/100	0.00%
- 3 BB/100	0.00%
- 2 BB/100	0.16%
- 1 BB/100	3.79%
0 BB/100	15.78%
- 1 BB/100	35.40%
- 2 BB/100	33.84%
- 3 BB/100	11.03%
Total	100%

We can see that here, our player is a heavy favorite to be a winning player, and a substantial one at that. However, as win rates of 5 BB/100 hands are absent from the population, Bayes' theorem properly adjusts his win rate to the more likely levels of the top 10% of all players.

Gathering more data should naturally cause these Bayesian estimates to converge on the observed win rates - the more consistently a player has demonstrated his ability to win at a particular rate, the more likely that that is his true rate. And if we recalculate the above considering a sample of 100,000 hands, we obtain:

Win Rate	$p(A B)$
- 5 BB/100	0.00%
- 4 BB/100	0.00%
- 3 BB/100	0.00%
- 2 BB/100	0.00%
- 1 BB/100	1.57%
0 BB/100	33.42%
- 1 BB/100	58.37%
- 2 BB/100	6.60%
- 3 BB/100	0.04%
Total	100%

By this sample size, we are much more confident that the observations match the reality, even though in the underlying distribution, only ten percent of players win at least 1 BB/100.

It is worth noting that there is somewhat of a schism between classical statisticians (sometimes called *frequentists*) and Bayesian statisticians. This disagreement centers (roughly) on the Bayesian idea that population parameters should be thought of as probability distributions. Bayesians estimate a prior distribution, observe a sample, and then adjust their prior distribution in light of the new evidence. Frequentists reject this approach, preferring to characterize population parameters as fixed, unchanging values, even if we cannot know their value. We are strongly oriented toward the Bayesian point of view because of its usefulness in poker analysis.

The method we have used here is not even a full-fledged Bayesian technique, but merely a simplification of techniques of Bayesian analysis intended to make the idea accessible. For more

detail on this and on the divide between practitioners of Bayesian techniques and classical statisticians, we suggest that you consult advanced statistics texts, particularly those that have thorough treatments of point estimation and Bayesian methods.

One topic that emerges from this analysis is the idea of *regression to the mean*. Suppose you take an observation - perhaps of win rates over a sample that is of a small size. The idea here is that observed win rates above the mean will tend to be lower if you repeat the observation in the future, and observed win rates below the mean will tend to perform better in the future. This is not because of any kind of statistical "evening out" where outlying trials from the past are offset by outlying trials (in the other direction) in the future - these events are independent. The principle that is at work here is that if you have outperformed the mean of the entire population, it is somewhat more likely that you have outperformed your expectation as well, while if you have underperformed it, you have likely underperformed your own expectation also. As a result, you will tend (sometimes very slightly) to regress a little toward the mean of all players - that is, the best prediction we can make contains a small amount of the overall mean mixed with observed results. We can see this at work in the Bayesian analysis of our hypothetical player - after 16,900 hands, his adjusted win rate prediction was still very heavily influenced by the population distribution, which dragged his win rate down toward the population mean of -1 BB/100.

Key Concepts

- When estimating population parameters from observed data, we can use one of two methods: the frequentist or classical method, or a Bayesian method.
- The classical method assumes that we have no information about the underlying distribution. This gives rise to a maximum likelihood estimate equal to the sample mean, and a confidence interval that includes all values for which the observed result is inside the significance level.
- Bayes' rule gives us a robust methodology for incorporating new information into prior probability estimates.
- Using a prior probability distribution and applying Bayes' theorem can yield better (Bayesian) estimators of population parameters when we have accurate information about the prior distribution.

Part II: Exploitive Play

*It is not enough to be a good player;
you must also play well.*

Siegbert Tarrasch

Chapter 4

Playing the Odds: Pot Odds and Implied Odds

The heart of poker is decision-making. Players who make better decisions than their opponents win; players who make worse decisions than their opponents lose. In Part I, we defined decisions with higher EV as "better" than their lower EV counterparts. In Parts II and III, we will examine the decision-making process as it applies to playing poker games and identify both techniques of analysis and methods of play that lead to making the best decisions at each juncture.

Part II deals with what we call *exploitive play*; this is play which seeks to maximize EV by taking the action at each decision point which has the highest EV in a particular situation, inclusive of whatever information is available about the opponent, such as his tendencies, tells, and so on. Virtually every player uses exploitive play in one form or another, and many players, even some of the strongest players in the world, view exploitive play as the most evolved form of poker.

Before we get into a discussion of exploitive play, we will introduce some terms and definitions. First, we have the concept of a *game*. In the poker world, we have different definitions of this term, and in Part III, we will be primarily working in the domain of *game theory*, the mathematical study of games. For now, we will define a game as containing the following elements:

- There are two or more players.
- At least one player has a choice of actions.
- The game has a set of outcomes for each player.
- The outcomes depend on the choices of actions by the players.

Normally in our poker discussions, there will be two or more players, and both players will have action choices. The set of outcomes for the game will be expressed in dollars won or lost.

Additionally, we call the "choice of action" a *strategy*. In game theory terms, a strategy is a complete specification of a player's actions choices at all possible paths the hand might follow. In poker, strategies are extremely difficult to specify, as we have what might be called a combinatorial explosion of paths. There are 1326 starting hands a player might hold. Then there are 19,600 different three-card flops that might hit, 47 different turns, and 46 different rivers. Even after factoring in some equivalences with regard to suit, we still have well over five million board/hand combinations to consider. Then we must specify how we will play—each of our hands on each street, how we will respond to raises, checks, bets, and so on.

This is basically impractical for any but the simplest toy games. As a result, we often use the term strategy a little more loosely in poker. Frequently, when we use this term we are referring to our expected play on this and perhaps one more street. The depth to which we specify the strategy is often tied to the convenience with which we can express it; simpler games and more static boards can often go deeper than more complex ones. It is normally just simpler to treat the game this way. We do, however, try to tie together the play on two or more streets as much as possible.

The concepts of Part I have particular meanings when we consider the play in terms of strategies. It is fairly meaningless to consider the expectation of a hand in a vacuum before the play begins, so instead we use the term "expectation of a hand" here to mean the expectation of a hand played with a given strategy against an opposing strategy. Likewise, the expectation of a hand distribution against a strategy is the weighted average expectation of the hands in that

distribution against the opposing strategy, and so on.

Maximizing expectation against the opponent's strategy is the goal of exploitive play. If our opponent plays a strategy *S*, we define the *maximally exploitive strategy* to be the strategy (or one of the strategies) that has the highest expectation against *S*. When playing exploitively, it is often our goal to find the strategy that is maximally exploitive and then employ it. By doing this, we maximize our expectation. We begin with a simple toy game to illustrate the process of finding this strategy.

Example 4.1

Two players play headsup limit poker on the river. Player A has either the nuts (20% of the time) or a valueless (or *dead*) hand (80% of the time), and Player B has some hand of mediocre value - enough to beat dead hands, but which loses to the nuts. The pot is four big bets, and A is first. Let us first consider what will happen if A checks. B could bet, but A knows exactly when he has B beaten or not; hence he will raise B with nut hands and fold at least most of his bluffing hands. B cannot gain value by betting; so he will check. As a result, A will bet all of his nut hands. A might also bet some of his dead hands as a bluff: if B folds, A can gain the whole pot.

We'll call the % of total hands that A bluffs with *x*. A's selection of *x* is his strategy selection. B loses one bet for calling when A has a nut hand, and wins five bets (the four in the pot plus the one A bluffed) when A has a bluff. B's calling strategy only applies when A bets, so the probability values below are conditional on A betting. Using Equation 1.11, the expectation of B's hand if he calls is:

$$\langle B, call \rangle = p(\text{A has nuts})(-1) + p(\text{A has a bluff})(+5)$$

$$\langle B, call \rangle = (0.2)(-1) + (5)x$$

$$\langle B, call \rangle = 5x - 0.2$$

If B folds, his expectation is simply zero.

$$\langle B, fold \rangle = 0$$

We consider a few potential values for *x*:

Situation	<i>x</i> value	$\langle B, call \rangle$	$\langle B, fold \rangle$
A never bluffs	0	-0.2	0
A always bluffs	0.8	+3.8	0
A bluffs 5%	0.05	+.05	0
A bluffs 4%	.04	0	0

B should choose the strategy that has higher expectation in each of these cases. If A bluffs often B should call all the time. If A bluffs rarely, B should never call.

To determine how often A will bluff, B might use his knowledge of A's play, tells, or some other information, perhaps using Bayes' theorem (A might have lost the last hand and therefore has a higher *a priori* chance to be playing this hand overaggressively because of emotion, etc.).

We can also graph these two functions:




Figure 4.1, Game Equity for various B strategies

Figure 4.1 shows the linear functions that represent B's strategy options. We can see that B's maximally exploitive strategy involves choosing the strategy that has the higher value at the x-value that A is playing. So when x is below 0.04, (meaning that A is bluffing 4% of the time), B should simply fold; above that B should call. One thing that is important to note about this is that exploitive play often involves shifting strategies rather drastically. If A changes his x-value from 0.039 to 0.041, that is, bluffing two thousandths of a percent more often, B changes his strategy from folding 100% of the time to calling 100% of the time.

Over the next several chapters, we will look, at some of the principles of exploitive play, including pot odds and implied odds, and then consider the play of some example hands. We will consider situations where the cards are exposed but the play is non-trivial and then play a single hand against a distribution of hands. We focus primarily on the process of trying to find the maximally exploitive strategy in lieu of giving specific play advice on any given hand, and especially on the process of identifying specific weaknesses in the opponent's strategy. This last is particularly valuable as it is generally quite repeatable and easy to do at the table.

Pot Odds

None of the popular forms of poker are *static games* (where the value of hands does not change from street to street). Instead, one common element of all poker games played in casinos is the idea of the *draw*. When beginners are taught poker, they often learn that four cards to a flush and four cards to a straight are *draws* and that hands such as pairs, trips, and flushes and straights are *made hands*. This is a useful simplification, but we use "draw" to mean a variety of types of hands. Most often, we use "draw" to refer to hands whose value if the hand of poker ended immediately is not best among the hands remaining, but if certain cards (often called outs) come, they will improve to be best. However, in some cases, this can be misleading. For example, consider the following two hands on a flop of T♣ 9♣ 2♦ in hold'em:

Hand A: Q♣ J♣

Hand B: A♠ 3♦

Hand A has more than a 70% chance of winning the hand despite his five-card poker hand being

worse at this point. In this hand, it may seem a little strange to refer to the hand that is more than a 7 to 3 favorite as "the draw" while the other hand is the "made hand," because of the connotations we usually associate with these terms. However, we will consistently use the term "draw" throughout this book to mean the hand that at the moment has the worse five-card poker hand and needs to catch one of its outs, no matter how numerous they may be, in order to win. In contrast, we will use the term "favorite" to mean that hand that has more equity in the pot and "underdog" to mean the hand that has less. In the above example, the Q♣ J♣ hand is both the favorite and a draw.

One of the fundamental confrontations between types of hands in poker is between made hands and draws. This confrontation is particularly accentuated in limit poker, where the made hands are unable to make bets large enough to force the draws to fold; instead, they simply extract value, while the draws call because the expectation from their share of the pot is more than the amount they must call. However, all is not lost in big-bet poker for the draws. As we shall see later on, draws are able to take advantage of the structure of no-limit games to create a special type of situation that is extremely profitable by employing a play called the *semi-bluff*.

Example 4.2

The game is \$30-60 holdem. Player A has A♣ A♦. Player B has 9♥ 8♥. The board is K♥ 7♣ 3♠ 2♥. The pot is \$400. Player A is first. How should the action go if both players know the full situation?

You can likely guess that the action goes A bets - B calls. It is valuable, however, to examine the underlying mathematics because it provides an excellent introduction to this type of analysis and to the concept of *pot odds*.

If Player A checks, then Player B will certainly check behind. If B were to bet, he would immediately lose at a minimum $\frac{3}{5}$ of his bet (because he only wins the pot $\frac{1}{5}$ of the time), plus additional EV if Player A were to raise. There will be no betting on the river, since both players will know who has won the hand. Rather, any bet made by the best hand will not be called on the river, so effectively the pot will be awarded to the best hand.

Since 35 of the remaining 44 cards give AA the win, we use Equation 1.11 to determine A's EV from checking to be:

$$\langle A, \text{check} \rangle = p(\text{A wins}) + (\text{pot size})$$

$$\langle A, \text{check} \rangle = \left(\frac{35}{44}\right)(400)$$

$$\langle A, \text{check} \rangle = \$318.18$$

Now let's consider B's options. Again, B will not raise, as he has just a $\frac{1}{5}$ chance of winning the pot (with 9 cards out of 44) and A will never fold. So B must choose between calling and folding.

$$\langle B, \text{call} \rangle = (p(\text{B wins})) (\text{new pot size}) - (\text{cost of a call})$$

$$\langle B, \text{call} \rangle = \left(\frac{9}{44}\right) (\$400 + 60 + 60) - \$60$$

$$\langle B, \text{call} \rangle = \left(\frac{9}{44}\right) (\$520) - \$60$$

$$\langle B, \text{call} \rangle = \$46.36$$

$$\langle B, \text{fold} \rangle = 0$$

Since B's EV from calling is higher than his EV from folding, he will call in response to from A. That makes A's EV from betting:

$$\begin{aligned} \langle A, \text{bet} \rangle &= (p(A \text{ wins})) (\text{new pot size}) - (\text{cost of a bet}) \\ \langle A, \text{bet} \rangle &= \left(\frac{35}{44}\right) (\$520) - \$60 \\ \langle A, \text{bet} \rangle &= \$353.64 \end{aligned}$$

Comparing A's equity from betting to A's equity from checking, we see that A gains a more than \$35 by betting. This result indicates the main principle of made hand vs. confrontations: ***In made hand vs. draw situations, the made hand usually bets.***

There are some exceptions to this rule, as we shall see later when we consider some: complicated draw vs. made hand games. Now let us consider B's play. We found that equity from calling was greater than the zero equity from folding. We can solve an inequity to find all the equities at which B would call, letting x be B's chance of winning the pot:

$$\begin{aligned} (\text{B's chance of winning the pot}) (\text{new pot size}) - (\text{cost of a call}) &> 0 \\ 520x - 60 &> 0 \\ x &> \frac{3}{26} \end{aligned}$$

So B ought to call if his probability of winning the hand is greater than $\frac{3}{26}$, which they are the given case. As we discussed in Part I, we might also say that the odds of B winning be at least 23 to 3. We often call this relationship B's ***pot odds***. B should call if his odd winning the pot are shorter than his pot odds. Pot odds are a shorthand for straightfor equity calculations-the pot lays X to our bet of Y ; our odds of winning are W to Z , then choose a decision based on the relationship between those two.

This leads us to the second principle of hand vs. draw situations.

In made hand vs. draw situations, the draw usually calls if it has positive equity in the after calling and subtracting the amount of the call.

There are a number of exceptions to these rules, and by studying individual hands w examine the situations in which simply relying on pot odds as a guide is incorrect or income. However, pot odds calculations are a useful and practical way of approximating and some calculating outright whether it's correct to call with a draw, particularly at the table.

Pot Odds on Multiple Streets

Poker, of course, is a multi-street game, and as such, we often confront situations where it simply action on the current street that needs to be considered, but action on future streets as Consider another perfect-information situation (both players know their opponents' cards).

Example 4.3 (\$75 Pot)

Is the flop of a \$30-\$60 holdem game. Player A holds $A\heartsuit K\heartsuit$, and player B holds $8\clubsuit$. The flop is $A\clubsuit Ks 4\clubsuit$. There is just \$75 in the pot (suppose that the AK limped preflop) the AK bets \$30. B has just eight outs in this case (as the $K\clubsuit$ makes AK a full house). So an immediate pot odds standpoint, B should call if his chance of winning the pot is go than 30/105, or 28.5%. So what is B's chance of making his flush on the next street? It's on the turn. This chance is only 17%, and so we can see that B's chance of making his is not high enough to justify calling.

Example 4.4 (\$195 Pot)

What if the pot were larger, say \$195? Now B has enough immediate odds to call with his flush

draw on the flop. His EV is then $(\frac{8}{45})(\$255) - \$30 = \$15.33$. On the turn, if the draw misses, there are two cases. In one case, an ace or king will have hit on the turn and he will be unable to win the pot, or be *drawing dead*. This occurs $\frac{4}{45}$ of the time. The remaining $\frac{41}{45}$ of the time he will have $(\frac{8}{44})(\$375) - \$60 = \$8.18$ of positive expectation from calling.

His total expectation on the turn, again using equation 1.11, is:

$$\langle B, \text{turn} \rangle = p(A \text{ or } K) (\langle B, A \text{ or } K \text{ hits} \rangle) + p(\text{other}) (\langle B, \text{other} \rangle)$$

$$\langle B, \text{turn} \rangle = (\frac{4}{45}) (0) + (\frac{41}{45}) (\$8.18)$$

$$\langle B, \text{turn} \rangle = \$7.45$$

So in this case, the play on the flop would go A bets - B calls. Then if the draw did not hit on the turn, A would bet again and B would be able to profitably call again (because calling is worth more than zero from folding).

Example 4.5 (\$135 Pot)

There is also another pot size of interest. If the pot were just \$135, B has enough equity to justify calling on the flop, but not on the turn.

$$\langle B, \text{flop} \rangle = p(\text{win on turn}) (\$135 + \$60) - \$30$$

$$\langle B, \text{flop} \rangle = \frac{8}{45} (\$135 + \$60) - \$30$$

$$\langle B, \text{flop} \rangle = \$4.67$$

$$\langle B, \text{turn} \rangle = p(\text{win on river}) (\$195 + \$120) - \$60$$

$$\langle B, \text{turn} \rangle = (\frac{8}{44}) (\$315) - \$60$$

$$\langle B, \text{turn} \rangle = -\$2.73$$

(Here we ignore the case where A hits an ace or king on the turn and wins the hand immediately. Since B can't call when we ignore this case, his equity is worse when we include those hands, and hence he can't call if we include it.)

This poses a difficulty with our previous characterization of pot odds. We stated earlier, "B should call if his odds of winning the pot are shorter than his pot odds." But here B's chance of winning the pot by the end of the hand is:

$$p(B \text{ wins}) = 1 - p(B \text{ loses})$$

By applying Equation 1.3:

$$p(B \text{ wins}) = 1 - [p(B \text{ loses on turn})][p(B \text{ loses on river})]$$

$$p(B \text{ wins}) = 1 - (\frac{37}{45}) (\frac{36}{44})$$

$$p(B \text{ wins}) = 0.327 \text{ or } 32.7\%$$

Here we use the joint probability $p(A \cap B) = p(A) p(B)$ of B not winning the pot and subtract it from one to find the probability of him winning the pot. We could likewise calculate the chance of him winning on the turn and also on the river, and then subtract the chance that he wins on both streets. However, the single step above is more straightforward for this type of calculation.

If B has nearly a one-third chance of winning the pot then why can we not call in all of these cases, since in each of them B is getting longer odds from the pot than two to one? For example, in 4.3, the smallest pot of the three, B is getting $(\$75 + \$60)$ to $\$60$, or 2.25 to 1. The answer to this question is that only situations where B wins on the next street count for his immediate pot

odds.

Example 4.6 (Why giving yourself odds by raising doesn't work)

Consider the case we looked at previously where the pot is \$135 on the flop. If B calls the flop, there is \$195 in the pot on the turn. If the draw misses, then when A bets \$60, B lacks the odds to call the bet on the turn. But what if B raises the flop instead? Now A will reraise, B will call, and B will have odds to draw on the turn as well. It turns out that this is a losing play for B compared to simply calling the flop and folding the turn. (If B hits one of his outs, he will bet and shut A out of the pot, so the full house redraw doesn't come into play.)

We can break B's expectation into two elements. One is his expectation if he wins on the turn, and the other is the expectation when the hand goes to the river. Let's call B winning on the turn card Bf and B not winning on the turn card Bt .

$$\langle B, \text{raise flop} \rangle = \langle Bf, \text{raise flop} \rangle + \langle Bt, \text{raise flop} \rangle$$

$$\langle Bf, \text{raise flop} \rangle = p(\text{B wins on the turn}) (\text{pot if B raises and A reraises}) - (\text{cost of the raises})$$

$$\langle Bf, \text{raise flop} \rangle = \left(\frac{8}{45}\right) (\$135 + 2(\$30 + \$30 + \$30)) - \$90$$

$$\langle Bf, \text{raise flop} \rangle = -\$34$$

$$\langle Bt, \text{raise flop} \rangle = p(\text{B doesn't win on the turn}) [p(\text{B wins on the river})(\text{pot after the turn bets}) - (\text{cost of the turn bet})]$$

$$\langle Bt, \text{raise flop} \rangle = \left(\frac{37}{45}\right) \left[\left(\frac{8}{44}\right) (\$315 + 2(\$60)) - \$60\right]$$

$$\langle Bt, \text{raise flop} \rangle = \$15.70$$

$$\langle B, \text{raise flop} \rangle = -\$34 + \$15.70$$

$$= -\$18.30$$

$$\langle B, \text{call flop} \rangle = p(Bf) (\$135 + 2(\$30)) - \$30$$

$$= \$4.67$$

So raising the flop, even though it makes it such that B has odds to call the turn and draw to his flush again, has over \$20 less expectation compared to simply calling. In real-life poker, both players don't have all the information. In fact, there will often be betting on future streets after the draw hits his hand. The draw can use this fact to profitably draw with hands that are not getting enough immediate pot odds.

Implied Odds

In the previous section, we discussed the concept of pot odds, which is a reasonable shorthand for doing equity calculations in made hand vs. draw situations. We assumed that both players knew the complete situation. This assumption, however, does not hold in real poker. After all, the purported drawing player could be bluffing outright, could have been drawing to a different draw (such as a straight or two pair draw), or could have had the best hand all along. As a result, the player with the made hand often has to call a bet or two after the draw comes in.

When we looked at pot odds previously, the draw never got any money after hitting his hand, since the formerly made hand simply folded to his bet once the draw came in. However, such a strategy in a game with concealed cards would be highly exploitable. As a result, the draw can anticipate extracting some value when the draw comes in. The combination of immediate odds and expected value from later streets is called *implied odds*.

Example 4.7

Consider the following example:

Imagine the situation is just as in the game previously (Player A holds A♦ K♦ and player B holds 8♣ 7♣. The flop is A♣ K♠ 4♣.), but Player A does not know for sure that Player B is drawing. Instead of trying to arrive at some frequency with which Player A will call Player B's bets once the flush comes in, we will simply assume that A will call one bet from B on all remaining streets.

Earlier we solved for B's immediate pot odds given that the hands were exposed. However, in this game, B's implied odds are much greater. Assume that the pot is \$135 before the flop. Now on the flop, A bets \$30, We can categorize the future action into three cases:

Case 1: The turn card is a flush card.

In this case, B wins \$195 in the pot from the flush, plus \$240 from the turn and river (one bet from each player on each street). Subtracted from this is the \$150 B puts into the pot himself. So B's overall value from this case is \$285. This case occurs $\frac{8}{45}$ (or 17.8%) of the time, for a total EV contribution of \$50.67.

Case 2: The turn card is not a flush card, but the river is.

In this case, B again wins \$285, as one bet goes into the pot on both the turn and river. This joint probability (using Equation 1.3) occurs $(\frac{37}{45})(\frac{8}{44})$ of the time, or about 14.9%. This results in a total EV contribution of \$42.61.

Case 3: Neither card is a flush card.

This occurs the remaining 67.3% of the time. In this case, B calls the flop and turn but not the river, and loses those bets. So he loses \$90. This results in an EV contribution of -\$60.55.

Outcome	$p(\text{Outcome})$	Value	Weighted EV
Turn flush	8/45	+\$285	\$50.67
River flush	$(\frac{37}{45})(\frac{8}{44})$	+\$285	\$42.61
No flush	$1 - [(\frac{8}{45}) + (\frac{37}{45})(\frac{8}{44})]$	-\$90	\$60.55
Total	1		\$32.43

Summing the weighted EV amounts, we find that this game has an expectation of about \$32.73 for B. Contrast this to the expectation of just \$4.67 in the same game if A never paid off B's flushes.

Effective Pot Size

Exploitive play in games such as no-limit holdem relies quite heavily on implied odds. It is frequently correct to take the opportunity to see a cheap flop with hands such as small or medium pairs or suited connectors or weak suited aces in order to attempt to flop a set or a strong flush draw. When doing this, however, it is important to take into account not just the amount that we will win when we do make a powerful hand, but the amount we might lose when that powerful hand loses anyway (as, for example, when the opponent flops a higher set than we). Also, we cannot assume that our opponents will simply pay us off for their entire stack, as some players do when they use the size of the stacks as a guide to what implied odds they have.

In the previous examples, the increased equity from implied odds allowed B to call at lower pot sizes and with less equity because it increases the size of what we might call "the effective pot." In Part III, we will see that "payoff amounts" are a very important part of games played between hand distributions containing various combinations of made hands and draws. In addition, these amounts are of critical importance when considering games where the draw is *closed*. By this, we mean the information of whether the draw has come in is not available to both players. This

is the case, for example, in seven-card stud, where the final card is dealt facedown. When this information is asymmetric, the player with the made hand is at a disadvantage because he must at least sometimes pay off the draw when the draw gets there. The reason for this is because if he does not, then the draw can exploit this by bluffing aggressively.

Bluffing

Bluffing is perhaps the most storied action in poker. One of the most exciting moments in any beginning poker player's career is when he first bluffs an opponent out of a pot. A lot of the drama that is captured in poker on television revolves around moments where a player makes a large bet and the other player must decide whether the player is bluffing or not when deciding to call. It is true in some sense, of course; bluffing is an integral part of poker and is the element of the game that differentiates it directly from games such as chess or backgammon.

A *pure bluff* is a bet with a hand that has no chance of winning the pot if called by the opponent's maximally exploitive strategy. A *semi-bluff* is a bet with a hand that might or — tight not be best at the moment, but which can improve substantially on later streets. Ordinarily, pure bluffs only occur on the last street of a hand, when no further improvement is possible, although it is possible for players to start bluffing with no hope of winning the pot on earlier streets. We often call this *snowing*. An example of an appropriate snow is when holding four deuces in deuce-to-seven lowball. Standing pat before the draw and bluffing is often a quite effective strategy because the opponent cannot hold a very strong hand (86543 being his best possible hand).

On the opposite end of the spectrum are *value bets*. These are bets that expect to have positive expectation even when called. In extreme cases, such as when a player holds the uncounterfeitable nuts on an early street, we can make pure value bets early in a hand. But like straight bluffs, it is frequently the case that with cards to come, most value bets can in fact be semi-bluffs (this depends on how strong the opponent's hand is).

However, between pure bluffs and pure value bets, there is a broad spectrum of bets, most of which are sometimes value bets and sometimes semi-bluffs. The clearest and easiest example of a pure semi-bluff with no value component is a weak flush draw, which might or might not have value if it pairs its cards, but has no chance of being called by a worse hand in terms of high card value. Still, this type of hand can frequently benefit from betting because the opponent may fold. We will return to semi-bluffing later in Part II. The next example, however, shows the exploitive power of bluffing.

Example 4.8

The game is \$40-80 seven-card stud. The hands are exposed, except for the river card, which will be dealt face-down. We call situations such as these *closed draws*: that is, the information about whether the draw has come in is available to only one player.

Player A: 6♥ A♣ A♠ 7♦ 9♦ K♣

Player B: 7♠ 8♠ 9♠ K♠ 2♣ 4♣

As you can see, this game mirrors to some extent the previous game. Player A has a made hand, while Player B has a pure flush draw with no secondary outs. The pot is \$655. The river card is dealt.

In this case, Player A is done betting for this hand. Either B has made his flush or he has not, and B will never call A's bet unless B beats A. Hence, A should simply check and then attempt to play accurately if B bets. If B checks, of course, A should expect to win the pot all the time. In

fact, A should expect to win the pot quite frequently here, far more than 50% of the time. However, he should still not value bet because he only loses expectation by doing so. His expectation in the large pot is still intact whether he bets or not. In this case, also, A's value bet, no matter what card he catches, will never extract value from B's hand, because B can't make a hand better than the open aces A has.

So A checks. Now B must decide what to do. Clearly, he should value bet his flushes. There are 40 cards left in the deck, and eight flush cards. So B will make a flush $\frac{1}{5}$ (or 20%) of the time. In addition to value betting his flushes, B might also want to bluff sometimes when he misses the flush. If he does this and A folds, he'll win a pot with more than eight bets in it with the worst hand.

If B bets, then A has to decide between calling and folding. This decision is actually affected by whether or not A caught a spade on the river, because if he did it is more likely that B is bluffing. But we'll neglect this for now; assume that A decides to play blind on the river, for our convenience.

This leaves us with two unknowns:

A's calling frequency: that is, given that B bets, how often A calls. Call this value x .

B's bluffing frequency: that is, what percentage of his total hands B bluffs with. Call this value y .

Then, from Equation 1.11, we have the following equations:

$$\langle A, \text{call} \rangle = p(\text{B has a flush}) (\text{lose one bet}) + p(\text{B is bluffing}) (\text{pot} + \text{one bet})$$

$$\langle A, \text{call} \rangle = (0.2)(-\$80) + (y)(\$655 + 80)$$

$$\langle A, \text{call} \rangle = \$735y - \$16$$

$$\langle B, \text{bluff} \rangle = p(\text{A calls}) (\text{lose one bet}) + p(\text{A folds}) (\text{pot})$$

$$\langle B, \text{bluff} \rangle = x(-\$80) + (1 - x)(\$655)$$

$$\langle B, \text{bluff} \rangle = \$655 - \$735x$$

So the value of calling for A is dependent on how often B will bluff; the value of calling for B is dependent on how often A will call.

For A, we can solve an inequality to see when calling will have positive value:

$$\$735y - \$16 > 0$$

$$y > \sim 2.2\%$$

This means that if B bluffs more than 2.2% of the time, A should call all the time, because he will have positive expectation by doing so. Also, we can see that at precisely this bluffing frequency, A's calling will have EV 0. This is a critical concept that we will revisit again and again in this book. At this frequency, A is *indifferent* to calling and folding—it doesn't matter what he does.

For B, we can solve an analogous inequality to see when bluffing will have positive value.

$$\$655 - \$735x > 0$$

$$x < \sim 89.1\%$$

This means that if A calls less than 89.1% of the time, then B should bluff all the time, because he will have positive expectation for doing so. We can again see that at precisely this frequency,

B's bluffs will have EV 0. That is, B is indifferent to bluffing or checking his non-flushes.

It might happen that Player A is a "skeptic," who thinks that Player B is a habitual bluffer and will therefore call all the time. If he does this, his expectation from calling is as above, about \$7.35 for every percent of dead hands that B bluffs (above the critical value). B's best response to A's skepticism is to simply stop bluffing. Or perhaps Player A is a "believer," who thinks that Player B wouldn't dare bet without having made a flush. So he folds quite often. Then Player B can simply bluff all his non-flushes, gaining the entire pot each time that Player A folds. Of course, Player B loses out on value when he bets his flushes, but the eight-bet pots he wins more than offset that.

Exploitive Strategies

The responses of Player A and Player B to their opponents' different bluffing and calling frequencies are prototypical exploitive plays. Both players simply maximize their equity against their opponent's strategy. We call these strategies *exploitive* because by performing these equity calculations and employing these strategies, these players identify the weaknesses in their opponent's strategy and exploit them by taking the actions that maximize EV against those weaknesses.

However, equity calculations can be difficult to do at the table, particularly when there are many additional factors, such as those that occur in the real world. In this book, we often present toy games as conceptual aids to the processes we describe. As we examine more complicated situations, where the play is distributions of hands against each other and complex strategies are employed by both sides, sometimes doing straightforward equity calculations is beyond a human player's reach in the time available at the table.

Nevertheless, the identification and exploitation of weaknesses in our opponents' strategies leads us to plays that increase EV against the strategy. And so if our ultimate goal is to maximize the results of equity calculation, identifying and exploiting weaknesses is an effective proxy for expected value calculation and is often more practical at the table.

Returning to Example 4.8, we presented this game as a sort of guessing game between A and B, where each attempts to guess the strategy that the other will employ and respond to that by moving their strategy dramatically in the other direction. However, we can briefly foreshadow things to come in Part III here. The indifference points we identified (~2.2% bluffs for B and 89.1% calls for A) are special values: they are the points at which each player's opponent cannot exploit the strategy any further, even if they know the strategy. If A and B each play these specific strategies, then neither player can improve his EV by unilaterally changing his action frequencies. These strategies are called *optimal*. Finding optimal strategies and playing optimally is the subject of Part III.

Strong exploitive play is essentially a two-part process. The first step is to gather information about the situation. This can include inferring the opponent's hand or hand distribution from the action or identifying situations in which the opponent plays poorly or in a manner that is easily exploitable. The second step is deciding on action in light of the information gathered in the first step. This is often simply taking the exploitive action called for based on the weakness identified in the first step. The second step is often simpler than the first, but this is not always true, as we shall see. We will consider these two steps of the process in turn, with our eye always turned to maximizing EV through the process.

Key Concepts

- Exploitive play is the process of maximizing expectation against the opponent's hands

and strategy. In practice, this often amounts to identifying weaknesses in the opponent's strategy and exploiting them because detailed calculations of expected value are too difficult at the table.

- Pot odds provide a metric for us when considering whether to call with a draw or not; if the size of the pot relative to our chance of winning is large enough, calling is correct; otherwise folding is often indicated.
- Pot odds across multiple streets must be treated together-sometimes if we lack pot odds on a particular street, we can still call because the total cost of calling across all streets is low enough to make the entire play profitable. The focus must be on the total expectation of the strategy rather than the play on any particular street.
- Implied odds involve additional betting that takes place after a draw is complete. Figuring implied odds into calling or folding decisions can help us to maximize EV.
- Two principles of hand vs. draw situations:
 - 1) In made hand vs. draw situations, the made hand usually bets.
 - 2) In made hand vs. draw situations, the draw usually calls if it has positive equity in the pot after calling and subtracting the amount of the call.
- A pure bluff is a bet with a hand that has no chance of winning the pot if called by the opponent's maximally exploitive strategy. A semi-bluff is a bet with a hand that might or might not be best at the moment, but which can improve substantially on later streets. Value bets are bets that expect to have positive expectation even when called.

Chapter 5

Scientific Tarot: Reading Hands and Strategies

Players are often fond, especially on television and in interviews, of the concept of "reading" opponents and the idea that they do this better than their opposition. Implicit in many of the characterizations of the process of "reading" is the idea that a "read" cannot be quantified and that unimaginative "math guys" somehow lack this ability. Unsurprisingly, we disagree. While the ability to accurately pinpoint the hand an opponent holds by "magic" would be in fact powerful, it is our view that reading opponents is largely a process of Bayesian inference -related to betting patterns, with some (normally small) adjustments made for physical tells. This is true whether the process of inference is explicit or subconscious. Using a mathematical model can greatly enhance one's intuition as a guide to proper play.

One popular method of "reading" opponents involves guessing at our opponent's hand, playing as if he held that hand, and hoping to be right. Sometimes this leads to the correct play for the situation. One example of this type of thinking are players who reraise early raisers in limit holdem with small to medium pairs in position, "putting their opponents on AX." It is true that AK. is (assuming the opponent will always raise this hand) the most likely single -and that an early raiser can hold; so frequently this "read" will turn out to be correct. Nonetheless, these players could be sacrificing equity both preflop and postflop by assuming that opponent holds a specific hand, as his distribution contains not only unpaired big cards, but also larger pairs.

One of the authors recently played a hand in a no-limit satellite that was illustrative of this point. There were five players left, the blinds were 150-300, and a player who had pushed his short stack all-in several times in the last ten or so hands raised all-in to about 2100 from the under-the-gun position. It was folded around to the author in the big blind, who had about 3200 in chips and held A♥8♥. The author called. The raiser held AKo, and the author caught an eight and won the hand.

After the satellite was over, the player who held AKo questioned the author, "What did you think I had that you were ahead?" The author responded, "Well, I figured you were jamming there with any ace. any pair, and some good kings. Plus, the stronger you are as a player, the more likely you are raising with hands like T9s and the like. Against that range. I had almost 50% equity, so I think it's an easy call."

The point of the story is that the questioner was trying to imagine the author's thought process as "putting him on a hand," and then evaluating his chances against that particular hand. But no such process occurred. Instead, your author created a distribution of hands with which the player could be raising and acted accordingly. As it happened, the player held one of the stronger hands in his distribution and the author was lucky to win the hand. But the question of "being ahead" of a specific hand was never part of the decision process for the *author*, and the call was still *fundamentally correct*.

Advanced Hand Reading

Other players "read" their opponents by using a combination of logical deduction and tells. These players rule out various options based on assumptions that their opponents play reasonably or that they have an idea of how their opponents play. One common error made by those who practice this type of hand reading is to strip hands from the range of the opponent too aggressively and thus improperly narrow the distribution of hands the opponent might hold.

Nevertheless, what these players do is much more akin to the hand-reading process that we

advocate; in its most sophisticated forms it is very much the same. In theory, the approach is as follows: We generally never assign one single hand to an opponent; instead, each opponent has a probability distribution of different possible hands. At the beginning of a hand, each opponent has a full distribution of random hands, adjusted for the Bayesian card-removal effects of our own hand. Then as each player acts, we adjust for the new information by modifying the various probabilities, both for the new cards revealed, and for our best estimates of what actions the players would take with each hand. We often include an auxiliary probability, which we might call a "lost his mind" probability, which reflects the (sometimes small) probability that our understanding of the player's style is incomplete or that he is simply deviating from his usual strategy. Players, even fairly strong ones, occasionally do things that are quite out of the ordinary, and pretending that these things are probability zero is simply incorrect and can lead to making incorrect "big layoffs" or bad calls and raises.

In practice, of course, we do not hold all the exact probabilities for each hand in our head; this would be overwhelming and much more trouble than it would be worth in terms of making decisions. However, we can usually reconstruct the betting and exposed cards in such a way that we can create this distribution for a player at a point where we need to make an important decision. Additionally, it is not difficult to incorporate physical tells into this framework; we simply apply an additional layer of Bayes' theorem to our existing distribution. That is, we ask "Given the probability distribution of hands that this player currently holds, what is the probability that he would exhibit the tell that I picked up on with each one?" and rebalance the probability distribution in light of this additional information.

The ultimate goal here is to incorporate all the information we can gather to *find the probability distribution of hands that our opponent holds*. This is generally a process of reduction and elimination; hands that the opponent would have played in some clearly different way should be reduced in relative probability within the distribution. Ideally, we could gather enough information to find a distribution that contains only one hand. In practice, though, gathering this much information is normally impossible.

A detailed example may illustrate the principles we're describing. In this example, we will make frequent assumptions about the meaning of play on certain streets. In many cases, these are assumptions that can be questioned and debated. The point of this exercise is not so much to argue about the proper way to play a particular hand but to illustrate the process of hand-reading described above. In this section, we will examine the hand-reading aspects of the hand; in a later chapter, we will consider the key decision that appears as the hand unfolds.

Example 5.1

The game is seven-card Stud eight-or-better. The antes are \$5, and the limits are \$30-60. A player to our right brings it in with the 2♣ for \$10. We complete the bet to \$30 with (5♣ A♥) 4♥. Two tens, the 7♦, and a king fold. The next player raises to \$60 with the 6♣. The bring-in folds, and we call.

(5♣ A♥) 4♥

(??) 6♣

Folded cards: T♣ T♠ 7♦ K♠ 2♣.

On third street, we begin to create our picture of the opponent's hand. Without really knowing anything about the player, we can still begin to construct a rough distribution of hands he might hold. Among the hands that would likely merit a re-raise from almost any player:

(AA) 6
(KK) 6
(66) 6

Three small clubs including the ace.

Our slightly more aggressive opponents' raising distribution would include:

(QQ-TT) 6
Any three small clubs
(A2) 6
(A3) 6

If the opponent were quite aggressive, then we might include:

(99-55) 6
(A4) 6
(A5) 6
(54) 6
(57) 6

Any three clubs with in ace, such as (A♣ J♣) 6♣

And your authors have both seen such re-raises with hands such as:

(QJ) 6

Unfortunately for our hand-reading methodology but fortunately for our bankrolls, the players who make this type of plays often lose their money too quickly to get a truly accurate read on their raising distribution.

At this early stage in the hand, we can only approximate the opponent's distribution. However, it is *worth* noting that certain of these hands are already fairly unlikely. For example, there are only two tens remaining in the deck, so (TT)6 is one-sixth as likely as it might have been had there been no folded cards. Also, a player with (TT)6 would likely see his tens out and fold his hand instead of re-raising.

On fourth street, we catch the 8s and the other player catches the T♦. This is a fortunate turn of events. The opponent now checks and we bet, expecting to occasionally pick up the pot right away. However, the opponent calls.

On this street, we pick up very little information. There are a couple of reasons why we cannot make strong inferences. The first reason is that the betting followed a predictable pattern; we caught a moderately good card, while our opponent caught a fairly bad one. However, the size of the pot and the relative strengths of our distributions (a topic we will return to at length later) make it reasonable for the opponent to flat call pending fifth street. The second is that the betting limits double on the next round. It is often correct for a player to defer raising on a smaller street when doing so will give away information and he will likely get to raise on the next street anyway.

On fifth street we catch the J♥. a rather poor card, although it does give us a three-flush in addition to our four low cards. The opponent catches the K♣ and checks again. We now bet,

expecting to pick up the pot fairly often.

The hands are now:

(5♣A♥) 4♥ 8♠ J♥
 (????) 6♣ T♦ K♣

Now let's consider what the implications of our opponent's various actions might be. Looking back to our candidate hands, especially those from the first two categories: We have probable calling hands:

(AA) 6♣ T♦ K♣
 (X♣Y♣) 6♣ T♦ K♣
 (QQ) 6♣ T♦ K♣
 (JJ) 6♣ T♦ K♣

Probably raising hands:

(66) 6♣ T♦ K♣
 (KK) 6♣ T♦ K♣

Probably folding hands:

(A2) 6♣ T♦ K♣
 (A3) 6♣ T♦ K♣

Note that for each of these categories, we only identify the hands as "probably" indicating some action; it's entirely unclear how our opponents will actually play these hands, or what percentage of the time they will take each action. However, it seems that raising as a semi-bluff with four clubs is the most likely deviation from these strategies. However, other such deviations cannot be ruled out. In creating distributions, it is rather important not to ignore the fact that opponents will play hands in ways we do not anticipate. Ruling out a hand based on prior action can be costly if doing so causes us to make supposedly "safe" raises later in the hand, only to be reraised and shown a better hand.

Now let's say that the opponent calls. We have tentatively narrowed his hand to the big pairs and to four-flushes. On sixth street, we catch the A♠, while the opponent catches the 3♣.

We have a decision to make.

Recapping the hand so far:

(5♣A♠) 4♥ 8♠ J♥ A♠
 (????) 6♣ T♦ K♣ 3♣

There is \$345 in the pot at this point. The question that appears here is: should we bet? Many players would send-automatically bet their aces and a low draw. But how strong is this hand given the assumptions we made above? And also importantly, how close is the decision? If we change the assumptions by a small amount, does it change the entire analysis?

To answer these questions, we must examine the opponent's assumed distribution at this point in the hand and our EV against his entire distribution.

First, let us identify the hands our opponent might hold. For now we will neglect the "lost his mind" category, but once we find an answer we will test it against distributions that include more of those hands.

We classified the opponent's hands into four categories:

(AA)	6♣	T♦	K♣	3♣
(X♣Y♣)	6♣	T♦	K♣	3♣
(QQ)	6♣	T♦	K♣	3♣
(JJ)	6♣	T♦	K♣	3♣

These were based on how he played the hand throughout. Now, however, we must further subdivide the QQ and JJ hands into hands that contain the appropriate club, and hands that do not. (All AA hands contain the ace of clubs, of course—we have the other aces in our hand!). No queens are dead, and of the six possible pairs of queens, three contain the queen of clubs. The jack of hearts is dead (in our hand), so there remain only three pairs of jacks and two of them contain the J♣.

One last step is to figure out how many "three low clubs" hands are possible. We know that the 2♣, 3♣, 5♣, and 6♣ are all outside of the opponent's hole cards, as the 2♣ was folded on third street, and the other cards appear in our hand or on the opponent's board. There are then four low clubs remaining: A♣, 4♣, 7♣, 8♣. This yields six different possibilities containing two clubs.

So we have the following possibilities, considering just the hole cards:

AA	-1
Q♣ Qx	-3
Qx Qy	-3
J♣ Jx	-2
Jx Jy	-1
X♣ Y♣	-6

In addition to this, there might be additional information available. For example, incorporating a tell that our opponent does not like his hand when the ace hits dearly indicates a *bet if we* think that he would exhibit this tell when he has a big pair in the hole but not otherwise. On the contrary, a tell that he is very pleased to see the 3♣ can be extremely valuable, and might save us at least one full bet if he has made a flush.

To recap, we began with a probability distribution that assigned an equal probability to every pair of hole cards that our opponent could hold. When he raised on third street, we were able to narrow that range to a subset of all hands that contained mainly big pairs, strong low hands, and low three-Hushes. On fourth street, we gained equity in the pot by catching a decent card while the opponent caught badly, but gained very little information.

On fifth street, we gained a great deal of information when he continued even after catching a second high card. It's entirely reasonable that an opponent might take another card off on fourth street with three unsuited low cards after reraising the previous street. However, once the K♣ comes off on fifth, we can effectively rule out all low draws that do not contain two clubs in the hole. If the opponent does call on fifth with just three low cards, we profit immediately from his mistake in calling. Even if our exploitive process leads to us sacrificing some expectation on later streets, we have already gained enough by his mistake to more than offset the expectation we lose.

On sixth street, our hand improves substantially because we make a high hand; at the same time

our opponent's distribution improves substantially by catching a small club that is likely to either make him a flush draw or a flush. We narrow his probability distribution down to approximately four classes of hands (aces, big pairs with flush draws, big pairs without flush draws, and made low flushes).

This example primarily made use of the idea of using *betting patterns* to read our opponent's hand. However, there is a second important element to reading hands, which is the ability to both spot and correctly interpret tells. As we stated previously, we will omit discussion on how precisely to identify tells that are being exhibited—this falls outside of the scope of this book. However, in the field of correctly interpreting tells, there is much that can be done mathematically. We will return to this topic later in the chapter. We will return to this hand in Chapter 8, at which point we will consider the bet-or-check decision we have on sixth street.

When we play exploitively, we are *figuring out a distribution of hands and an opponent's strategy* (or a weighted distribution of strategies) against which to maximize our EV. So far in this chapter, we have discussed the process of "reading hands," which for us is effectively a process of Bayesian inference of hand distributions. This process in turn depended on our estimation of how the opponent would play certain hands. In the examples we looked at previously, we either assumed that opponents would play straightforwardly, or we simply assumed that we knew their strategy to facilitate the analysis.

Reading Strategies

An important and extremely difficult part of playing exploitively is accurately assessing the way in which opponents will play the various hands that are present within their distributions. One possible assumption would be that the opponent will play all the remaining hands in his distribution very well. If we make this assumption and find a corresponding exploitive strategy, that strategy will exploit the shape of his distribution of hands. In cases where the opponent has made a serious error in reaching a given point in the hand, exploiting the shape of his distribution will be quite profitable.

We can additionally profit, however, if our opponents continue to make mistakes in how they play their distribution of hands at decision points to come in the hand. We have several sources of information that are useful in estimating the strategy with which our opponents will play their various hands, a process we will call "strategy reading."

- **Direct hand evidence**
Sometimes, we get information about the hand that was played that is known to be truthful because players are forced to show their hands at showdown in order to win the pot. When it is available, this is perhaps the most valuable type of evidence because many of the other forms of evidence can be falsified in some way.
- **Direct hand evidence (opponent-controlled)**
As a special subclass of the above, opponents will sometimes show their hands on purpose. This category is a double-edged sword; we are in fact being shown truthful evidence. However, the fact that this information is being revealed voluntarily by the opponent makes it less valuable because it is likely designed to produce a psychological effect. However, assuming that we are appropriately giving weight to the evidence that we see, we can often use this information to positive effect.
- **Indirect hand evidence**
In addition to the hand value information that must be disclosed at showdown, we can also tabulate frequencies of actions, even for hands that are not shown. This information can be useful in both hand and strategy reading, although it must be combined with additional information about the *types* of hands that are being played.

- ***Player classification correlation***

This is perhaps the most commonly used method of strategy reading both for players about whom we have not gathered much information and for reading strategies generally in situations that come up infrequently. Essentially this consists of a broad classification of players into different categories or category combinations. Some authors have suggested axes such as loose-tight and passive-aggressive. Others have used animal archetypes as the model for their classifications. The core of all these methods, though, is this. If we can accurately characterize the play of a group of players and then accurately assign a particular player to the proper group, then we can make inferences about a player's play even in situations in which we have never seen him play.

The following discussion will consider primarily players who play reasonably; that is, they play in a manner that is not obviously and easily exploitable, even though they make frequent mistakes. At the end of this discussion, we write about some of the common types of very weak players, whose strategies are much easier to read.

Even though we have these sources of information available, it is important to recognize the limitations of our ability to read our opponents' strategy. It is our view that many players, especially those who are skilled at gathering and retaining the above information, overestimate the degree of certainty with which they can read their opponents' strategies, and as a result - overestimate the predictive value of their player classifications. Even over the course of many hands, we gain relatively little direct evidence of how a player plays.

Consider a sample of 1,000 full ring game limit holdem hands (equivalent to about thirty hours of brick and mortar play and perhaps ten to fifteen hours online). This seems like a fairly large sample of hands with which to "get a line" on an opponent's play. But consider that a typical opponent will likely show down only a hundred of those hands. These hands will be split unevenly amongst the nine table positions with the blinds and the button having a disproportionate share of the total hands shown down. Even further, the hands shown down will be played on quite different textures of flops and against different opponents, in different situations.

Indirect hand evidence is not much more help here, as the sample sizes for preflop actions only slowly reach reliable levels. For example, consider 121 hands played under the gun. If a player raises 10% of hands in the observed sample from this position, then over the entire sample, a 95% confidence interval is that he plays roughly between 4% and 16% of hands. 4% of hands is {TT+, AK, AQs}. 16% of hands is {66+, AT+, KJ+, QJ, JT, T9} or some variation. Even seeing a few strong hands that this player raised in this position doesn't allow us to narrow the distribution much, because hands such as QQ are in all the various distributions. In addition, hands like these are more likely to reach the showdown because of their inherent strength. So we have some difficulty in inferring much of value from this indirect evidence as well.

So it seems that the primary source of information about the opponents' play comes from applying the direct evidence not to the problem of directly estimating how he will play in various scenarios, but from applying it to the problem of accurately classifying the player and making conclusions about his play based on his classification. Some of the threats to the validity of this method include:

- Players do not employ a single, fixed strategy in the short run: they in turn might be playing exploitatively in the observed hand. Using an observed hand as a piece of information toward classifying the opponent relies on the observed hand being characteristic of the opponent's overall strategy. What is often important for exploitive play is not how

the player plays in general but how the player plays against you. In a later chapter we will discuss counter-exploitation and how the process of opponent adaptation causes problems in exploitive play and eventually leads us to look for other solutions.

- Players do not employ a single, fixed strategy over a longer period; they read books, talk to other players, and come to realizations on their own. Because of this, even a carefully honed read on an opponent's strategy might become invalid in a fairly short period of time. Even such variables as a player winning for a while at a higher limit than he is used to playing might lead him to play more confidently and drastically change his strategy.
- There is comparatively little opportunity to validate and confirm our hypotheses about an opponent's play. For example, many hands are played in a similar manner by many different strategies. If a player plays one of these hands in the typical manner, it only slightly confirms a read on the player's play. It is a rare circumstance where our read that forecasts a *deviation* from typical play is confirmed. In order for this to occur, our strategy read must predict a deviation from typical play with a particular hand or hands, the opponent must hold one of those particular hands, and we must somehow get direct information about the hand he held (usually through a showdown). Such an occurrence would be a powerful piece of confirming evidence—but this is rare.

The human mind is quite adept at seeing patterns, even in random sequences. Many studies in the field of psychology have confirmed that people often describe underlying patterns in data presented to them, even when the data is randomly generated and no such pattern exists. As a result of this tendency, we should aggressively seek out information to confirm or contradict our hypotheses. But as we saw above, it is often quite difficult to obtain such information because there are not many chances to observe it.

In this vein, we must also be careful not to assign confirming value to evidence that seems to be confirming evidence but actually is not. For example, consider the case of a holdem player who folds each hand for the first few orbits, then plays a hand and at the showdown it turns out to be AA. Many players might take this as confirming evidence that this player plays tightly preflop. However, the fact that the player played AA is not confirming evidence of his tightness—in fact, it is only evidence that he was dealt AA. After all, all reasonable strategies would play AA preflop. The thirty hands he folded before *are* evidence of his preflop tightness; but the fact that the player played AA is not evidence that he is tight. Of course, it makes for a good story..., "This guy folded thirty hands in a row. Then he played a hand and it's aces." But in fact, the fact that he played aces doesn't add anything to the evidence available.

Our discussion here may give the impression that we are quite negative on the process of reading our opponents' strategies. It is not that we deem this process to be without value; instead, we believe that the high requirement for effort expended, and the sometimes dubious value of the information gained are in sharp contrast to the perceptions of many players, many of whom believe that they can quickly and accurately assess a reasonable player's strategy based on just a few hours at the table with them. The process of characterizing a player's play from the very limited data available often produces a vague sense of how a player might play without the attendant details that make exploitation extremely profitable.

Reading Tells

Many players tend to overvalue their ability to read a hand or spot a physical tell. This frequently occurs because we have selective memories about these types of things; we remember times when we had a dead read on someone's hand, but forget times when we slightly surprised at the showdown. The authors try to train our intuition to accurately assess the value of tells observed. One method is called "the hand reading game." At the showdown of a hand in which we are not involved, just after the river betting concludes, we try to name the hand that one of the

players in the hand holds. Hitting the hand precisely is a “win,” while missing the hand is a “loss.” By doing this, we can get a better sense of how often our reads are accurate and how often we're simply wrong. After doing this exercise for a session, many players are surprised at how infrequently they can actually call their opponent's hands and how often they are incorrect.

With regard to tells, the additional information that we can gain from physical tells in terms of reading hands or strategies is of substantial value. However, we have some difficulties here (which are related to the difficulties we have examined so far). Like historical information, it is difficult to observe repeated confirming instances; there is the same parlay of having an appropriate hand, exhibiting the tell, and then providing confirming evidence in the form of a showdown. However, when we have a confirmed tell, this evidence can be extraordinarily valuable. Often, tells can be applied (in a slightly weakened form) by player classification correlation; that is, we observe a particular tell in a broad group of players, and therefore infer that an unknown player exhibiting the tell is in a similar situation to others who exhibit it.

We can attempt to quantify some of these effects as well. Reading physical tells has much in common with diagnosing a disease, which is a fundamentally Bayesian process. We observe certain effects, and from these effects we attempt to infer the cause. A cold causes a stuffy nose and fatigue; amyotrophic lateral sclerosis causes the progressive loss of motor function. We observe these symptoms and try to ascertain their cause. In the same way, there is a cause and effect relationship between hand strength and tells. The player moves his chips in quickly because he's excited about the strength of his hand; the player holds his breath when bluffing, etc.

The skill here is in inverting the Bayesian network and accurately assessing what the *a priori* distribution looks like in order to properly interpret the new information. Cancer patients report fatigue-but a typical physician would not normally diagnose an otherwise healthy patient with cancer based on that symptom. Nor would his mind leap to that on the small list of possible causes; this is because there are many much more probable causes for fatigue.

Suppose that A is an event, such as "has a bluff," and \mathcal{T} is a tell that we have observed. What we are looking for ultimately is $p(A | \mathcal{T})$, the probability that A is true, given that we observe \mathcal{T} . We know from Bayes' theorem (Equation 3.1):

$$p(\mathcal{T}|A) = \frac{p(A \cap \mathcal{T})}{p(A)}$$

and also from Equation 1.5:

$$p(A \cap \mathcal{T}) = p(\mathcal{T})p(A|\mathcal{T})$$

Hence,

$$p(\mathcal{T}|A) = \left(\frac{p(\mathcal{T})p(A|\mathcal{T})}{p(A)} \right)$$

$$p(A|\mathcal{T}) = p(A \cap \mathcal{T})/p(\mathcal{T})$$

We also know that $p(A|\mathcal{T}) = p(A \cap \mathcal{T}) + p(\bar{A} \cap \mathcal{T})$ that is, the probability of \mathcal{T} occurring is the probability of A and \mathcal{T} occurring together plus the probability of A not occurring but \mathcal{T} occurring.

This second term is very important; what it represents is the probability of a *false positive*. This is how often we observe this tell and it doesn't have the meaning we ascribe to it. Tells that are very seldom false positives are very valuable, because the probability of A given T approaches 100%. One example of this type of tell is when a player who has not yet acted fails to protect his cards and, for example, turns away from the table. This type of tell is almost never a false positive, because a player with a strong hand would not exhibit this behavior.

Likewise, this leads to the error of a typical tell-reading practice; that is, ascribing meaning to tells that have alternate meanings that contradict the proposed meanings. Suppose we discuss why a player who pushes in his chips quickly might do that. Some arguments can be made on either side of this discussion-the player might have a weak hand and be attempting to look strong; he might be anxious to put his chips in because his hand is strong but vulnerable, and so on. The point is that even if the observation of the tell is very clear- that is, we can be quite sure we have seen this tell exhibited, the tell is of weaker value because the $p(\bar{A} + T)$ term is so large that our conditional probability isn't close to 1 or 0.

The real problem that we want to solve here is a much more complex one. We want to find the probability of some event A , given a sequence of observations that includes some mixture of tells T_n , some observed hands A_n , some unobserved hands, and some observed tell T for the current hand. But we lack many observations with which to accomplish this. Some players' brains seem to do a better job of addressing these conditional probabilities than others-it is likely this effect that causes us to identify players with stronger "reading" abilities.

On a more positive note for exploitive play, there are certain common types of players who play quite poorly about whom we can gain information easily and immediately. Some examples are:

- **Maniacs** - We use the term maniac to describe players who play extremely loosely and aggressively in all situations, often playing very weak hands very aggressively and putting in extra bets and raises when it is dearly incorrect to do so. It should be noted that players who play very well are often mistaken for maniacs, especially in short-handed or very tight games, because they too play weak hands aggressively and tenaciously. Exploiting these players often requires identifying how they respond to counter-aggression. If they back off when raised, then it is often correct to simply call with fairly strong hands, enticing them to bluff as many times as possible. On the other hand, if they ignore aggression or simply continue to bash away, then pushing strong hands is the proper exploitive response.
- **Rocks** - We use the term rock to describe players who play extremely tightly and often quite passively, entering the pot only with very strong hands and generally failing to extract value from those hands. Rocks rarely bluff, preferring the safety of their strong hands. Exploiting rocks is essentially about stealing their blinds and stealing pots on the flop when both players miss.
- **Calling Stations** - This term denotes a player who plays loosely but passively, often calling bets on many or all streets with weak hands or weak draws. Calling stations share the tenacity of a strong player, but fail to extract value from their hands appropriately. Exploiting calling stations is normally accomplished by value betting additional hands that are substantially weaker than the normal betting ranges, because these bets will now gain value against the expanded calling range of the calling station.

The reason that these strategies are so easy to spot and characterize is that they reflect rather quickly in the primary sources of information we discussed earlier. While players who raise 8% of their hands in a given spot are virtually indistinguishable from players who raise 12% of their hands in that spot after 100 hands, it is easy to identify a player who raises, for example, 50% of

his hands. And in the case of a player who raises just *one* hand out of hundred, gets to showdown with it, and it happens to be a wired pair of aces, we can use Bayes' theorem to make a conclusion about that player's tendencies. Obviously, players play various strategies to degrees; the extreme maniac might simply be a raise-bot, who simply raises as many chips as he can at every opportunity. Such a player is relatively easy to beat. The more that the maniac tempers his aggression in spots where he is clearly the underdog, the closer he moves to playing well, the more difficult he is to exploit, and the less valuable that exploitation is.

Caution should be used, though, when we make a read that a player plays poorly based on one or two examples. There are many dangers inherent in this process. First, the player may have simply made an uncharacteristic bad play. Poker is a complex game, and all players make mistakes of one form or another. Second, the player may have improved from when you first saw the weak play. This occurs frequently, especially in situations such as online where we might make a "note" about a play we saw. Months later, attempting to exploit the player who may have improved a lot in that time is dangerous. Third, the play may not be as weak as you believe—this occurs often when people are faced with loose, aggressive play that is actually quite strong but does not conform to conventional ideas about preflop "tightness." By automatically characterizing loose aggressive players as weak, many players ignore a severe danger that costs them significantly, as well as setting themselves up for frustration at their inability to beat these "weak" players.

The weaker the game, the more that players fall into easily exploitable categories. It is primarily in weak games that we think that the process of reading opponents' strategies pays off the most by quickly identifying and exploiting players who play very poorly. In the case of players who play fairly well, without any very large leaks, we often consider it to be more efficient and profitable to spend energy on identifying leaks in our own play and preventing others from exploiting us.

Key Concepts

- Maximizing EV by playing exploitively requires that we formulate accurate and useful information about the opponents' distributions of hands and their strategies for playing those hands.
- *Some evidence we acquire is direct evidence—a hand revealed at showdown, for example.* This evidence is extremely valuable because it is known to be true.
- Indirect evidence, such as action frequencies, can also be a basis for inferences about play and strategy.
- Reading tells is essentially a Bayesian process—the value of a tell is related directly to both its frequency and reliability. False positives reduce the value of a tell substantially.
- Exploitive play reaches its maximum effectiveness against players who play very poorly.
- Our ultimate goal in hand-reading is to incorporate all the information we can gather to find the probability distribution of hands that our opponent holds.

Chapter 6

The Tells are in the Data: Topics in Online Poker

Online poker has emerged as a major poker venue. Players of all skill and bankroll levels play online, and the increased pace of the game and the ability to play multiple tables makes it particularly attractive to skilled players as a vehicle for profitmaking. In this chapter, we will examine a number of ideas that differentiate online play from brick and mortar play, as well as scratching the surface of the important and difficult problem of reliable data naming.

Online play presents a different set of information to be analyzed.

Clearly, the most obvious difference at a glance between online play and traditional casino (*brick and mortar*) play is the absence or limited presence of physical tells. Instead of the large amount of information that could conceivably be gathered from the mannerisms and actions of a player at the table in brick and mortar, players are generally confronted with some sort of flashing icon or avatar. Pace of play can be an indication, but an uncertain one at best, as players often are distracted from the game by external factors that have no relationship to the contents of their hand. As a result, information based on betting patterns is the primary source for exploitive play online. This makes playing exploitively more difficult, particularly for players who rely on tells in the brick and mortar setting.

Online play also produces an opportunity for information gathering that is unheard of in the brick and mortar world. Almost all of the online sites provide the ability to request via email or otherwise access detailed hand histories of all hands on a site, including (usually) hands mucked at showdown, all actions taken, etc. A player who is willing to request and process a large number of hand histories can rapidly build up a database on his own play as well as on his opponents. A number of commercial programs have been created to manage these hand history databases.

The widespread availability of these programs is a major change in the way that players approach the game. Many of the difficulties with gathering strategy information, particularly small and unreliable sample sizes and perception bias, can be minimized because of the ability to look at real, complete data for a set of hands, rather than relying on memory- and the ability to classify on-the-fly as guides in characterizing opponents. Some commercial programs even provide support in-game, popping up statistics, such as prelop raise percentage, flops seen percentage, and so on, that have been gathered in the database over time on a particular player.

This increased access to data does help to rectify some of the difficulties with gathering data, but at the same time we must not fall into the common error of assuming that since we can acquire a large amount of raw data, that the most difficult part of the analysis has been completed. In fact, this is quite untrue-gathering the raw data is almost always the easiest part of a data analysis project. In analyzing data we have gathered through processes like these, there is typically significant selection bias that makes gathering information about our opponent's strategies difficult, in addition, even what appears to be a relatively large number of hands is often quite insufficient to assess a particular situation- or there could be a significant bias in the hands you are viewing.

For example, suppose that you are trying to identify from empirical data a win rate for a particular opponent, and you have gathered 10,000 hands on that opponent, how accurate will your sample mean be as a predictor? Ignoring for a moment the Bayesian aspects of the problem, suppose this opponent has been a break-even player over the course of the hand sample, with a variance of perhaps $4 \text{ BB}^2/\text{hand}$. Then a 95% confidence interval on his win rate would be -0.02

big bets (BB) to +0.02 BB per hand. But if you are a winning player, then this is likely an underestimate of this player's "true" win rate, since you won't have gathered very much data from hands where you did not play. So his "overall" win rate against the field would generally be higher than his win rate in all the hands where you are a player.

This is a very important idea in mining data; *the only player for whom you have an unbiased record of all play is you*. Some groups (and even some commercial services) have attempted to gather a more global view of data; however, some sites have created policies against such behavior and even so, many hands are omitted. If the hands that are omitted occur randomly, then there shouldn't be a very strong bias, but this is something to consider.

The key to using the hand history data is in finding useful, unbiased statistics that can enable us to make accurate decisions. However, finding these statistics is non-trivial. For example, consider a commonly cited statistic—"voluntarily put \$ in pot %." This is a single percentage figure that indicates how often a player put money in the pot that was not a blind. This may serve as a very rough guide to a player's looseness, but there are many factors that might influence such a number. One of the authors has a database of about 15,000 hands of his own play that has a "VIP" over 50%; this is primarily because this play was largely headsup and three-handed. Playing so loosely in a full game would be highly unprofitable.

But the difficulty of the requirement to make decisions quickly when playing online prevents us from very accurately assessing anomalies like the above. In addition, using a single number there neglects the problem of small sample size—that is, if we have observed 20 hands with a player, and that player has entered the pot 5 times, that is hardly the same thing as a player we have observed 2,000 times with 500 pot entries. Even with larger sample sizes, however, we must take care to understand the effect of position and to understand the nature of the data we have collected—players who frequently play in a mixture of shorthanded and full games create a murkiness about global statistics that is difficult to overcome.

This problem is exacerbated when trying to look at play on a particular street; it is quite difficult to find a methodology that does a good job of matching situations from historical hands to the current situation, giving higher weight to hands that meet the same profile as the current hand (number of players in, pot size, action sequence, etc). Clearly “% of the time player bets on the turn” is a terribly misleading metric because it ignores the context of the hand—the strength of distributions, the previous action, and so on. Such a statistic would only serve well in the most extreme of circumstances, such as when facing a complete maniac, for example. However, if we fit past data to the current situation, we often narrow the sample size too much to obtain information of any use.

This is not to say that it is impossible to gain insight from studying hand histories—instead, it to say that it is decidedly non-trivial, and we have run into many players who think that it should be easy once the hand histories have been gathered to mine out meaningful information about their opponent's play. We think that this overlooks the difficulties of data mining, and that most players gain little additional advantage from trying to use these techniques.

The major gain from hand history collection, in our view, is in analyzing our own play. Here, we lack the biases that are present in our analysis of other players' play as well as having the largest sample sizes and the most complete set of hands. In fact, we highly recommend to you (if you haven't already) to either obtain a commercial product or perhaps write your own program to gather and parse hand histories. This will enable you to analyze your own play in-depth, including the ability to subdivide play by hand, by position, by limit, and so on.

In doing this, it is generally possible to eventually drill any subdivision down to the hands that make it up; we find that anecdotally, occasionally looking at the way hi which we play our hands empirically on a one-by-one basis is helpful in understanding leaks or patterns of play. Additionally, by reviewing your play, you will likely find many mistakes that you didn't know you made; that is, places in which you played hands in ways that you wouldn't have thought you would. The gap between how we intend to play and how we play in practice is often much larger than we would believe.

One of the best things (for the winning player) about online play is the pace of the game. Online play is fast!

Gathering statistically significant data on your own play might take place in a matter of months rather than years online.

Suppose that we want to know our win rate to within 0.01 BB/h with 95% confidence when our variance is $4 \text{ BB}^2/\text{h}^2$. Using equation 2.2, $\sigma = \sqrt{V}$, we know that $\sigma = 2 \text{ BB}/\text{hand}$. Then we apply Equation 2.4, $\sigma_N = \sigma/\sqrt{N}$, and find that our standard deviation for N hands will be $2/\sqrt{N}$. Thus, $2/\sqrt{N} = 0.005$, or $N = 160,000$ hands to gather statistically significant data.

Before online play, the Limit of hands that could be reasonably played on a sustained basis was something around 35 hands per hour. So suppose that a workaday professional played 2,000 hours a year at 35 hands per hour. It would take him a little over two years to play this many hands.

Online players can, with a small amount of practice, play as many as four (some report playing up to eight or twelve games at once). Supposing that hill tables get about 90 hands per hour, an online player can play approximately 350 hands per hour on average. Working the same 2,000 hour year, the online player can play this many hands in just three months. The most obvious effect of this is that if a player's win rate is the same per hand from brick and mortar play to online play, he can win as much as ten times as much money online. However, this is not necessarily possible.

Online players often play better at comparable limits than their brick and mortar counterparts.

There are a number of factors that contribute to this:

- In order to sustain a poker economy that contains some winners and a rake, there must be a substantial number of net losers. Often, net losers have some breaking point, where they are no longer willing to continue to lose. In the brick and mortar setting, some net losers can play for quite a long time before reaching this breaking point. Online, due to many more hands being played, these players reach their thresholds and often quit much faster.
- There are barriers, both psychological and practical, to being large losers at high limits online. It is typically necessary to get large sums of money onto the site through electronic means. Additionally, it's harder to sustain the illusion that one is "doing all right" by not keeping records online than hi brick and mortar; the requirement to keep transferring more money is always present.
- These two factors tend to keep the number of "whales" online smaller than it is in a comparable brick and mortar setting. Additionally, these effects also tend to weed out the weaker players who are modest losers. Players who are small net losers can often play with a little luck for a long time in brick and mortar. Online these players are bankrupted or pushed back down to lower limits much faster because they reach the long run more quickly.

- Players who are interested in profit at poker often prefer the online setting because they can play so many more hands. Only the biggest games in brick and mortar offer a comparable edge and often have much larger variance. Making money at poker is essentially a volume business, and players who are serious about making money can get the most volume online.

As a result of all these factors, it is typical that an online game for a particular stake will be tougher than the comparable game in a brick and mortar casino. In summary, online poker provides a different type of challenge or opportunity for players; different skills are emphasized in the online setting compared to the brick and mortar setting.

Key Concepts

- Online poker presents a different set of information to incorporate into our hand and strategy reading process. Instead of physical tells, we are able to more accurately observe what happened at the table by analyzing hand histories.
- It is sometimes possible to use hand history data to characterize our opponents' play—however, small sample size problems can frustrate this effort. We can, however, use hand history data to analyze our own play very effectively.
- Play online is generally tougher than corresponding play in a traditional casino because of the speed of play and psychological barriers that may not exist offline.
- The only player for whom you have an unbiased record of all play is you.
- Gathering statistically significant data on your own play might take place in a matter of months rather than years online.
- Online players often play better at comparable limits than their brick and mortar counterparts.

Chapter 7

Playing Accurately, Part I: Cards Exposed Situations

Exploitive play is a two-step process. The first step is the processing of reading our opponents' hands and strategies, which we discussed in Chapter 5. The second step is deciding on an action in light of the information from the first step. The process here is simply to calculate the EV of the various action options and select the one that has the highest expectation.

In calculating the EV against our opponent's distribution, we will consider each hand he can hold in turn and weight the EV values against his distribution. Since in each case we know his cards, and we are playing just a single hand against his known hand, we may be able to gain some insight by having the players simply turn their cards face up and examine the play from there. This may seem to be a frankly uninteresting game; in many cases it would be, with made hands betting while draws chase or fold depending on pot odds, as we saw in Chapter 4. But some interesting situations can arise where the correct play is neither obvious nor intuitive.

Example 7.1

Consider the following example:

The game is \$30-\$60 7-card stud.

Player X has: A♥ K♦ Q♥ J♠ T♣

Player Y has: A♠ T♠ 8♠ 5♠ 2♣

The pot contains \$400. Player Y has only \$120 in chips left; player X has him covered, How should the play go?

In considering this, we can look at each player's situation separately. For player X, the initial decision is between checking and betting. However, we prefer to characterize strategic options in a more complex manner that reflects the thought processes behind the selection. So instead of two options-checking and betting-X actually has five options:

Check with the intention of folding to a bet.	(check-fold)
Check with the intention of calling a bet.	(check-call)
Check with the intention of raising a bet.	(check-raise)
Bet with the Intention of calling a raise.	(bet-call)
Bet with the intention of folding to a raise.	(bet-fold)

X can select from among these options the one with the highest value. We present X's options in this manner because the consideration of future actions on the same street is an important part of playing well. Many errors can be prevented or mitigated by considering the action that will follow when taking an initial action.

We can immediately calculate the expected value of the fourth and fifth options. The expected value of betting if Y folds or calls is the same no matter what X's intentions were if Y raises, so we can ignore those cases. If Y raises, there is no more betting on future streets, and the expected value of the bet-call option is:

$$\langle X, \text{bet-call} \rangle = (p(X \text{ wins})) (\text{new pot size}) - (\text{cost of action on this street})$$

The pot is currently \$400; if both players put their remaining \$120 in on this street, the new pot will be \$640. Y will make a flush on sixth street $\frac{8}{42}$ of the time, and when he misses, he will

make a flush on the river $\frac{8}{41}$ of the time. X cannot improve, so X's chance of winning the pot is:

$$p(X \text{ wins}) = 1 - p(X \text{ loses})$$

$$p(X \text{ wins}) = 1 - \{p(Y \text{ wins } 6^{\text{th}} \text{ street}) + [p(Y \text{ misses } 6^{\text{th}} \text{ street})][p(Y \text{ wins river})]\}$$

$$p(X \text{ wins}) = 1 - [(\frac{8}{42}) + (\frac{34}{42})(\frac{8}{41})]$$

$$p(X \text{ wins}) = 65.16\%$$

So the expected value of X's bet-call option is:

$$\langle X, \text{bet-call} \rangle = p(X \text{ wins}) (\text{new pot size}) - (\text{cost of action})$$

$$\langle X, \text{bet-call} \rangle = (0.6516)(\$640) - \$120$$

$$\langle X, \text{bet-call} \rangle \approx \$297$$

If X bets and folds to a raise, then his expected value is:

$$\langle X, \text{bet-fold} \rangle = - \$60$$

These are not the total EVs of adopting these strategies-however, the EVs of the cases where Y doesn't raise are identical, hence we can compare these directly. It's clear from this that X's strategy of bet-call has higher expectation than his strategy of bet-fold. We can therefore eliminate bet-fold from consideration as a possible strategy for X. This should be intuitive-after all, X has almost two-thirds of the pot! Folding in a fairly large pot when his opponent is merely drawing to a flush would be disastrous.

In a like manner, we can examine the three options that begin with a check. If Y checks behind, then these three options have identical EV. If Y bets, then we have the following:

$$\langle X, \text{check-fold} \rangle = \$0$$

If X check-raises and Y calls, then X's EV is the same as in the bet-call case where Y raised. And if X check-raises and Y folds, X's EV is even higher:

$$\langle X, \text{check-raise} \mid Y \text{ calls} \rangle \approx \$297$$

$$\langle X, \text{check-raise} \mid Y \text{ folds} \rangle = \$460$$

From this, we can eliminate the check-fold option. X can make more money by substituting the check-raise option for the check-fold option anywhere that he would play the latter.

With simple analysis, we have eliminated three of X's possible strategies, leaving us with:

check-call
check-raise
bet-call

To go further, however, we must consider what Y will do if X checks or bets. Suppose that we take as our guess at Y's strategy that he will check behind if X checks, and call a bet if X bets. On sixth street, what will happen? If Y has made a flush, then all betting will end-X will no longer call any bets because he has no chance of winning the pot. If Y has not made a flush, the hand has simplified to the case we looked at in Chapter 4. X will bet and Y will call or fold based on pot odds-in this case, Y will call a bet.

Using this information, we can calculate the expected values of the three options for X.

$$\begin{aligned} \langle X, \text{check-call} \rangle &= p(\text{flush on } 6^{\text{th}})(\$0) + p(\text{flush on } 7^{\text{th}})(-\$60) + p(\text{no flush})(\$460) \\ \langle X, \text{check-call} \rangle &= \left(\frac{8}{42}\right)(\$0) + \left(\frac{8}{41}\right)\left(\frac{34}{42}\right)(-\$60) + \{1 - [\frac{8}{42} + \left(\frac{8}{41}\right)\left(\frac{34}{42}\right)]\}(\$460) \\ \langle X, \text{check-call} \rangle &= \left(\frac{8}{41}\right)\left(\frac{34}{42}\right)(-\$60) + \left(1 - \frac{8}{42} - \left(\frac{8}{41}\right)\left(\frac{34}{42}\right)\right)(\$460) \\ \langle X, \text{check-call} \rangle &\approx \$290.24 \\ \langle X, \text{check-raise} \rangle &\approx \$290.24 \end{aligned}$$

The equity of check-raising and check-calling are the same because if X does check, Y will check behind, such that X doesn't get to Follow through with his call or raise intentions.

$$\begin{aligned} \langle X, \text{bet-call} \rangle &= p(\text{flush on } 6^{\text{th}})(-\$60) + p(\text{flush on } 7^{\text{th}})(-\$120) + p(\text{no flush})(\$520) \\ \langle X, \text{bet-call} \rangle &= \left(\frac{8}{42}\right)(\$60) + \left(\frac{8}{41}\right)\left(\frac{34}{42}\right)(-\$120) + \{1 - [\frac{8}{42} + \left(\frac{8}{41}\right)\left(\frac{34}{42}\right)]\}(\$520) \\ \langle X, \text{bet-call} \rangle &= \left(\frac{8}{42}\right)(\$60) + \left(\frac{8}{41}\right)\left(\frac{34}{42}\right)(-\$120) + \{1 - [\frac{8}{42} - \left(\frac{8}{41}\right)\left(\frac{34}{42}\right)]\}(\$520) \\ \langle X, \text{bet-call} \rangle &\approx \$308.43 \end{aligned}$$

Based on this, then, X should bet.

We have two final confirmations to make. Recall that we guessed at Y's strategy of checking behind X's check and flat calling X's bet. If Y bets behind X's check, X can simply flat call Y's bet and get \$308.43. So Y does better by checking behind. We can calculate the EV of Y raising X's bet as well. Recall that we previously found that X's EV from the bet-call option when Y raised was \$297. This is lower than \$308.43; hence Y should instead raise. X can still do no better than his bet-call.

So the play on fifth street goes (if both players maximize their EV):

X bets, and Y raises, and X calls.

This may look a little odd, as the draw raises the made hand's bet. X (the made hand) bets and charges Y to draw-however, it is in turn correct for Y to raise, while a 2-1 dog to make his hand! This occurs because the stacks are limited -Y is all-in and can't be punished by more betting. Since he has odds to call on 5th street for the 6th street card alone, he might as well make sure that all the money goes in on 5th street, so that if he should make his flush on 6th street, he'll still get the last bet hum the straight. This is because his draw is *open* - that is, X knows when he is beaten. In cases where the draw is *closed*, X may still have to pay off on sixth if Y makes his flush.

Example 7.2

Next we have a different example with limited stacks. In this case we allow the starting stack sizes to vary and examine the effect of the initial stack sizes on the play.

The game is pot limit holdem, but with a special rule that only pot-sized bets (or allin bets if the player has less than a pot remaining) may be made. We call this the *rigid pot limit game*; it is substantially simpler than the full pot limit game (where players may bet any amount between a minimum bet and the pot).

Player X has: A♥ A♦

Player Y has: 8♣ 7♣

The flop is: 9♣ 6♣ 2♦

(We'll ignore runner-runner full houses for the AA and runner-runner two pair for the 87s for the sake of discussion. Assume that on each street the 87 simply has 15 outs and either hits or docs

not.)

We can immediately calculate Y's equity if the cards are simple dealt out:

$$\begin{aligned} \langle Y \rangle &= 1 - p(\text{miss twice}) \\ \langle Y \rangle &= 1 - \left(\frac{30}{45}\right)\left(\frac{29}{44}\right) \\ \langle Y \rangle &= 56.06\% \end{aligned}$$

The pot contains \$100, Player X is first to act. How should the play go for different stack sizes?

Case 1: Small stacks.

Let's first assume the stacks are \$50. In this situation, player Y is the favorite if we simply dealt the cards out—he has 56.06% equity with his straight-flush draw. If Player X bets, then Yearly Player Y will call, yielding an EV of:

$$\begin{aligned} \langle X, X \text{ bets}-Y \text{ calls} \rangle &= p(X \text{ wins}) (\text{new pot value}) - (\text{cost of bet}) \\ \langle X, X \text{ bets}-Y \text{ calls} \rangle &= [1 - p(Y \text{ wins})](\text{new pot value}) - (\text{cost of bet}) \\ \langle X, X \text{ bets}-Y \text{ calls} \rangle &= [1 - 0.5606](\$200) - \$50 \\ \langle X, X \text{ bets}-Y \text{ calls} \rangle &= (0.4394) (\$200) - \$50 \\ \langle X, X \text{ bets}-Y \text{ calls} \rangle &= \$37.88 \end{aligned}$$

This is the equity of all the money getting in on the flop, no matter who bets first. It should be clear that if Player X checks, then Player Y can guarantee that X's equity is no greater than this number by simply betting. X will have odds to call, and this same equity will be achieved.

If the play went check-check on the flop and Y failed to make a straight or flush ($\frac{30}{45}$ of the time) X could then bet the turn. In that case, X would have $\frac{29}{44}$ chance of winning. Y still has a clear call, getting more than 3 to 1 from the pot. X's equity, then is:

$$\begin{aligned} \langle X, \text{bet turn} \rangle &= [p(Y \text{ misses flop})][p(X \text{ wins})(\text{new pot value}) - (\text{cost of bet})] \\ \langle X, \text{bet turn} \rangle &= \left(\frac{30}{45}\right)\left[\left(\frac{29}{44}\right)(200) - (50)\right] \\ \langle X, \text{bet turn} \rangle &= \$54.55 \end{aligned}$$

This expected value is higher for X than the expected value when both players were all-in on the flop. And if it went check-check on the flop and Y did make his hand on the turn, X would simply fold. So X would prefer to have the action go check-check on the flop.

However, Y knows this as well. Since Y can limit X's equity to \$37.88 by betting, the action should go check-bet-call on the flop. Notice that Player X's equity in the pot based on the showdown value of his hand is actually $(\$100) (1 - 0.5606) = \43.94 , so the post-flop betting here reduces his equity by more than \$6.

Case 2: Medium Stacks.

Now let's assume the stacks are \$400. Again, player Y has a small edge if the money goes all-in on the flop with 56.06% equity.

If X bets, then Y has three choices—folding, raising all-in to \$400 or calling \$100. If he folds, then X's equity is \$100.

$$\langle X, \text{bet}; Y, \text{fold} \rangle = \$100$$

If he raises all-in, then X's equity is:

$$\begin{aligned} \langle X, \text{bet; Y raise} \rangle &= p(X \text{ wins}) (\text{new pot value}) - (\text{cost of bet}) \\ \langle X, \text{bet; Y raise} \rangle &= [1 - p(Y \text{ wins})](\text{new pot value}) - (\text{cost of bet}) \\ \langle X, \text{bet; Y raise} \rangle &= (1 - 0.5606)(\$900) - (\$400) \\ \langle X, \text{bet; Y raise} \rangle &= (0.4394)(\$900) - \$400 = -\$4.55 \\ \langle X, \text{bet; Y raise} \rangle &= -\$4.55 \end{aligned}$$

If Y calls, then there are two things can happen:

$$\begin{aligned} &15/45 \text{ of the time Y hits and X loses } \$100 \text{ net.} \\ \langle X, \text{bet; Y call} \mid Y \text{ hits} \rangle &= \$ -100 \end{aligned}$$

$$\begin{aligned} &30/45 \text{ of the time Y misses and the game simplifies to a pot odds game:} \\ \langle X, \text{bet; Y call} \mid Y \text{ misses} \rangle &= p(X \text{ wins}) (\text{new pot value}) - (\text{cost of bet}) \\ \langle X, \text{bet; Y call} \mid Y \text{ misses} \rangle &= (29/44)(\$900) - \$400 \\ \langle X, \text{bet; Y call} \mid Y \text{ misses} \rangle &= \$193.18 \end{aligned}$$

At this point, X will bet \$300 and Y will be forced to call with $15/44$ equity in the pot. X's EV in this case will be

$$\begin{aligned} \langle X \rangle &= [p(X \text{ wins})](X\text{'s net when winning}) - [p(Y \text{ wins})](X\text{'s net when losing}) \\ \langle X \rangle &= (29/44)(\$500) - (15/44)(\$400) \\ \langle X \rangle &= \$193.18 \end{aligned}$$

so X's overall equity will be:

$$\begin{aligned} \langle X, \text{bet; Y call} \rangle &= (15/45)(-\$100) + 30/45 (\$193.18) \\ \langle X, \text{bet; Y call} \rangle &= \$95.45 \end{aligned}$$

We calculated these EVs in X's terms-since this is a two player game and no one else can claim any part of the pot, Y can seek to either maximize his own expected value or minimize X's, and these will have the same effect. So it's clear that if X does bet \$100, Y should raise all-in as this is the best of these three options.

Alternatively, X could check. If X checks, then Y can check or bet.

If Y checks, then again we have two outcomes. $15/45$ of the time Y hits and X nets \$0. $30/45$ of the time, Y misses, and X bets the pot and is called.

X's overall EV, then is:

$$\begin{aligned} \langle X \text{ check; Y check} \rangle &= [p(Y \text{ wins}) (\text{pot value to X})] + [p(X \text{ wins}) (\text{pot value to X} - \text{cost of bet})] \\ \langle X, \text{check; Y check} \rangle &= (15/45)(\$0) + (30/45)[(29/44)(\$300 - \$100)] \\ \langle X, \text{check; Y check} \rangle &= \$65.15 \end{aligned}$$

Alternatively, Y can bet. If X raises all-in, he has -\$4.55 again, while if he calls, he again makes \$95.45.

Summarizing these outcomes (EVs from X's perspective):

X's action	Y's action		Call	Raise	Fold
	Check	Bet			
Bet			\$95.45	-\$4.55	\$100
Check	\$65.15				
Check-raise	\$65.15	-\$4.55			
Check-call	\$65.15	\$95.45			

From this table, we can see that if X checks, the worst he can do is +\$65.15 (Y checks behind.) So at this stack size (\$400), the action should actually go check-check on the flop, Neither player can bet; if X bets, Y simply raises all-in and takes his edge on all the money, while if Y bets. X calls and gets the majority of the money in on the turn when he has a significant edge, Also notice that as the stacks have grown, X has gained because of the betting. In the small snack case, X actually lost money as a result of the betting (versus his showdown equity in the initial pot.) But here, X gains \$65.15 - \$43.94, or \$21.19 from the postflop betting.

Case 3: Deep Stacks

The last rigid pot-limit case we'll consider is stacks of \$1300, or three pot-sized raises. This time, let's subdivide the EV calculations into four subcases:

Subcase a). *No money goes in on the flop.*

Since neither player will raise on the turn, X's equity in this subcase is the same as the receding case when no money went in on the flop, or \$65.15

Subcase b). *One pot-sized bet (\$100) goes in on the flop.*

This subcase, too, is similar to the preceding with X's equity at \$95.45.

Subcase c). *Two pot-sized bets (\$400) go in on the flop.*

In this subcase, X's equity is -\$400 in the case where Y hits on the turn. When Y misses, X bets \$900 and Y calls. $\frac{29}{44}$ of the time, X wins the pot and the rest of the time Y does.

$$\langle X, 2betsonflop \rangle = \left(\frac{15}{45}\right)(-\$400) + \left(\frac{30}{45}\right) [(\$2700)\left(\frac{29}{44}\right) - \$1300]$$

$$\langle X, 2betsonflop \rangle = \$346.21$$

Subcase d). *Three pot-sized bets (\$1300) go in on the flop.*

In this subcase, A simply gets his showdown equity from the pot:

$$\langle X, 3 \text{ bets on flop} \rangle = (1 - 0.5606) (\$2700) - \$1300$$

$$\langle X, 3 \text{ bets on flop} \rangle = (0.4394) (\$2700) - \$1300$$

$$\langle X, 3 \text{ bets on flop} \rangle = -\$113.62$$

We can formulate some logical rules for the players' play based on these equities:

- 1) **X will never put in the second bet on the flop.**
If he does, Y will raise all-in. obtaining subcase d), the worst outcome for X.
- 2) **Y will never put in the second bet on the flop.**
If he does, X will flat call and we have subcase c), the worst outcome for Y.
- 3) **X prefers subcase b) to subcase a) –**
If he has a choice between zero or one bets, he will choose one.

4) *Y prefers subcase a) to subcase b) –*

If he has a choice between one and zero bets, he will choose zero.

From these rules, we can obtain the strategies for X and Y. Neither player will put in the second bet. Hence, both players have the option to put in the first bet if they wish. Since X prefers this option, he will bet, and Y will call. This results in an EV for X of \$95.45. It is worth noting that X's expected value from betting if Y folded is just \$100. In this case X gains a lot of value from the postflop betting-in fact the value of his hand has nearly doubled from \$43.94 to \$95.45.

When a made hand competes with a good draw, the draw generally does best by getting all the money into the pot as early in the hand as possible, while there are many cards to come. The made hand, by contrast, only wants to put in enough action such that he still has bets with which to punish the draw when the draw misses on the next street. In the case where there were only two bets left, if the made hand bet, the draw could get all-in. Instead, it's preferable for the made hand to delay betting so that he can extract value on the turn when the draw misses. But when three bets are left, the made hand can afford to bet the flop, knowing that the draw cannot prevent him from making a pot-sized bet on the turn.

Change the hands slightly by reducing player Y's outs by one, to 14, and now if there are three pot-sized bets left, player Y has to fold despite being the favorite to win the hand. Not only is there a pot-sized bet on the turn ($\frac{14}{44}$ being less than $\frac{15}{45}$), but also a pot-sized bet on the flop.

To this point, we've considered only the rigid pot-limit case, where the only bet size allowed is a pot-sized one. But in real pot-limit, either player can bet any amount up to the pot-how does this change the expected equities?

To examine this case, let's go back to the case where both players have two pot-sized bets remaining.

Example 7.3

The game is pot-limit holdem.

Player X has: A♥ A♦

Player Y has: 8♣ 7♣

The flop is: 9♣ 6♣ 2♠

(Again, we'll ignore runner-runner full houses for the AA and runner-runner two pair for the 87s for the sake of. discussion-assume that on each, street the 87 simply has 15 outs.)

The pot contains \$100, and both players have \$400 in front of them.

We've seen previously that X's EV from betting \$100 is -\$4.55, while his EV from betting \$0. (checking) is \$65.15.

What if we allow X to bet any amount between those values? For example, suppose X bets \$5.

We know that the action on the turn is deterministic-if Y hits on the turn, there will be no more betting, while if Y misses, X will bet either the pot or all-in. As a result, we should be able to express X's overall EV as a function of how much money goes into the pot on the flop total). Let x be the amount of money per player that goes in on the flop (up to \$100 - we will address the case where more than \$100 goes in momentarily). The pot will be (2x+\$100) on the turn.

X's expectation is as follows:

Outcome	$p(\text{Outcome})$	$\langle X \rangle$ from outcome
Y wins on turn	15/45	-x
Y wins on river	$(30/45)(15/44)$	$-x - (2x+100)$
X wins	$1 - (15/45) - (30/45)(15/44)$	$3x + 200$

$$\langle X, x \text{ on flop} \rangle = \left(\frac{15}{45}\right)(-x) + \left(\frac{30}{45}\right)\left(\frac{15}{44}\right)(-3x - 100) + \left(1 - \left(\frac{15}{45}\right) - \left(\frac{30}{45}\right)\left(\frac{15}{44}\right)\right)(3x + 200)$$

$$\langle X, x \text{ on flop} \rangle = \left(\frac{1}{3}\right)(-x) + \left(\frac{2}{3}\right)\left(\frac{15}{44}\right)(-3x - 100) + \left(\frac{2}{3} - \left(\frac{2}{3}\right)\left(\frac{15}{44}\right)\right)(3x + 200)$$

$$\langle X, x \text{ on flop} \rangle = \left(\frac{1}{3}\right)(-x) + \left(\frac{5}{22}\right)(-3x - 100) + \left(\frac{2}{3} - \frac{5}{22}\right)(3x + 200)$$

$$\langle X, x \text{ on flop} \rangle = (-x/3) - 15x/22 - 500/22 + \left(\frac{29}{66}\right)(3x + 200)$$

$$\langle X, x \text{ on flop} \rangle = -x/3 - 15x/22 - 500/22 + 29x/22 + 2900/33$$

$$\langle X, x \text{ on flop} \rangle = -x/3 - 1500/66 + 7x/11 + 5800/66$$

$$\langle X, x \text{ on flop} \rangle = 10x/33 + \$65.15$$

We can check this by comparing our values from the last analysis; $\langle X, \$0 \rangle = \65.15 , and also $\langle X, \$100 \rangle = \95.45 . This expression for $\langle X, \$x \rangle$ is only valid for values from \$0 to \$100. Above \$100, X can no longer bet the entire pot on the turn because the stacks are too small. Therefore, for values of x from \$100 to \$400:

Outcome	$p(\text{Outcome})$	$\langle X \rangle$ from outcome
Y wins on turn	15/45	-x
Y wins on river	$(30/45)(15/44)$	-400
X wins	$1 - (15/45) - (30/45)(15/44)$	500

$$\langle X, \$x \rangle = \left(\frac{15}{45}\right)(-x) + \left(\frac{30}{45}\right)\left(\frac{15}{44}\right)(-400) + \left(1 - \frac{15}{45} - \left(\frac{30}{45}\right)\left(\frac{15}{44}\right)\right)(400)$$

$$\langle X, \$x \rangle = \left(\frac{1}{3}\right)(-x) + \left(\frac{5}{22}\right)(-400) + \left(\frac{2}{3} - \frac{5}{22}\right)(400)$$

$$\langle X, \$x \rangle = (-x/3) + \left(\frac{5}{22}\right)(-400) + \left(\frac{2}{3} - \frac{5}{22}\right)(400)$$

$$\langle X, \$x \rangle = -x/3 + \$128.79$$

For $x = \$400$, for example, we get $-\$4.55$ from this expression.

So now we have a two-part function that is X's EV for x getting into the pot on the flop:

x on [\$0, \$100]:	$\langle X, \$x \rangle = 10x/33 + \65.15
so	$\langle X, \$0 \rangle = \65.15 and $\langle X, \$100 \rangle = \95.45
x on [\$100, \$400]:	$\langle X, \$x \rangle = -x/3 + \128.79
so	$\langle X, \$100 \rangle = \94.45 and $\langle X, \$400 \rangle = -\4.55




Figure 7.1. X's equity in Example 7.3 for various pot sizes

So let us again consider the case where X bets \$5. Y has an infinite number of choices of raise amounts, but now that we can see the behavior of X's EV function, it is easy to discard many of the choices. Y can make it a total of \$10 to \$115, but if he chooses any value from \$10 to \$100, he simply adds to X's EV. Likewise, he can choose any amount from \$100 to \$115, but since all the additional money above \$100 decreases X's EV, Y would choose to raise to \$115 if he were going to raise an amount about \$100.

So all that is left is to compare X's EV on these two actions:

$$\begin{aligned} \langle X, \$x \rangle &= 10x/33 + \$65.15 \\ \langle X, Y \text{ calls } \$5 \rangle &= (10/33)(\$5) + \$65.15 \\ \langle X, Y \text{ calls } \$5 \rangle &= \$66.67 \\ \langle X, Y \text{ raises to } \$115 \rangle &= (-1/3)(\$115) + \$128.79 \\ \langle X, Y \text{ raises to } \$115 \rangle &= \$90.46 \end{aligned}$$

Thus, Y will choose to simply call \$5.

By betting \$5, X has increased his equity by \$1.50. Up to some point, for each dollar X bets he gains $10/33$ of a dollar in equity, If he bets \$10, Y will again choose to call, but his equity is increased an additional \$1.52 to \$68.19. In fact, X can continue to bet larger and larger amounts until the values from Y's two options are equal. This will be the point at which X's expectation is maximum. When X bets the critical amount, Y has two options: to make the amount that goes into the pot \$x, or to make a pot-sized raise and make the amount that goes into the pot $(3x + 100)$.

$$\begin{aligned} \langle X, Y \text{ calls } x \rangle &= \langle X, Y \text{ raises to } (3x + 100) \rangle \\ 10/33 (x) + \$65.15 &= -1/3(100+3x) + \$128.79 \\ 10/33 (x) &= -x + \$30.30 \\ x &= \$23.26 \end{aligned}$$

This is the proper amount For X to bet; it maximizes his expectation. No matter what action takes, X is guaranteed at least \$72.20 in expectation:

$$\langle X, x = \$23.26 \text{ on flop} \rangle = 10/33 x + \$65.15 = \$72.20$$

The upshot of all tills is that the stack sizes are of critical importance in determining how to play even when the cards are faceup. Additionally, as we can see from the results we obtained, deeper stacks favor the made hand because they allow the made hand to punish the draw on the turn when he misses. This runs somewhat counter to the usual wisdom about draws and made hands, which holds that draws benefit from deeper stacks and implied odds.

Another important concept from these examples is the following:

Good draws benefit a great deal from getting all the money in the pot early in the hand.

Example 7.4

Let's look at a common no-limit case.

No-limit holdem.

Player X has: A♥ K♦.

Player Y has: Q♥ Q♣.

It is preflop. Player Y posted a blind of \$100. The stacks are \$800.

Let's say that the hands are revealed sequentially, such that player X sees his AK first, acts, and then player Y reveals his queens. Assume that player X raised to \$300 (an arbitrary but reasonable amount). What should player Y's strategy be?

AK benefits from seeing all five cards- it will only pair a third of the time on the flop, but will win 43% over all against QQ. The deeper the stacks, then, the more that the AK should want to get all the money in and the QQ should try to see the flop and push the AK out after the Hop if no A or K flops, right?

Wrong.

First, let's look at the shallower stacks:

Y with QQ is more than a 57% favorite; if he raises to \$800, player X will be forced to call, getting 3 to 2 on his money. Raising to some lesser amount will not be better, as X can simply reraise to \$800 to guarantee his equity mere. Jamming costs Y \$700.

$$\langle Y, \text{jam} \rangle = (0.5717) (\$800 + \$800) - \$700$$

$$\langle Y, \text{jam} \rangle = \$214.72$$

On the other hand, let's say that Y with QQ calls. Then there will be \$600 in the pot. The flop will contain an ace or a king but no queen about 30% of the time. In these cases, X with the AK can simply jam the flop, and Y will fold. The remaining 70% of the time, Y with the QQ will jam and X will fold with the AK.

$$\langle Y, \text{call} \rangle = p(\text{A or K flops}) (Y \text{ loses his call}) + p(\text{A or K doesn't flop}) (\text{pot size})$$

$$\langle Y, \text{call} \rangle = (0.3) (-\$200) + (0.7) (\$400)$$

$$\langle Y, \text{call} \rangle = \$220$$

So Y does slightly better here by calling and seeing if X hits the flop, then jamming if the AK misses.

Increase the stacks to \$1800, however, and now the tables are turned;

$$\langle Y, \text{jam} \rangle = (p(Y \text{ wins})) (\text{pot size}) - \text{cost of jam}$$

$\langle Y, \text{jam} \rangle = (0.5717) (\$3600) - \$1700$

$\langle Y, \text{jam} \rangle = \358.12

$\langle Y, \text{call} \rangle = \text{still } \220

So Y does much better by jamming immediately.

This runs counter to many players' intuitions about how the QQ/AK matchup might play out - with deep stacks, the QQ wants to get all the money in because of its significant equity advantage, while with shallower stacks it prefers to push the AK out of the pot on favorable flops. This conclusion is a little artificial-here the QQ can fearlessly push all his chips in the pot because he knows he has the advantage. In real holdem, however, he must fear AA or KK, and so he can't play in quite this way. However, many players overvalue seeing the flop, forgetting that their hand is not a simple "coin flip," but a substantial favorite over AK.

Going back to our original hypothesis that the action on each street should go "the made hand bets, and the draw either calls or folds depending on whether it has pot odds," we've seen that there are a number of situations where this is not at all the case. In pot-limit, for example, the size of the remaining stacks has a profound impact on the proper strategy for each player, and this is true in limit as well. In no-limit we saw a case where the stack sizes were integral to deciding whether to jam or just call.

Despite the occasionally interesting situations that occur in games where cards are exposed, you are unlikely to happen upon such a game. We present these situations, however, as primer material for the cases to follow; all the hand-reading skills in the world are of no use if one cannot play accurately when the cards are face-up, and while principles are often valuable as rules of thumb, recognition of situations where play might deviate from what is intuitively apparent can be worth equity.

Key Concepts

- Even when playing with the cards face-up, counterintuitive situations can occur.
- When considering the last raise all-in, draws should consider being more aggressive than is indicated by simple pot odds analysis. Further betting cannot hurt them and the made hand cannot get away if the draw hits when the money is all in.
- Stack size is of critical importance in big bet poker; even with the cards face-up changing the stacks changes the nature of the play drastically. Situations can even occur where the favorite has to fold to a bet from the underdog because of the potential for betting on later streets.
- Good draws (those which have close to 50% equity in the pot) benefit greatly from getting all the money in on the flop. However, if they cannot get all the money in (or enough that they cannot be hurt by later betting), they prefer to get as little as possible.
- Made hands often want to carry through a full bet for a later street when playing against a draw because of the heavy value extraction that can occur when the draw misses.

Chapter 8

Playing Accurately, Part II: Hand vs. Distribution

Concealed hands make for a much better game: hence the popularity of poker games with hidden information. We now consider situations where only one hand is revealed, while the other hands are some sort of distribution. For these cases we can use the expected value techniques previously discussed to arrive at the best play, if we can come up with an accurate characterization of the opponent's strategy. For the time being, we will not concern ourselves - with our opponent's changing strategies, but simply assume that we are able to estimate reliably how he will play his various hands. We will return to the idea of players adapting their strategies and the value of an entire distribution changing as the result of different strategies with individual hands at length in later chapters.

But before we get too deeply into analyzing particular games, we can discuss a concept that plays a valuable role in attempting to analyze games of more complexity. This concept is of particular value in games such as holdem. Frequently, two hands put in some action before the flop, and then one or the other of the hands calls while there are still a significant amount of chips to play with postflop. When the play stops with a call in this manner, we frequently want to know what the expectation for the hand will be for both players in order to compare that expectation with the expectation of, for example, raising again.

This value has two components; one is what we frequently call *showdown* equity, which is the expectation that each hand has if we simply stopped the betting right now and dealt all the cards out immediately. The other is what we call *ex-showdown* equity, the expectation that each hand has in the post-flop betting. Ex- is a Latin prefix meaning "outside"—hence ex-showdown equity is the equity outside of the current pot. The sum of these two values is the total expectation a player has in the pot. These two values are actually related because they are both dependent on the strength of the distributions that each player holds. It's quite possible for a player's showdown equity to be positive while his ex-showdown equity is negative. One example of this is if a player holds an extremely weak draw and his opponent bets. His showdown equity is equal to his chance of winning times the pot before the bet. However, since he must fold, his ex-showdown equity is actually the negative of that value since he is losing), making his total equity from the hand zero.

We would love to be able to project out the entire play of the hand over all possible boards and the players' strategies in order to see what the EV of playing the hand out in that manner would be. Unfortunately, this is a fairly intractable problem, especially at the table. What we try to do is train our intuition to guess at these equities within a certain band. For example, consider a hand where both players' distributions are approximately equal and both players of equal skill. Then the showdown values will be equal, and the most important influence on the ex-showdown value for these two players will be position; the player who has the button will have a higher ex-showdown value than the other because of his positional advantage.

Other types of situations might change this; for example, consider a situation where one player is known to hold a quite strong distribution, such as [AA, KK, QQ, AK] and the other player holds a random hand. In this case, the player with the strong distribution will have much higher showdown equity than his opponent. However, his advantage ex-showdown will be smaller than if he held one of those hands in a situation where he could hold a wide distribution. We can also state the equity of a particular hand within a distribution. For example, assuming a typical button-raising range in limit holdem, it might be the case that against a strong player in the big blind, holding aces on the button might be worth as much as four or five small bets ex-showdown.

When analyzing games where players might call preflop, we often refer to a player having "X% of the pot" after the call. In some cases, such as when the money is all-in preflop, this is strictly showdown equity. In other cases, however, we are trying to capture the effect of future betting on the players' expectation. We will see an example of this in a moment.

Many types of exploitation are quite straightforward. For example, against an opponent who bluffs too often on the end, the most profitable play is to call with additional hands that can beat only a bluff, while against opponents who fold too often to bets on the end, the highest expectation play with weak hands is generally to bluff. We looked at an example of this type at the outset of Part II. These cases are rules of thumb that often adequately represent the expectation analysis appropriate to those situations. We can analyze more complicated situations, however, by simply applying a methodical approach—evaluate the EV of each potential strategy, and choose the one that is the highest.

Example 8.1

We're the big blind in a \$5-\$10 blind no-limit holdem game. The button, a player we know well, has about \$200 in front of him. It's folded around to him and he opens for \$30. The small blind folds.

We know the following things about the button (through the magic of toy game technology):

- The button's raising distribution here is {22+, A2+, KT+, K9s+, QTs+, QJ+, JTs, T9s}.
- If we jam (re-raising \$170 more), he will call with AA-JJ and AK and fold all other hands.

If we call, we can expect to have some portion of the pot, depending on how many hands we call with. If we called, for example, with the same range of hands that he is raising, we could expect to have approximately 45% equity in the resultant pot (the shortfall due to position) from the postflop play.

What is our maximally exploitive strategy?

The opponent will raise with a total of $\frac{350}{1326}$ hands:

6 ways to make each of thirteen pairs = 78

16 ways to make each of twelve ace-high hands = 192

12 ways to make each of four other unsuited hands = 48

4 ways to make each of eight other suited hands = 32

He only calls, however, with 40 of those hands. These proportions are influenced slightly by *card removal* effects (if we hold an ace, he has fewer ways to make AA, etc). This is the basic proportion, and we will use it in this analysis.

We are risking \$190 (the remainder of our stack after posting the blind) to win the \$45 in the pot. If we simply jam with all hands, we will win the pot immediately ($\frac{310}{350}$) times, or 88.57% of the time. In addition, when we are called, despite having poor equity in the pot, we do still win sometimes. The equity of our random hand against his distribution of AA-JJ and AK is about 24.95%. Note that even against a very strong range, we still have nearly a quarter of the pot with two random cards.

This results in an overall EV for this jamming play of:

$$\begin{aligned} \langle \text{jam} \rangle &= [p(\text{he folds})] (\text{pot}) + (p(\text{he calls})) (p(\text{we win}) (\text{new pot}) - \text{cost of jam}) \\ \langle \text{jam} \rangle &= (0.8857) (\$45) + (1 - 0.8857) ((0.2495) (\$200 + \$200 + \$5) - \$190) \\ \langle \text{jam} \rangle &= \$29.69 \end{aligned}$$

This is a fairly important point. We have positive equity by simply jamming any hand here, no matter how weak. Our opponents strategy has a significant flaw; he folds too often to reraises.

We can also look at specific hands-take for example, 32o. Against {JJ+, AKs, AKo}, 32o has 21.81% of the pot.

Then our EV of jamming is:

$$\begin{aligned} \langle 32o, \text{jam} \rangle &= (p\{\text{he folds}\})(\text{pot}) + (p(\text{he calls}))p(\text{we win})(\text{new pot}) - \text{cost of jam}) \\ \langle 32o, \text{jam} \rangle &= (0.8857) (\$45) + (0.1143)((\$405)(0.2181) - \$190) \\ \langle 32o, \text{jam} \rangle &= \$28.24 \end{aligned}$$

We can contrast this to calling, where we call \$20 and wind up with some percentage x of the pot.

$$\begin{aligned} \langle \text{call} \rangle &= x(\text{new pot}) - \text{cost of call} \\ \langle \text{call} \rangle &= x(\$65) - \$20 \end{aligned}$$

In order to make calling as good a play as jamming with any specific hand, this EV has to be greater than \$29.69.

$$\begin{aligned} x(\$65) - \$20 &> \$29.69 \\ x &> 76.45\% \end{aligned}$$

It should be clear that no player, no matter how impossibly gifted, has more than 76% of the pot postflop with a random hand out of position against a button raiser.

But what about the case where we have aces in the blind? Holding two of the aces makes some significant difference in the hands that are raised, so we adjust the calling frequency using Bayes' theorem. After adjusting for the card removal by removing two aces from the deck, the button raises 249 hands and calls with just 27 of them, folding $\frac{222}{249}$. Against his calling hands, our equity is much higher - 83.43%.

$$\begin{aligned} \langle \text{AA}, \text{jam} \rangle &= p(\text{he folds})(\text{pot}) + [p(\text{he calls})][p(\text{we win}) (\text{new pot}) - \text{cost of jam}] \\ \langle \text{AA}, \text{jam} \rangle &= (\frac{222}{249})(\$45) + (\frac{27}{249})(.8343)(\$405) - \$190) \\ \langle \text{AA}, \text{jam} \rangle &= \$56.16 \end{aligned}$$

To make calling with aces correct, we need at least x% of the pot where:

$$\begin{aligned} X(\$65) - \$20 &> \$56.16 \\ x &> 117.20\% \end{aligned}$$

It is not impossible that this could be the case, especially against an opponent who is aggressive postflop. But this is the very best case out of all the possible hands that we could hold.

So, in summary of this example, jamming with any hand is a substantially stronger play against this type of button raiser than either folding or calling. This holds except when the blind has a very strong hand; in which case it's likely close, depending on how well the players play

postflop.

What this type of analysis leads us to are the flaws in our opponents' strategies—the place where they give up ground. In the above example, the flaw in the button raiser's strategy was that he raised with a wide range of hands but surrendered too often to reraises. In truth, the only way to be sure that this is the case is to do an EV analysis as above, but at the table, it is fairly easy to do off-the-cuff analysis such as the following:

"This guy is raising 25% of his hands, but he's only calling with few. If I jam, I'm putting in about 4 pots to win 1. If he calls me less than a fifth of the time, I'm making money on the jam immediately, and sometimes I'll win when he calls too!"

Finding the weaknesses in opponents' strategies is the core of exploitive play. In the above case, it was an imbalance between the number of raising hands and hands that would stand a reraise in the button player's strategy. This imbalance created the opportunity for exploitation by an aware big blind, who could profit by jamming with any two cards. If the button player wanted to prevent this type of exploitation, he would need to do one of two things. He could either tighten his initial raising requirements or loosen the range of hands with which he would call a jam from the big blind.

Another way to exploit one's opponents is to take advantage of advantageous flops that don't appear to be so at first glance.

Example 8.2

Consider the following situation in a limit holdem game:

You are dealt 99 under the gun in a full game and raise (stipulate that your distribution of hands for this action is AA-99, AK, AQ, and AJs). A player several seats to your left reraises (you judge his distribution to be AA-TT, AK and AJs), and the field folds. You call.

The flop comes A72 with no suits. Your estimation of the other player's likely strategy for playing the hand from the flop on is roughly as follows: (We'll ignore the possibility of hitting a set on the turn for simplicity)

If you check the flop, he will bet the flop with all hands.

If you bet the flop, he will call with all hands.

If you check-raise the flop, he will call the check-raise with all hands.

His turn strategy is as follows:

Your flop action	Your turn action	AA	KK/QQ	JJ/TT	AK	AJs
Check-raise						
	Bet	Raise	Fold	Fold	Raise	Call
	Check-call	Bet	Bet	Bet	Bet	Bet
	Check-raise	Bet/Reraise	Bet/Fold	Bet/Fold	Bet/Call	Bet/Call
Check-call						
	Bet	Raise	Call	Call	Raise	Raise
	Check-call	Bet	Check	Bet	Bet	Bet
	Check-raise	Bet/Reraise	Check	Bet/Fold	Bet/Call	Bet/Call

What is your plan?

A typical thoughtful player would begin by noting that given the action and distributions - stipulated, 99 is never ahead, and against all the hands in the opponents distribution, has just two outs to improve. This fairly bad situation seems to argue for check-folding as the proper course of action.

But a stronger method of evaluating the problem is to consider the comparative equities of different types of strategies. The check-fold strategy clearly has an EV of 0, so that's a baseline against which we can compare other strategies.

First, we should consider how many hands of each type the opponent will hold. Given our hand and the flop, the opponent holds the following hands with the given frequency:

AA: $\frac{3}{42}$
 KK, QQ, JJ, TT: $\frac{6}{42}$ for each pair
 AK: $\frac{12}{42}$
 AQs: $\frac{3}{42}$

Let's consider just a few candidate strategies:

- 1) Check-folding the flop.
- 2) Check-calling the flop and check-raising the turn, giving up if the opponent pushes back or stays until the river.
- 3) Check-raising the flop and playing aggressively through the turn, giving up if the opponent pushes back or stays until the river.
- 4) Betting out on the flop and turn.

Strategy 1) has an EV of 0.

Strategy 2) performs as follows against the opponent's hands:

vs AA, AK, AQs: -5 small bets

vs JJ-TT: +10.5 small bets (7.5 from the pot, one on the flop, two on the turn)

vs KK-QQ: -3 small bets

$$\langle \text{strategy 2} \rangle = p(\{AA, AK, AQs\})(-5) + p(\{JJ, TT\})(10.5) + p(\{KK, QQ\})(-3)$$

$$\langle \text{strategy 2} \rangle = (\frac{3}{7})(-5) + (\frac{2}{7})(10.5) + (\frac{2}{7})(-3)$$

$$\langle \text{strategy 2} \rangle = 0$$

Strategy 3) performs as follows against the opponent's hands:

vs. AA, AK, AQs: - 4 small bets

vs. KK-TT: +9.5 small bets

$$\langle \text{strategy 3} \rangle = p(\{AA, AK, AQs\})(-4) + p(\{KK, QQ, JJ, TT\})(9.5)$$

$$\langle \text{strategy 3} \rangle = (\frac{3}{7})(-4) + (\frac{4}{7})(9.5)$$

$$\langle \text{strategy 3} \rangle = 3.71$$

Strategy 4) performs as follows against the opponent's hands:

vs AA, AK, AQs, KK, QQ: -3 small bets

vs JJ, TT: +8.5 small bets

$$\langle \text{strategy 4} \rangle = p(\{AA, AK, AQs, KK, QQ\})(-3) + p(\{JJ, TT\})(8.5)$$

$$\langle \text{strategy 4} \rangle = (\frac{5}{7})(-3) + (\frac{2}{7})(8.5)$$

$$\langle \text{strategy 4} \rangle = 0.29$$

So interestingly enough, the immediately intuitive solution seems to have the worst EV of any of these possibilities!

The clear winner, however, is the strategy of check-raising the flop and betting the turn. And the value of this strategy (given this strategy by the opponent) is almost 4 small bets! Against an opponent who plays this way, it is a gigantic mistake to check and fold on the flop in this situation; we know this because we calculated the equities. But why, intuitively, is this strategy so good?

The re-raiser has made the pot six bets before the flop. Post-flop, however, when an ace flops, he is folding more than half his hands to only four bets worth of aggression. The original raiser is putting in four small bets to win six and the re-raiser is folding more than half the time. What's happened here is that a flop that appears at first glance to be poor for the nines (and which would be if the hands were face up) is actually a strong one *because of the shape of the re-raiser's distribution of hands and the way in which he will play those hands*. If the re-raiser wanted to prevent this type of exploitation, he would need to call down with hands such as KK and QQ instead of folding them.

Example 8.3 (Example 5.1 revisited)

As a final example of this process of attempting to maximize EV, we return to the stud eight or-better hand from Chapter 5 (Example 5.1).

To recap the hand:

We: (5♣A♥) 4♥8♠ J♥A♠

Opponent: (??) 6♣ T♦ K♣ 3♣

The antes are \$5, and the limits are \$30-60. There is \$345 in the pot and we are first on sixth street. We have narrowed the opponent's range to the following hands:

AA	-1
Q♣ Qx	-3
Qx Qy	-3
J♣ Jx	-2
Jx Jy	-1
X♣ Y♣	-6

We consider the play against each of these hands in turn. In each of these cases, we consider the most likely way the play will go, assuming that there are not many branches of equal value. The strategy on later streets here is a little more complicated in practice with very strong players because of small fractions of semi-bluff raises and so on. We will also ignore the check-raising possibility; this hand in this position should almost certainly not check raise.

As we proceed through the cases, there are a fair amount of places where we speculate as to how the action will go on later streets. It is often not that easy to predict the shape of future action, and so a lot of estimates will be made. However, as long as we do not introduce a systematic bias toward overestimating or underestimating our equity, the various errors introduced by these speculations should at least partially balance each other; each error is individually quite small in the whole scheme of the hand.

Against AA:

When deciding whether to bet against a hand of approximately the same strength as our own, one important consideration is that often the hand result is the same either way. If we bet, it would seem overly aggressive for the opponent to raise with AA, since we could already have a low made, in which case we're freerolling to scoop (perhaps with a wheel draw). In that case we could also immediately re-raise for value. So if we bet, it seems clear that AA should simply call. But if we check, the hand with aces and a flush draw will bet; so it doesn't matter what we do. One bet goes into the pot in either case. This situation occurs often between hands of approximately equal strength.

$$\langle \text{bet} \rangle = \langle \text{check} \rangle$$

Note that we're not really interested in what the absolute expectations of the various actions are; we're primarily interested in the difference between the candidate actions. It should be fairly clear that we have enough equity in this situation that we don't want to fold, so all we need to pay attention to is the difference in expectation between the two reasonable actions.

Against $Q\clubsuit Qx$ and $J\clubsuit Jy$:

Now we have the best hand for high and the best draw for low. If the game were stud high, where there was no low possibility, it would be clear to bet and extract value from the draw. But the situation here is much stronger, as even if the flush comes in or the opponent makes two pair, we still have 12 out of 39 remaining cards that give us half the pot. So this case seems clear; we should bet and the opponent will call with his fairly strong high draw. We can likely use the hand equities as a proxy for the total equities, including betting on the river. Occasionally the opponent will make a flush and a better low or we will miss our low. Against $Q\clubsuit Q\spadesuit$, we have about 70% equity in the pot.

$$\langle \text{bet} \rangle = p(\text{win})(\text{new pot value}) - \text{cost of bet}$$

$$\langle \text{bet} \rangle = (0.70) (\$345 + \$120) - \$60$$

$$\langle \text{bet} \rangle = \$265.50$$

$$\langle \text{check} \rangle = p(\text{win})(\text{pot value})$$

$$\langle \text{check} \rangle = (0.70) (\$345)$$

$$\langle \text{check} \rangle = \$241.50$$

for a expected gain of approximately \$24 from betting.

Against $Qx Qy$ and $Jx Jy$:

In this case, we are against an even worse hand. This is the same as the previous situation, but with no flush draw.

This case is even clearer; we should bet. Now the opponent has a difficult decision about whether to fold or not. He would consider our possible distribution of hands (aces up, aces with a low draw, a made low with only a weak high draw, etc.). But with the strong likelihood that we have made at least a pair of aces, the opponent will probably have to fold in this case. If we had checked, the opponent will probably call all the time on the river, figuring us for perhaps a pair lower than aces.

$$\langle \text{bet} \rangle = \$345$$

$$\langle \text{check} \rangle = p(\text{win})(\text{new pot value}) - \text{cost of bet}$$

$$\langle \text{check} \rangle = (0.85)(\$345 + \$120) - \$60$$

$$\langle \text{check} \rangle = \$335.25$$

Here we gain about \$10 by betting because we are able to claim the entire pot immediately rather than allowing the opponent to draw.

Against X♣ Y♣:

In this case, if we bet, we have the difficulty. The opponent has a made flush and a low draw, so he will raise. Now we have to consider calling, his all likelihood, our low draw and the chance that the opponent is semi-bluffing will tip the scales toward calling. However, it's dear that against this hand, we do not wish to bet. If we check, of course, the opponent will bet. On the river, in either case, we will call if we hit the low draw. Occasionally we will still be scooped when the opponent wins low anyway. So in the betting case, we will put in two bets on sixth street and one on seventh when we hit a low card or aces up, and just two bets on sixth when we miss completely.

We have about 13% equity against the opponent's hands when we hit a calling card. Our calling cards are two aces, three of each card from deuces through eights, and three jacks. We fold on two tens, two kings, and four queens. So $\frac{26}{34}$ of the time we have the following equities:

$$\begin{aligned} \langle \text{bet} \rangle &= p(\text{win})(\text{new pot value}) - \text{cost of bet} \\ \langle \text{bet} \rangle &= (0.13)(\$345 + \$360) - \$180 \\ \langle \text{bet} \rangle &= -\$88.35 \\ \langle \text{check} \rangle &= p(\text{win})(\text{new pot value}) - \text{cost of bet} \\ \langle \text{check} \rangle &= (0.13)(\$345 + \$240) - \$180 \\ \langle \text{check} \rangle &= \$43.95 \end{aligned}$$

And the following occurs the other $\frac{8}{34}$ of the time:

$$\begin{aligned} \langle \text{bet} \rangle &= -\$120 \\ \langle \text{check} \rangle &= -\$60 \end{aligned}$$

Weighting these by their probabilities, we have:

$$\begin{aligned} \langle \text{bet} \rangle &= p(\text{don't hit calling card})(\text{bet equity}) + p(\text{hit calling card})(\text{bet equity}) \\ \langle \text{bet} \rangle &= \frac{26}{34}(-\$88.35) + \frac{8}{34}(-\$120) \\ \langle \text{bet} \rangle &= -\$95.80 \\ \langle \text{check} \rangle &= p(\text{don't hit calling card})(\text{check equity}) + p(\text{hit calling card})(\text{check equity}) \\ \langle \text{check} \rangle &= \frac{26}{34}(\$43.95) + \frac{8}{34}(-\$60) \\ \langle \text{check} \rangle &= -\$47.73 \sim -\$48 \end{aligned}$$

Summarizing all of these possibilities:

Opp. Hole Cards	Probability	$\langle \text{Bet} \rangle - \langle \text{Check} \rangle$	Weighted EV
AA	1/16	+	0
Q♣ Qx or J♣ Jx	5/16	+\$24	\$7.50
Qx Qy or Jx Jy	4/16	+1\$0	\$2.50
X♣ Y♣	6/16	-\$48	-\$18
Total	1		-\$8/hand

So it turns out, in this analysis, that betting is actually about \$8 worse per hand than checking in this spot! Now we could add additional hands in our "lost his mind" category; for most of these

hands, we are better served by betting than by checking. But any reasonable distribution of hands will still result in the finding that the “automatic” bet that would be made by many players is actually worse than the alternative strategy of checking.

This example had many assumptions throughout. You may disagree with our characterizations of how players would play, what outcomes we could expect, and so on. The central point here is not to quibble over the play of a specific stud eight-or-better hand. Instead we presented it first under the topic of hand reading, to show the Bayesian framework by which a hand unfolds, and how appropriately adding each newly revealed piece of information throughout a hand, making the proper inferences based on the streets that have gone before, and forecasting what the action will be on later streets, can provide information. We then continued the example under the topic of accurate play and used that information to make a correct EV analysis of our hand against the assumed distribution our opponent held.

Key Concepts

- After the process of hand reading and strategy reading is complete, we can take the information gathered by those processes and make decisions about which actions to take by comparing the EV of different actions.
- Identifying the weaknesses in our opponents' play qualitatively will often help us to identify the right exploitive play even if we cannot do all the calculations at the table.
- Developing intuition for situations that commonly call for exploitation is extremely helpful in playing exploitively at the table.
- When our opponents hold distributions of hands, often we will want to take different actions against their different hand types. Weighting our expected values by the Bayesian-adjusted probabilities that they hold each hand type can allow us to make an accurate decision, even against their entire distribution of hands.
- Showdown equity is the expectation that each hand has if we simply stopped the betting and dealt all the cards out immediately. Ex-showdown equity is the expectation that each hand has in the post-flop betting.

Chapter 9

Adaptive Play: Distribution vs. Distribution

In the preceding four chapters, we have laid out a straightforward approach to playing single hands against a static or predictably dynamic opponent; we assign the player a distribution of hands, we attempt to plausibly infer what his strategy will be with the hands he can hold, and then we attempt to maximize our EV against his distribution of hands and the strategy he will employ. This approach will extract the maximum on a particular hand from any opponent who plays either a fixed strategy or a strategy that deviates in a predictable way.

But on the other hand, let's say that the opponent is also a strong player. Then he will be attempting to do the same to us (to read our hands and play, and attempt to exploit the weaknesses in our strategy). This means that we must take care in how we attempt to exploit the opponent, because if we exploit him in too predictable a way, he can easily alter his strategy to exploit us in turn. We call this *counter-exploitation*.

Let us look at an example. Suppose some situation comes up, and our information gathering finds that the opponent bluffs about 10% too seldom in this particular situation. The maximally exploitive strategy in this situation, then, is to fold all the time with our hands that can beat a bluff, since each hand with which we call will win less money by picking off bluffs than it will lose to value bets.

Our strong opponent, however, will not be ignorant of our change in strategy. By the same process that we were able to read his hands and strategy and identify this weakness, he now sees a weakness in our strategy; that in this situation, we fold too often. So he in turn can change his strategy to bluff much more often in this situation. And we in turn notice this and swing our calling frequency back to counteract this, and so on. The maximally exploitive strategies in each of these cases call for a sudden swing of changing strategies; from calling all the time in a particular situation to never calling in a situation, and so on. But by creating such severe swings, we open ourselves up to being counter-exploited.

An alternative strategy in this circumstance might be to adjust our strategy by a smaller amount. This would delay detection by our opponent, but also reduce the amount of expectation gamed through the exploitation process. Another strategy would be to simply attempt to play the game of exploitation and counter-exploitation, trying to stay one step ahead of the opponent. In this circumstance, then, the player who is best able to accomplish this (staying ahead of the opponent) will gain value and the other will lose it.

Example: Playing Strong Draws

With these concepts in mind, then, we consider the main subject of this chapter. In previous chapters, we examined the play of a single hand against an opponent's distribution of hands. However, in reading our opponent's hands and strategy, the distribution we hold is often of prime importance. When we exploit our opponent, we take advantage of the information that is available to us through his betting patterns. In the same way, when we play against an opponent who will be attempting to exploit us, we must be aware of how he might use the information that our strategy gives up and attempt to minimize the effect of his counter-exploitation as we construct our strategy.

We will look at a special type of hand in no-limit holdem; the "strong" flush draw. This is the flush draw that also contains secondary draws, such as overcards, gutshot straight draws, or open-ended straight draws. These types of hands normally have anywhere from 12 to 15 outs against top pair, and sometimes fewer against very powerful made hands (such as sets).

However, they do normally need to hit one of their outs in order to win. A hand with twelve outs on the flop has approximately 45% equity in the pot.

You will recall that in the section on cards exposed situations, we showed how high-equity draws frequently benefited from getting all the money into the pot early in the hand. This occurred because when they missed on the first card, they could be punished by the made hands before the river card came. However, when they got all the money in the pot on the flop, they simply were able to claim their equity.

In future chapters, we will look at bet sizing in an in-depth manner. However, for the purposes of this discussion, we will simply state a principle and justify it intuitively. First, we should normally not bet a large multiple of the pot for value: our opponent can then wait for very strong hands and call, folding all his weaker hands. If we risk, for example, ten pots to win one, the opponent only has to call occasionally to make our play unprofitable. Perhaps you have played against players who frequently jam in the early rounds of no-limit tournaments for hundreds of times the blind with hands such as AJ. Picking such players off with strong hands is the obvious and correct exploitive response. We will claim that instead the bettor should bet a smaller amount with a wider range of hands, responding to a raise by jamming with a fraction of those hands, and so on.

It would seem, then, that our strong flush draws might make up a portion of that "wider range of hands." But here we have a problem. By playing these draws in this manner, they lose a great deal of their value, because the opponent can simply call the flop. When we have a strong made hand, such as two pair or trips, this is not a problem; we simply bet again on the next street, hoping to extract even more value. But for these strong flush draws, it is a poor situation. In fact, as we saw previously, this is precisely the way that the made hand wants the action to go. Instead, our best play with a flush draw specifically is often to jam, even with a stack where we would not jam with other hands.

Consider our opponent's dilemma if we jam a flush draw. Normally, his response to our large overbet would be to pick us off with his best hands. But now he has a problem. When he calls with his best made hands (presumably sets or perhaps straights), the flush draw still has at least 30% equity in the pot. This will likely create a problem for him. For example, presume that we jam a bet of three into a pot of one, and we'll assume that we have just 33% equity when he calls. If he picks us off with his best 25% of hands, for example, we still gain equity by jamming because we pick up one unit 75% of the time, and the remainder, we lose less than two thirds of a unit. Now, perhaps the opponent will call much more frequently, with one pair-type hands. However, this is where the description of these draws as "strong" flush draws matters. Against one-pair hands, ace-high flush draws often have additional equity, as hitting an ace on a later street is good. Straight draws also add value, and so on. As the opponent calls more frequently, the equity of the flush draw when called increases, and so jamming is still positive EV, as long as the opponent folds a fair amount of hands (how many is dependent on the specific situation).

So at a stack size of perhaps three times the pot, it seems we would prefer to be able to jam our strong flush draws and bet smaller amounts with our made hands and weaker semi-bluffs. But this strategy is still exploitable! Our opponent can switch to calling our jams with all hands that beat our likely draws, and to playing a normal strategy against our smaller bets, knowing that we have fewer strong flush draws in our distribution for making the smaller bets.

But now we can move to exploit this in turn. If the opponent will call our jams with mediocre hands because they are all high-equity flush draws, then we can exploit that by adding some of our bigger hands (sets and two pair, for example) to our jamming range. Now our opponent has

to choose between paying off our monsters and getting value from our flush draws, or dodging our monsters, but losing a large fraction of the pot against our flush draws. Our opponent might try in turn to exploit this strategy; however, it seems that he has no obvious path to do so in a macroscopic way (as he did against prior strategies).

This brief strategy construction highlights two important points that are worth noting about playing distribution against distribution. The first is *information hiding*. When the candidate strategy played different hands in different ways, it was much more exploitable because the opponent could trim the hand-reading tree much more effectively. Strong strategies play many hands in the same way, making it difficult for the opponent to read the distribution of hands remaining. Hiding information is the most important part of preventing counter-exploitation. In the example above, playing both strong hands and draws in the same way made it difficult for the second player to respond exploitively.

Another way to see the importance of information hiding is by considering the information that we desire in order to exploit the opponent. Generally, the narrower a distribution of hands and the more obvious a strategy, the easier it is to find an exploitive strategy. However, against opponents who play a wide range of hands in exactly the same way, we have much more trouble finding exploitive strategies, because the distribution of hands they hold is so poorly defined.

The second principle may seem obvious, but it is also important. Suppose we have equally numerous hands A and B, and two ways of playing them, X and Y. We've decided for information hiding reasons that we want to play A and B in the same way. Hand A does a little better being played in way X than in way Y while Hand B does a lot better being played in way Y than in way X. Then we should, in the absence of a compelling reason not to do so, play both types of hands in way Y. This occurred in our example above when we considered playing strong made hands. These hands would probably do slightly better (in a vacuum being played by betting a smaller amount. However, when we jammed them (within the context of the distribution), the entire distribution gained value because the opponent could no longer exploit us. Another good example of this point is the preflop play of AK or AA. In many situations (especially a little later in a tournament), AK has very strong equity for a preflop all-in but can suffer when played for a third of one's stack. AA, by contrast, does well in both scenarios. If we decide we want to play AA and AK in the same way preflop, then by this principle we should jam with both.

Attempting to play the game of counter-exploitation often assumes that we have fairly good information about the opponent's strategy. In real play, exploitive play relies on highly incomplete information: a balance must be struck between trying to exploit specific weaknesses of the opponent's strategy and playing in a manner that extracts value from a wide variety of strategies when no specific weakness is known or when we have insufficient information about the opponent's play to justify playing purely exploitively. In Part III, we will study the process of trying to find *optimal* strategies, which are strategies that maximize their equity against the opposing strategies that maximally exploit them. But it is not necessary to solve a game completely in order to create what we call a *balanced* strategy.

Balance

This term, unlike some of the other terms we use when referring to strategies, is not a technical term from game theory. To define it, though, we introduce one additional term that is from game theory. The *value* of a strategy is the expectation of that strategy against the opponent's maximally exploitive strategy. We sometimes also call this the *defensive value* of a strategy, because value can be easily confused with expected value. Balanced strategies are those strategies whose defensive value is close to their expectation against the optimal strategy.

Essentially, balance is a measure of how exploitable a strategy is.

A short example may help to illustrate this. Consider a game where it is the river, and Player A has either completed a flush draw or has no pair and where Player B has two pair. The draw was closed. We've seen previously that Player A will bet all his Bushes and then some fraction of his weak hands as bluffs. The fraction of bluffing hands will depend on the size of the pot. We'll assume the pot is three bets, and we'll assume that player A has made his flush 20% of the time.

We consider a number of candidate strategies for A and find the expectation of B's maximally exploitive strategy (MES) against them. B's MES in this case is to fold all the time if A bluffs less than 5%, obtaining an ex-showdown equity of $3x$, where x is A's bluffing frequency

$$\langle B, \text{fold} \rangle = 3x$$

When A bluffs more than 5%, B switches to calling all the time:

$$\langle B, \text{call} \rangle = x - 0.2$$

% of total hands A bets as a bluff	$\langle B, \text{MES} \rangle$ (ex-showdown)
0%	0
2%	- 0.06
4.8%	- 0.144
5%	- 0.15
5.2%	- 0.148
10%	- 0.1
20%	0
50%	0.3
80%	0.6

In this game, the optimal strategy is for A to bet 5% of his hands as a bluff. But as you can see, betting amounts slightly above and slightly below this value don't change A's expectation by a very large amount. To achieve these expectation gains, B must also employ the maximally exploitive strategies (that is, he must call all the time or fold all the time appropriately).

In this example, there is only one strategy variable, and so the degree of balance of a strategy is directly tied to how close that variable is to optimal. However, in real poker, we often have many different types of hands. As a result, there are additional variables to be balanced. For example, we may bet our best 30% of hands for value, and balance those value bets with 10% of our hands as bluffs. The optimal strategy for such a game might in fact be to bet 50% of our hands for value, and to bluff 17%. In this case, our original strategy sacrifices value, but it is still balanced and fairly difficult for the opponent to exploit.

In actually playing poker, it is often too difficult (due to computational infeasibility, etc.) to find the optimal strategy. However, we can often create strategies that are balanced and difficult to exploit; we can then play those strategies in situations where we do not have enough information to reliably exploit the opponent. This achieves two things. First, it enables us to extract most of the value that is available from playing exploitively. Second, it prevent; the opponent from exploiting us in turn; that is, he can normally do no better than to respond with a strategy that is

balanced in turn. If the opponent attempts to exploit our strategy, he will often sacrifice equity for no return.

Key Concepts

- Maximally exploiting opponent strategies often requires extreme swings in our own strategy. Many opponents, especially astute ones, will observe this and take steps to counteract our exploitation. When constructing strategies, we must be aware of the potential for counter-exploitation by our opponents.
- There are special strategies called optimal strategies. These strategies do not always extract the maximum value from our opponents, but the opponents cannot exploit them, even if they knew the strategy we were playing beforehand.
- Information hiding is an important principle of poker play that prevents the opponent from exploiting us as easily.
- We use the term balance to describe the "exploitability" of a strategy. Strategies that can be exploited heavily are unbalanced; strategies that cannot be exploited are balanced. Balance is not always correlated with profitability, but balanced play is often strong.
- Optimal strategies are perfectly balanced.

Part III: Optimal Play

Now the general who wins a battle makes many calculations in his temple ere the battle is fought. The general who loses a battle makes but few calculations beforehand. Thus do many calculations lead to victory, and few calculations to defeat.

Sun Tzu, The Art of War

Chapter 10

Facing The Nemesis: Game Theory

To this point, we have explored various methods and examples of exploiting weaknesses in our opponent's strategies that are weak in one way or another. It may be the case that all your opponents are ripe for exploitation and will never adapt or attempt to exploit you in turn. If this is the case, then exploitive play will serve to generate profit. However, many games contain a mixture of strong and weak players; and in a related vein, tournaments often preclude game selection decisions by providing players with no choice of opponents. In the extreme case, you may actually find yourself in a game with players stronger than you are.

In these circumstances, what can you do? One approach is to continue to try to exploit the opponent's mistakes, playing a regular exploitive game and effectively concede an edge to the opponent. This is the approach favored, or perhaps we should say used, by many players, often for lack of a better alternative.

And a related situation occurs fairly frequently - that is, playing against new opposition with whom we are not familiar. One approach to playing in this situation is to simply play a "default" exploitive game, assuming the new player meets a typical profile and playing accordingly, while attempting to gain information about his play and adjusting as indicated. Again, this is the approach favored by the majority of players, often because they lack another reasonable alternative.

We, however, advocate a different approach in both of these circumstances. Instead of relying solely on an exploitive strategy, we attempt to find strategies that are fundamentally sound; that yield no benefit to an opponent's attempts to exploit us. We may in this way give up a small amount of potential profit that we could have seized by playing exploitively. But by doing this, we attempt to make ourselves in turn unexploitable.

This approach to the game finds its roots in game theory, a subject that occupies the remainder of Part III. In Part II, we introduced the concepts of *games* and *strategies*, which are the main topics with which game theory is concerned. This branch of mathematics is a large and growing field with applications in a broad array of disciplines. However, since this is a poker book, we will not attempt to provide a comprehensive overview of the many interesting situations game theory addresses except for those that lend insight into poker. Instead, we will introduce the various concepts of game theory as they are needed for our discussion, without attempting to be rigorously formal.

In part II, we defined a game as containing the following elements:

- There are two or more players.
- At least one player has a choice of actions.
- The game has a set of outcomes for each player.
- The outcomes depend on the choices of actions by the players.

We call the choice of actions a *strategy*, and the outcomes *payoffs*. Each action that a player may take is called a *strategic option*; the specification of all the choices he makes at each possible decision point is the strategy.

As a simple introduction to the structure of games, let us consider one of the simplest games, Odds and Evens. Two players play. Each conceals either a penny or no penny in his hand, and the hands are revealed simultaneously. By prior agreement, one player (Player A) win 1 unit if

the total number of pennies in the hands is even, and the other player (Player B) wins 1 unit if the total number of pennies is odd.

The following is what is called a *payoff matrix* for this simple game. Player B's strategies are listed left to right and Player A's from top to bottom. Once the strategies have been revealed, looking up their intersection in the chart gives the payoff due to each player (A's payoff, B's payoff).

Example 10.1 - Odds and Evens

	Player B	
Player A	0 pennies	1 penny
0 pennies	(1, - 1)	(-1, 1)
1 penny	(-1, 1)	(1, - 1)

There are a number of properties of interest that games can have. Some games, such as Odds and Evens or rake-free poker, are what are called *constant-sum* games. The total of all the payoffs to the various players must total a constant. In poker, this constant happens to be zero, and it is common to speak of *zero-sum* games as well. All constant-sum games can be converted to zero sums by simply scaling the payoffs appropriately.

A second important property of games is whether the play is *sequential or simultaneous*. In a game like Odds and Evens, for example, the play is simultaneous. In a game like checkers, for example, play is sequential; first one player plays, then the other. In some sense, we can convert all games to simultaneous play by simply requiring all players to specify their complete strategy (including all possible game trees) beforehand. However, the distinction between simultaneous and sequential play is still useful because of what it tells us about mixed strategies (more in a moment).

Another important property of games is whether there is *hidden information*. Hidden information is information available to one player but not the other. First consider a game such as chess, in this game, the entire game situation is available to either player at any time, and the game is completely deterministic; that is, a player or computer with enough calculating power could examine every possible game tree and find the best move. This is an example of a game with no hidden information. Backgammon is an example of a game where there is no hidden information, but the game is not deterministic. On each roll of the dice the game tree forks, but the mechanics of this probabilistic selection process are known to each player. So in this case, even though perfect play cannot guarantee victory for either player because of the random element, there is no hidden information; that is, there is no information that one player has that the other does not.

Poker, on the other hand, is a game that relies on hidden information. Each player knows his own hand, and the knowledge of his opponent's hands would be a great boon to him. As a result, the task of poker is to extract value by taking advantage of the hidden information. From a game theory standpoint, the most important thing about games with hidden information is that their optimal strategies can contain *mixed strategies*. Mixed strategies are situations where strategies contain two or more actions in a single situation. For example, "raise half the time, limp- reraise half the time with aces in a full game" is an example of a mixed strategy. Only games with either hidden information or sequential play can contain mixed optimal strategies.

In using game theory to study poker, **zero-sum two-player games** (such as Odds and Evens) are normally the most important and most interesting. In zero-sum two-player games, only two players have payoffs and those payoffs always sum to zero; that is, whatever player A loses, player B wins, and vice versa. The comments that follow about "optimal strategies" apply only to zero-sum two-player games.

We can discuss **strategy pairs**; that is, a combination of a strategy for the first player and a strategy for the second player. For a zero-sum two-player game, a strategy pair is **optimal** if neither player can improve his expectation by unilaterally changing his strategy. A strategy for one player is optimal if it is part of any optimal strategy pair.

Imagine playing against a super-opponent. We call this super-opponent the **nemesis**. The nemesis always knows your strategy and always plays the maximally exploitive strategy against it. If you change your strategy, the nemesis changes instantly to counter, always playing the strategy that maximally exploits yours. An optimal strategy is the strategy that has maximum EV against the nemesis. Another way of stating this is:

An optimal strategy pair consists of two strategies that maximally exploit each other.

Please note that the term "optimal" in this context is carefully defined - many disciplines use this term, sometimes with varying definitions or meaning. We will use the term only in the narrow sense given above and only in reference to strategies.

In both zero-sum two-player and other types of games, either variable-sum or multi-player, we can still satisfy the condition that no player be able to increase his expectation by acting unilaterally. Strategy sets for each player that satisfy this condition are called **Nash equilibria**. It is proven that all multiplayer games with finite payout matrices have at least one such equilibrium; some games have multiple equilibria. In fact, some games have multiple equilibria that makes them difficult to analyze. We will revisit this in Part V.

Mathematics tells us that the optimal strategies to a **zero-sum two-player game** have the following properties:

- As long as mixed strategies are allowed (i.e., each player can use strategies such as "do X 60% of the time and Y 40%"), optimal strategies always exist.
- As a corollary to this, if an optimal strategy contains a mixed strategy, then the expectation of each strategic alternative must be equal against the opponent's optimal strategy. Thus, optimal strategies in poker do not contain "loss leaders" or other plays that sacrifice immediate expectation for deceptive purposes. If a hand is played in different ways, then each way of playing the hand will have the same expectation. If this were not the case, then the player could simply shift all the hands from the option with lower expectation to the option with higher expectation and unilaterally improve.

Frequently, especially in very simple games, the optimal strategy is simply the one that guarantees zero EV. This is because simple games are frequently totally symmetric; in this case the optimal strategy must yield zero for the player employing it, as his opponent could simply play the optimal strategy in response.

Looking back at Odds and Evens, it is clear that Player A's objective is to match Player B's strategy. The two could play this game as an exploitive guessing game and try to outthink each other; in that case, the player better at outguessing his opponent would have an edge. Another option is available, however. Suppose that B felt he was inferior at the guessing game to A. B

could instead try to play optimally. One way of doing this is to try to find the nemesis strategy for any strategy he himself utilizes and then maximize his equity.

B can play any mixed strategy which consists of playing 0 pennies X% of the time and 1 penny (1-X)% of the time. Note that pure strategies are simply a mixed strategy where one of the options is given 100% weight.

We can directly calculate the expectation of A's counter-strategies, but we should know from our work on exploitive play that the nemesis will best exploit B by playing a pure strategy. If B plays 0 pennies more than 50%, the nemesis will play 0 pennies all the time. If B plays 1 penny more than 50%, the nemesis will play 1 penny all the time.

From Equation 1.11, B's EV for 0 pennies > 0.5 is:

$$\langle B, x > 0.5 \rangle = (-1)(x) + (1)(1-x)$$

$$\langle B, x > 0.5 \rangle = 1 - 2x$$

His EV for 1 penny > 0.5 is:

$$\langle B, x < 0.5 \rangle = (-1)(1 - x) + (1)(x)$$

$$\langle B, x < 0.5 \rangle = 2x - 1$$

We can see that in both of these cases, the expectation of the strategy is negative. When $x > 0.5$, $1 - 2x$ is negative, while when $x < 0.5$, $2x - 1$ is negative.

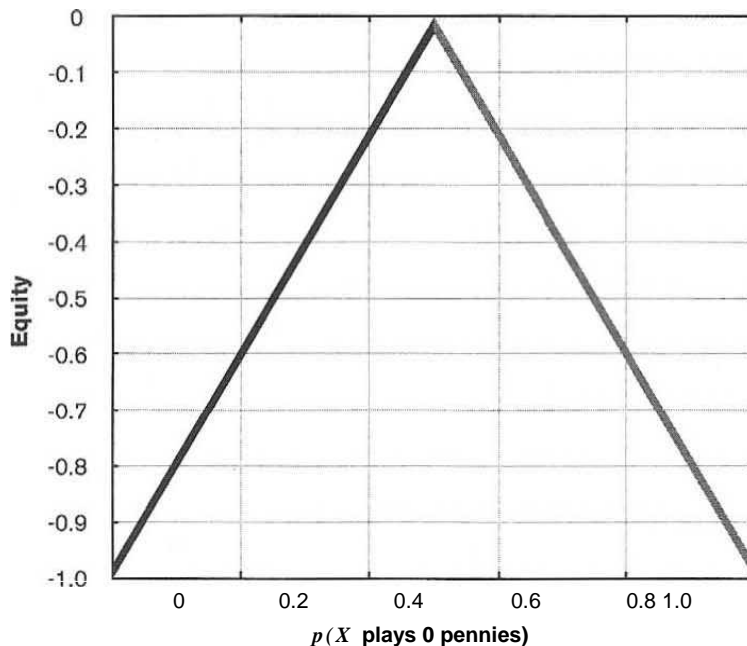


Figure 10.1. Odds and Evens equity vs. nemesis

At precisely $x = 0.5$, the nemesis can do anything and achieve the same equity.

$$\langle \text{nemesis, 0 pennies} \rangle = (-1)(0.5) + (1)(0.5)$$

$$\langle \text{nemesis, 0 pennies} \rangle = 0$$

$$\langle \text{nemesis, 1 penny} \rangle = (-1)(0.5) + (1)(0.5) = 0$$

$$\langle \text{nemesis, 1 penny} \rangle = 0$$

This is the highest expectation B can have against the nemesis. Therefore, $x = 0.5$ is B's optimal strategy. B can guarantee himself zero expectation by following a strategy where he randomly selects between 0 pennies (50%) and 1 penny (50%). To do this, he could use any source of external randomness; a die roll, another coin flip, or perhaps quantum measurements ■i radioactive decay, if he were sufficiently motivated to be unexploitable.

Example 10.2 - Roshambo

A second, slightly more complicated game that illustrates the same concept is the game of Roshambo (also known as Rock, Paper, Scissors). In this game, the players choose among three options: rock, paper, or scission, and the payoff matrix is as follows:

	Player B		
Player A	Rock	Paper	Scissors
Rock	(0, 0)	(-1, 1)	(1, -1)
Paper	(1, -1)	(0, 0)	(-1, 1)
Scissors	(-1,1)	(1, -1)	(0, 0)

In general, throughout this book we will be referring to zero-sum games. In cases where we have games that are non-zero sum, we will denote it explicitly. The following matrix is equivalent to the above, except that only A's outcomes are in the matrix cells.

	Player B		
Player A	Rock	Paper	Scissors
Rock	0	-1	+1
Paper	+1	0	-1
Scissors	-1	+1	0

If B wants to ensure 0 EV, he can play the strategy $\{1/3, 1/3, 1/3\}$ a similar way to the calculation in Odds and Evens, no matter what option A chooses, he will win 1 unit $1/3$ of the time, tie $1/3$ of the time, and lose 1 unit $1/3$ of the time. Any other choice on the part of B will lead to the nemesis exploiting him by playing a pure strategy in response. So this is an optimal strategy for this game.

What occurs frequently in these types of games is a situation where the nemesis has two or more possible exploitive strategies to employ. B selects his strategy, and the nemesis selects his exploitive strategy in response. At some strategy S that B chooses, the nemesis strategy is *indifferent* between his exploitive strategy choices; that is, it makes no difference to his expectation which one he chooses. This *indifference threshold* is the optimal strategy for B. The reason that this concept is so powerful is that it enables us to ensure that the strategies we find cannot be improved by moving them a tiny bit in a different direction. Remember that one of the conditions for a strategy being optimal is that neither player be able to improve by changing his strategy unilaterally. When a player is indifferent between strategy choices, this prevents this condition from being violated.

To see how this allows us to more easily find optimal strategies, let us consider a slightly modified version of Roshambo, where there is a bonus for winning with scissors. We'll call this game Roshambo-S, and its payoff matrix is as follows:

Example 10.3 - Roshambo - S

	Player B		
Player A	Rock	Paper	Scissors
Rock	0	-1	+1
Paper	+1	0	-2
Scissors	-1	+2	0

Now the strategy $\{1/3, 1/3, 1/3\}$ is easily foiled by the opponent playing $\{0, 0, 1\}$, in which case the opponent wins $1/3$ of a unit overall. However, A can again reduce the game to 0 EV by attempting to make his opponent indifferent between the various choices.

Say A's strategy will be $\{a, b, c\}$. Then A wants to make B's EV from the various strategy options equal. Using Equation 1.11:

$$\langle B, \text{rock} \rangle = (0)(a) + (-1)(b) + (1)(c)$$

$$\langle B, \text{rock} \rangle = c - b$$

$$\langle B, \text{paper} \rangle = (1)(a) + (0)(b) + (-2)(c)$$

$$\langle B, \text{paper} \rangle = a - 2c$$

$$\langle B, \text{scissors} \rangle = (-1)(a) + (2)(b) + (0)(c)$$

$$\langle B, \text{scissors} \rangle = 2b - a$$

Setting these three things equal and solving the resultant system of equations, we find that $a = 2b = 2c$. This indicates a strategy of $\{1/2, 1/4, 1/4\}$. We can verify that this strategy is optimal indirectly. First of all, the game is symmetrical, so any strategy B can utilize, A can utilize also. Since the payoffs sum to zero, no strategy can have positive expectation against the nemesis. If such a strategy existed, then both players would play it and the zero-sum nature of the game would be violated. So the strategy must be optimal.

Note that increasing the amount gained by winning with scissors shifted the optimal strategy to play less scissors. We often find this type of defensive maneuver in situations like this. When the rules change to benefit some specific type of action, optimal strategies often move toward counter-acting the benefited strategy instead of utilizing it. We will see that in poker, this manifests itself in that when the pot is larger, we bluff less, because a successful bluff is more profitable.

In the cases we have considered so far, all the strategy options have been equally "powerful," so to speak. Each one represented a key part of a strategy or counter-strategy, and no option was so weak that it should never be played. However, this does happen; very frequently, in fact, in more complex games. A strategy S is said to be **dominated** if there is a strategy S' such that $\langle S' \rangle \geq \langle S \rangle$ against all opponent strategies and $\langle S' \rangle > \langle S \rangle$ for at least one opponent strategy. S' is said to **dominate** S . (This term should not be confused with the flop game term "domination," which refers to situations where one player is reduced to three or less outs.)

This should make intuitive sense; if we have **1** strategic option, but there's some other option that does at least the same and sometimes better against all the opponent's strategies, we should always pick the better option. We can see this by looking at another variant of Roshambo. This time, instead of changing the payout matrix, we include an additional strategy option, "flower." Flower loses to each of rock and scissors and ties with paper and other flowers.

Example 10.4 - Roshambo - F

	Player B			
Player A	Rock	Paper	Scissors	Flowers
Rock	0	-1	+1	+1
Paper	+1	0	-1	0
Scissors	-1	+1	0	+1
Flowers	-1	0	-1	0

It should be clear that no strategy that includes flower can be as good as the same strategy that substitutes paper for flower in the frequencies. Therefore, flower is dominated by paper. We also have a second terminology, *strictly dominated*, which means that a strategy S' performs better than S (not equal) against all possible opponent strategies. This often occurs when all strategic options are gone: a poker example of this is when an opponent pushes all-in preflop in a cash game (with blinds in the pot) and you are the last remaining player in the hand with AA. No matter what your opponent's strategy is, you have higher expectation from calling than from folding. Hence calling strictly dominates folding. It should be obvious that an optimal strategy cannot contain strictly dominated options: if it did, then the player could simply substitute the higher-expectation option and unilaterally improve.

Strategic options are usually valuable; in fact, we can go so far as to say that in zero-sum two - player games, *strategic options have non-negative value*. Adding an additional strategic option to one player's choices in a game cannot reduce his expectation in the game, because he could simply pretend that the option was not added and play his old strategy to get the old expectation. However, often he can incorporate the new option into his strategy to gain value.

One other term we frequently use to describe strategies is *co-optimal*. Sometimes there can be many strategies that are part of an optimal strategy pair. This can include dominated strategies in some cases. This may be surprising to some readers, but an example should clarify. Suppose that to Roshambo-F, we add yet another strategic option: "safety scissors." (Safety scissors have blunt tips and cut paper just as well as regular scissors, but other materials more poorly). Safety scissors plays just as regular scissors (i.e., loses to rock and cuts paper), but loses to flower. So in this variant, playing safety scissors is always dominated by playing regular scissors. However, strategies involving safety scissors can still be co-optimal, because they can be part of an optimal strategy pair. This is because optimal strategy pairs will never contain flower; hence, safety scissors are as good as regular scissors.

When we search for optimal strategies in poker, however, we are primarily concerned with finding undominated strategies, because picking up value from opponent's mistakes is important and occurs frequently.

When there are dominated strategic options present, we can use a recursive process to simplify the game. In the above example, we begin with the full payoff matrix:

	Player B				
Player A	Rock	Paper	Scissors	Flower	Safety
Rock	0	-1	+1	+1	-1
Paper	+1	0	-1	0	-1
Scissors	-1	+1	0	+1	0

Flower	-1	0	-1	0	+1
Safety	-1	+1	0	-1	0

We can identify safety as a dominated strategic option by comparing the payoff values to the payoff values of scissors. In each case, the payoff value for scissors is greater than or equal to the payoff value for safety scissors. As a result, safety scissors is dominated by scissors. When we are searching for the optimal strategy, we can then reduce the game to a subgame where neither player can play safety scissors:

	Player B			
Player A	Rock	Paper	Scissors	Flowers
Rock	0	-1	+1	+1
Paper	+1	0	-1	0
Scissors	-1	+1	0	+1
Flowers	-1	0	-1	0

This matrix is, of course, Roshambo - F. We can now apply this process again, identifying Flower as the dominated strategy to remove:

	Player B		
Player A	Rock	Paper	Scissors
Rock	0	-1	+1
Paper	+1	0	-1
Scissors	-1	+1	0

And we have returned to basic Roshambo. Thus, the solution to the expanded game, including flower and safety scissors, is the same as the solution to this last sub-game. This recursive process can be used to simplify many games. This is often helpful because it reduces the number of strategic options that must be considered. To summarize:

A game G can be reduced to a sub-game G' by removing dominated strategic options from both sides. An optimal strategy pair for G' will be an optimal strategy pair for G.

The last simple game we will look at in this chapter, before going on to games directly related to poker, is called Cops and Robbers. In this game, one player is the "Cop" and the other player is the "Robber." The cop has two strategy options: to patrol or not to patrol. The robber has two strategy options as well: to attempt to rob a store, or not to. If the robber decides to rob, he wins 1 unit if the cop is not patrolling, but loses 1 unit if the cop patrols (because he gets arrested!). If the robber stays home, the cop loses 1 unit (resources, etc.) by patrolling, but nothing if he doesn't patrol. And because the game is zero-sum, the robber gains that unit when he chooses not to rob and the cop patrols.

The payoff matrix for this game is as follows:

Example 10.5 - Cops and Robbers

	Robber	
Cop	Rob	Don't Rob
Patrol	(1, -1)	(-1,1)
Don't Patrol	(-1, 1)	(0, 0)

To find the optimal strategies to this game, we pick a player. Consider the cop's strategies. As before the cop will choose a percentage x of the time to patrol, and the rest of the time he will not. From Equation 1.11:

$$\langle \text{robber, rob} \rangle = (-1)(x) + (1)(1 - x)$$

$$\langle \text{robber, rob} \rangle = 1 - 2x$$

$$\langle \text{robber, don't} \rangle = (1)(x) + 0(1 - x)$$

$$\langle \text{robber, don't} \rangle = x$$

The robber is indifferent between his choices when these two values are equal.

$$1 - 2x = x$$

$$x = 1/3$$

So the cop's optimal strategy is to patrol $1/3$ of the time.

Next we consider the robber. If he chooses a percentage x of the time to rob, the cop's expectations are as follows:

$$\langle \text{cop, patrol} \rangle = (1)(x) + (-1)(1 - x)$$

$$\langle \text{cop, patrol} \rangle = 2x - 1$$

$$\langle \text{cop, don't} \rangle = (-1)(x) + 0(1 - x)$$

$$\langle \text{cop, don't} \rangle = -x$$

The cop is indifferent between his choices when these values are equal.

$$2x - 1 = -x$$

$$x = 1/3$$

So the robber's optimal strategy is to rob $1/3$ of the time.

These two values (cop patrols $1/3$, and the robber robs $1/3$) are the optimal strategies for this game. We consider this game as an example of optimal strategies leading to indifference. Mixed strategies will occur if both sides have the ability to exploit each other's pure strategies. In this game, if the robber plays a pure strategy of always robbing, the cop can play a pure strategy of always patrolling. But if the cop does this, the robber can switch to a strategy of always staying home, and so on. These oscillating exploitive strategies tell us that the optimal strategies will be mixed.

This is true generally for a zero-sum two-player game; suppose two players X and Y are formulating their strategies for a game. X utilizes the pure strategy A . Y then exploits X by playing the pure strategy B . But X in turn exploits B by playing C . Y then exploits C by playing

D. *X*'s best response to *D* is *A* again. This recursive exploitation is a sign that the strategies are going to be mixed between *A* and *C* for *X*, and between *B* and *D* for *Y*.

If we know what components are to be mixed in the optimal strategies, we can easily solve the game systematically by writing and solving the equations that make both sides indifferent to the various mixed options. If we don't know which components are mixed - this occurs in different types of games, we often have to guess at the structure of the solution. Our guesses are sometimes called *parameterizations*, a topic we will revisit in discussing the solutions to [0,1] games.

In this game, however, since both the cop and the robber will be utilizing mixed strategies, we know that their expectation from each option will be equal. By setting the expectations equal to each other, we use this fact to find their strategies.

This game may seem simple and fairly unrelated to poker; but as we shall soon see, this game shares a structure and important features with a very simple poker game.

Key Concepts

- The optimal strategies to a zero-sum two-player game have the following properties: 1) As long as mixed strategies are allowed optimal strategies always exist. 2) If an optimal strategy contains a mixed strategy, then the expectation of each strategic alternative must be equal against the opponent's optimal strategy.
- If a strategy-pair is optimal, neither player can improve his expectation unilaterally by changing strategies.
- An optimal strategy pair consists of two strategies that maximally exploit each other.
- Optimal strategies maximize expectation against the nemesis.
- Optimal strategies do not contain strictly dominated alternatives.
- When the players have oscillating exploitive pure strategies - that is, if *X* exploits *Y* by playing a pure strategy, *Y* can exploit *X*, after which *X* exploits *Y* again, and so on...then *X* and *Y*'s strategies will be mixed.
- If a player employs a mixed strategy at any point, both of the mixed strategic options must have equal expectation against the opponent's optimal strategy.
- We can obtain simpler games that have equivalent optimal strategies by recursively removing dominated strategies from a larger game.
- In both zero-sum two-player and other types of games, strategy sets for each player that satisfy the condition that no player be able to increase his expectation by acting unilaterally are called Nash equilibria. All multiplayer games with finite payout matrices have at least one such equilibrium; some games have multiple equilibria.

Chapter 11

One Side of the Street: Half-Street Games

We begin our investigations into poker toy games by considering a class of simplified games, which we call *half-street games*. Half-street games have the following characteristics:

- The first player (conventionally called X) checks in the dark.
- The second player (conventionally called Y) then has the option to either check or may bet some amount according to the rules of the game.
- If Y bets, X always has the option to call, in which case there is a showdown, and may additionally have the option to fold (but not raise). If Y checks, there is a showdown.

We will make reference to the *value* of the game. This is the expectation for the second player, Y, assuming both sides play optimally. We will also only consider money that changes hands as a result of the betting in our game. This value is the *ex-showdown* value. This can include bets and calls that X and Y make on the streets of our game. It can also include the swing of a pot that moves from one player to the other as the result of a successful bluff.

The reason that we use ex-showdown value in studying these games is that we're not particularly concerned with the value to the player of action that occurred before our game. We're trying to capture the value of the betting that takes place on the half-street we are examining. The fact that the pot is five or ten or eighteen bets isn't important except as it relates to the way that the action plays in our game. We should also note that Y can't have a negative value in a half-street game, since he can simply check behind and achieve a value of zero.

Half-street toy games have at least one feature that is exceptionally valuable; we can solve them - in fact, we can solve just about any half-street game - as the complexity remains manageable throughout. As we will see, multi-street games often turn out to be intractable, or we can only obtain approximate answers. Being able to solve a game is of great benefit, as we can use the solution to gain insight into more difficult problems.

The first game we consider is a quite simple game involving a concept called *clairvoyance* - throughout Part III, players will be referred to as *partially clairvoyant* if they can see their opponents' cards (but not cards to come, if there are any), and simply *clairvoyant* if they can see both their opponent's cards and the cards to come. In this example, there are no cards to come, so either definition serves.

Example 11.1 - The Clairvoyance Game

One half-street.

Pot size of P bets.

Limit betting.

Y is clairvoyant.

Y's concealed hand is drawn randomly from a distribution that is one-half hands that beat X's hand and one-half hands that do not.

In this game, X and Y each have only one decision to make. This is a common element of all limit half-street games; Y must decide on a range of hands to bet, while X must decide on a range of hands with which to call. X's "range" in this game is going to consist of all the same hand, since he is only dealt one hand. This doesn't pose a problem, however - betting and calling distributions can just as easily contain mixed strategies with particular hands as they can pure

strategies.

We can immediately see that Y has the informational advantage in this game; he can value bet all his winning hands (call these hands *the nuts*) and bluff some of his losing hands (call these *dead hands*) with perfect accuracy.

The ex-showdown payoff matrix for this game is as follows:

	Player X	Check-Call	Check-Fold
Player Y			
Nuts	Bet	+1	0
	Check	0	0
Bluff	Bet	-1	+P
	Check	0	0

In finding the optimal strategies for this game, we should keep in mind some of the principles from the last chapter. First, we should consider pure strategies. Presume that Y value bets all his strong hands but checks all his weak ones. Then X can exploit him by folding all the time. If X does this, however, Y can switch to betting all his hands (value bets and bluffs). If Y does that, X switches to calling all the time, in which case Y can exploit him by value betting all his strong hands and checking weak ones.

So we have the familiar pattern of oscillating pure strategies; this indicates that the optimal strategies will be mixed. There are two strategy choices - one for X (how often to call a bet from Y), and one for Y (how often to bluff). Y should always bet his nut hands because this option dominates checking. X will call with some fraction of his hands. We will call this fraction *c*. Likewise, Y will bluff with some fraction of his dead hands. We will call this fraction *b*. If we find values for *b* and *c*, then we will have the optimal strategies for the game.

We can consider each of these decisions in turn.

First we consider how often X must call. When X plays his optimal strategy, Y will be indifferent to bluffing (that is, checking and bluffing his dead hands will have the same expectation). When bluffing, Y wins *p* bets by bluffing successfully (if X folds) and loses one bet when X calls. Recalling that *c* is the calling frequency, and $1 - c$ is the frequency that X folds, we have:

$$\begin{aligned} (\text{pot size}) (\text{frequency X folds}) &= (\text{bluff bet})(\text{frequency X calls}) \\ P(1 - c) &= c \\ c &= P/(P+1) \end{aligned}$$

Notice that as the pot size grows, the calling ratio grows. This is in some sense an extension of the pot odds principle; the more there is in the pot, the more often X must call to keep Y from bluffing.

Y, likewise, must bluff often enough to make X indifferent to calling or folding. When calling, X loses 1 bet by calling a value bet, and gains $P + 1$ bets by calling a bluff. If *b* is the ratio of bluffs to bets, then we have:

$$\begin{aligned} 1 &= (P + 1)b \\ b &= 1/(P + 1) \end{aligned}$$

This value $1/(P + 1)$ is an extremely important value in analyzing poker. It is so important that we

assign it its own Greek letter, α (alpha).

$$\alpha = 1/(P+1) \text{ (limit cases)} \tag{11.1}$$

α represents two things in this game. First, X must call enough to make Y indifferent to bluffing and checking his weak hands. X's calling frequency is $P/(P+1)$, or $1 - \alpha$. Therefore α is X's **folding frequency** when faced with a bet from Y. In addition, Y bluffs with α of his dead hands. Since he will bet with 100% of his nut hands, this makes α equal to the **ratio of bluffs to value bets** to keep the opponent indifferent between betting and calling.

Likewise, the ratio $1 - \alpha$ is of some interest:

$$\begin{aligned} 1 - \alpha &= 1 - 1/(P+1) \\ 1 - \alpha &= P/(P+1) \quad \text{(limit cases)} \end{aligned} \tag{11.2}$$

Even though we have only introduced cases with limit betting, we can generalize the expression of α to include games with different bet sizes:

For games with variable bet sizes:

$$\alpha = s/(1+s) \text{ (variable bet sizes)} \tag{11.3}$$

where s is the bet size in pots. (For example, betting \$50 into a \$100 pot makes $s = 1/2$)

Notice that here, as the pot size grows, α , the bluffing ratio becomes smaller: that is. Y bluffs less often into larger pots. This may seem counter-intuitive to you as there is more to gain by successfully bluffing into a large pot, but this is an important principle of optimal play:

Bluffing in optimal poker play is often not a profitable play in and of itself. Instead, the combination of bluffing and value betting is designed to ensure the optimal strategy gains value no matter how the opponent responds. Opponents who fold too often surrender value to bluffs while opponents who call too often surrender value to value bets.

We can now state that the optimal strategies for this game are:

Solution:

Y bets all his nut hands and bluffs α of his dead hands, or $\alpha/2$ of his total hands.
X calls with $1 - \alpha$ of his hands total.

The ex-showdown expectation of the game to Y is as follows. He wins 1 bet when he value bets and gets called. Recall that he holds half nut hands and half dead hands; hence, his value betting frequency will be $1/2$. He wins the entire pot when he bluffs successfully, and loses 1 bet when he bluffs and gets called. We know that X will call frequently enough to make him indifferent to bluffing. His EV from bluffing is therefore zero and his expectation from the game is:

$$\begin{aligned} \langle Y \rangle &= (\text{frequency of value betting}) (\text{one bet}) (\text{frequency of getting called}) \\ \langle Y \rangle &= (1/2)(P/(P+1)) \\ \langle Y \rangle &= P/2(P+1) \end{aligned}$$

Notice that as the pot grows. X's disadvantage grows; he must pay off the second player's value bets more and more often, and the number of bluffs that Y is required to make in order to keep X

calling is smaller and smaller. Also note that Y cannot have expectation of a half bet, no matter how large the pot is - this is because Y would have to be called all the time when he value bet to have such a value.

The important features of the clairvoyance game are the existence of betting to bluffing and calling ratios, their dependence on the pot size, and the idea that Y's ability to bet is worth more as the pot grows. Changing the size of the pot and seeing how it affects strategy is a useful tool in understanding poker. We can even examine what happens when we make the pot extremely large, such that neither player will fold any hand with any value at all.

The [0,1] Distribution

We introduce here for the first time the ***[0,1] distribution***. Throughout Part III, we discuss various games using this distribution. In [0,1] games, instead of discrete cards, both players are "dealt" a uniformly random real number between 0 and 1. This means that each player has an equal chance of being dealt any number between 0 and 1. These numbers function for these players as poker hands. The convention that we use is that if there is a showdown, then the lowest hand wins (this simplifies calculations). That means that 0 and numbers approaching it are the strongest hands, while 1 and numbers approaching it are the weakest.

Strategies for this type of game have a different structure than the strategy for the clairvoyant game. The primary way in which they are different is that in most cases, mixing is not strictly required. Since there are an infinite number of different hands, mixing on any particular hand is irrelevant, and for any interval, we can play more effectively with pure strategies, separated by a ***threshold***. These are points along the [0,1] continuum that mark the boundaries between one action and another,

When solving [0,1] games, we will often employ the following technique:

- 1) Guess at the structure of the solution,
- 2) Solve the game as if the guessed structure were the correct one,
- 3) Verify that our guess was correct by showing that neither player can improve his or her strategy unilaterally.

We use the term ***parameterization*** to refer to a guess at the structure of the solution of a game. For example, one possible solution to a hypothetical [0,1] game might involve a player Y betting his best hands between 0 and y_1 , checking his intermediate hands between y_1 and y_0 , and betting his worst hands between y_0 and 1. A different parameterization might be that Y bets his intermediate strength hands while checking his best and worst hands. Many structures of a solution are possible; if we guess at a parameterization, we can then solve the game for that parameterization only.

Once we have identified a parameterization, we will find the optimal solution for that particular parameterization. We generate equations based on principles of being indifferent at all the thresholds. Instead of the indifference for hands that utilize mixed strategies (as we saw before), we find that at the thresholds between one action and the next, the player will be indifferent.

It may not be obvious why this is true; here is an intuitive argument. First, we know that as we move very small distances along the line (from one hand to a very slightly stronger hand) that our showdown equities are continuous - that is, the showdown equity of one hand is very close to the showdown equity of its close neighbors. If this is true, then the value of playing hands must be continuous also, since a hand's value will be composed of its equity when there's a showdown plus its equity when there is no showdown (which is constant). For example, the value of check-

calling with the hand 0.6 should be very close to the value of check-calling with the hand 0.60001. Now consider the threshold value. Suppose the equity of one option at the threshold is higher than the equity of the other option. Then in the region of the lower-valued option, we can find a very small region very close to the threshold where we could switch from the lower-valued threshold to the higher-valued one and gain equity. If we can do this, the strategy isn't optimal. Hence optimal strategies are indifferent at the thresholds.

In most of our parameterizations, there are a fixed number of thresholds (between different strategies). For each threshold, we write an equation that relates the strategy elements that force that threshold to be indifferent. We call these *indifference equations*. By solving these systems of equations, we can solve for each of the threshold values and find the optimal strategy for a given parameterization.

In some cases, a better strategy can be found that utilizes a different parameterization. Often, if we try to solve the equations generated from a particular parameterization that is incorrect, we encounter inconsistencies (impossible thresholds, for example). Therefore, it is necessary for us to confirm the strategies after solving for the solution to a particular parameterization in order to show that it is indeed the best one.

Example 11.2 - [0,1] Game #1

This very simple game consists of a single half-street of betting with no folding allowed, where player X is forced to check and must call Y's bet, if Y chooses to bet. For this type of game (where no folding is allowed) the pot size is irrelevant.

X has no decisions to make. Y's strategy consists of a single decision - to bet or not to bet? Y knows X's response will be to call any bets, so Y can simply bet all hands that have positive or zero expectation and check all hands with negative expectation.

We will frequently create tables such as the follow to show various types of outcomes as we have done at other points in the text.

Y's EV of betting a hand of value y is:

X's band	Result
[0, y]	1 (Y bets and X calls with a better hand)
[y, 1]	+1 (Y bets and X calls with a worse hand)

Recall that X's chance of having any hand is uniformly distributed, so the probability of each line of the table is equal to the size of the interval on which X holds that hand.

$$\langle Y, \text{bet} \rangle = p(\text{X has a better hand})(-1) + p(\text{X has a worse hand})(+1)$$

$$\langle Y, \text{bet} \rangle = (y - 0)(-1) + (1 - y)(1)$$

$$\langle Y, \text{bet} \rangle = 1 - 2y$$

Now we find all the hands where Y's EV is greater than 0:

$$1 - 2y >= 0$$

$$Y <= 1/2$$

So Y should bet the best half of his hands; X is forced to call. When X has a hand in the range $[0, 1/2]$ Y breaks even overall, and when X has a hand in $[1/2, 1]$, Y gains 1 bet.

Half the time, Y bets, and half that time, he gains a bet, making his overall equity from this game $\frac{1}{4}$.

Viewing this graphically:




Figure 11.1 Game #1--Half Street, no fold

The no-folding game is a very simple one, but it introduces the idea of $[0,1]$ games, as well as illustrating some key principles. The first of these principles is that if the opponent has no decisions to make, simply maximizing expectation against his static strategy is optimal. The second is that if you hold a distribution similar to your opponent's, you cannot be raised, and your opponent must always call, you should bet half your hands. We will see this in all types of cases where there is no more betting.

Example 11.3 - $[0,1]$ Game #2

In the preceding example, we made a special rule that does not normally apply to poker situations - that no folding was allowed. We found that in this "no fold" game there was no bluffing - the betting region for Y was simply his top half of hands by value. Poker is often perceived to be fundamentally a game of bluffing, and so we turn our attention to games where one or the other player may fold.

$[0,1]$ Game #2 is identical to Game #1 (Example 11.2), except that instead of being forced to call, X may instead fold. It is in this case that the pot size again becomes relevant, as X's decision about whether to call depends on the value he stands to gain by calling. We use the convention (as we generally will in limit poker) of calling the pot size P and the size of the bet 1 unit.

In Game #1, we expressed Y's strategy as a single threshold; however, as games become more complex, we will have multiple thresholds to consider. We will assign these values to variables in a consistent way:

- x_n as the threshold between strategies which put in the n th bet and strategies which only put in the $(n - 1)$ th bet for value. (When more bets and check-raising are allowed, we will see that the threshold x_2 , for example, is a threshold between checking and betting

initially.)

- x_0 as a specific threshold (the threshold between bluffing and checking).
- X_n^* as the threshold between calling the n th bet and folding to it.

By the first of these rules, the Game #1 strategy for y would be expressed as y_1 , since Y is betting hands above that for value. There is no y_0 value, or perhaps it can be thought of as being at 1. Since X was forced to call all the time, x_1^* was also 1.

In Game #2, X actually has a strategy, and it is made up of a single threshold value x_1^* , which separates the region where X calls from the region where X folds to a bet from Y . The calling region for X will be an interval $[0, x_1^*]$, and the folding region an interval $[x_1^*, 1]$.

We can show that Y plays a strategy that includes some bluffs as follows:

If X and Y are playing optimally, x_1^* will be chosen such that it maximizes expectation against Y 's strategy. Let's see what happens if Y uses a strategy similar to his strategy from Game #1 - using a betting region from 0 to some threshold y_1 and a checking region from y_1 to 1.

X 's best response is to call with hands that have positive equity against Y 's betting range. X 's calling threshold x_1^* will then be chosen so that X 's equity is positive.

X 's equity is determined by his calling threshold. If he calls at x , he will beat Y 's hand and win the pot when Y is between y_1 and x_1^* , and lose a bet when Y is between 0 and x_1^* . So his overall equity is:

$$\langle X \rangle = (\text{pot value when he wins})(\text{chance of winning}) - (\text{one bet})(\text{chance of losing})$$

$$\langle X \rangle = (P+1)(y_1 - x_1^*) - 1(x_1^*)$$

When this is positive, he will call.

$$(P+1)(y_1 - x_1^*) - 1(x_1^*) > 0$$

$$(P+1)(y_1) - (P+2)x_1^* > 0$$

$$x_1^* > y_1(P+1)/(P+2)$$

So X would call with a fraction of Y 's betting hands, such that the ratio of Y 's bets to X 's calls would be $(P+1)/(P+2)$,

Intuitively, this makes sense. Let's say that the pot is 1, so the bet is pot-sized. Then X would need to have a $1/3$ chance of having the best hand in order to call. The threshold such that he would have such a chance is $2/3$ of the way from 0 to y_1 .




Figure 11.2. Anon-optimal strategy for Y in [0,1] Game#2

However, if X played this way, Y could unilaterally improve his equity by checking hands from x_1 to y_1 and instead bluffing an equal amount of hands near 1. By doing this, he improves the equity of those bluffs because X will more often fold the winning hand against those bluffs than against the hands on x_1, y_1 .

Since Y can unilaterally improve his equity by changing his strategy, then his strategy cannot be optimal. But the assumption that led to this was that Y's strategy was optimal. So something must be wrong with our assumption that Y should only value bet. Y must have a bluffing region.

We can parameterize Y's strategy as follows:

Y bets in a region of strong hands from 0 to some threshold y_1 .

Y checks in an intermediate region from y_1 to some threshold y_0 .

Y bets weak hands in a region from y_0 to 1.

We also guess for this parameterization that x_1^* lies to the right of y_1 .




Figure 11.3. Optimal strategy for [0,1] Game #2

We know that if X and Y are playing optimal strategies, then at x_1^* , X will be indifferent between calling and folding.

The indifference at x_1^* (X indifferent between calling and folding):

Y's hand	$p(\text{Y's hand})$	$\langle \text{X, call} \rangle$	Product	$\langle \text{X, fold} \rangle$	Product
$[0, y_1]$	y_1	-1	$-y_1$	0	0
$[y_0, 1]$	$1 - y_0$	$P + 1$	$(P + 1)(1 - y_0)$	0	0
Total			$-y_1 + (P + 1)(1 - y_0)$		0

Note that we omit Y's hands on $[y_1, y_0]$ as with those hands he has not bet and so we are not faced with this decision. The sums of the "Product" columns must then be equal - these are the weighted average equities across all of Y's hands.

$$\begin{aligned}
 -y_1 + (P + 1)(1 - y_0) &= 0 \\
 y_1 &= (P + 1)(1 - y_0) \\
 1 - y_0 &= y_1 / (P + 1)
 \end{aligned}$$

$$1 - y_0 = \alpha y_1 \tag{11.4}$$

In our parameterization, $1 - y_0$ represents the length of the interval of bluffing hands, just as y_1 represents the number of hands that will be value bet. The relationship between these two quantities is the same as it was in the clairvoyant game - the betting to bluffing ratio is still α .

We can then look at the indifference at the Y's two thresholds.

The indifference at y_1 (Y indifferent between betting and checking):

X's hand	$p(\text{X's hand})$	$\langle \text{Y, bet} \rangle$	Product	$\langle \text{Y, check} \rangle$	Product
$[0, y_1]$	y_1	-1	$-y_1$	0	0
$[y_1, x_1^*]$	$x_1^* - y_1$	+1	$x_1^* - y_1$	0	0
$[x_1^*, 1]$	$1 - x_1^*$	0	0	0	0
Total			$x_1^* - 2y_1$		0

Setting the expected values equal:

$$x_I^* - 2y_I = 0$$

$$y_I = x_I^* / 2 \tag{11.5}$$

The indifference at y_0 (Y indifferent between betting and checking):

X's hand	p(X's hand)	<Y, bet>	Product	<Y, check>	Product
$[0, x_I^*]$	x_I^*	-1	$-x_I^*$	0	0
$[x_I^*, y_0]$	$y_0 - x_I^*$	+P	$(y_0 - x_I^*)P$	0	0
$[y_0, 1]$	$1 - x_I^*$	0	0	0	0
Total			$Py_0 - (P + 1)x_I^*$		0

$$Py_0 - (P + 1)x_I^* = 0$$

Recalling from Equation 11.3 that $1 - y_0 = \alpha y_I$, we have:

$$P(1 - \alpha y_I) = (P + 1)x_I^*$$

$$x_I^* = P/(P+1) (1 - \alpha y_I)$$

Remembering that $(1 - \alpha) = P/(P + 1)$, we have:

$$x_I^* = (1 - \alpha) (1 - \alpha y_I) \tag{11.6}$$

This result is of equal importance as it determines the calling ratio in continuous cases. To make Y indifferent to bluffing at y_0 , X must call with a fraction of his hands that can beat a bluff equal to $P/(P + 1)$, or $(1 - \alpha)$. Alternatively, we can say that X can fold α of his hands that can beat a bluff and that will make Y indifferent to bluffing.

Combining the three indifference equations, we can find a solution to the game.

$$x_I^* = (1 - \alpha) (1 - \alpha y_I) \tag{11.6}$$

$$y_I = x_I^* / 2 \tag{11.5}$$

$$x_I^* = 2y_I$$

$$2y_I = (1 - \alpha y_I) (1 - \alpha)$$

$$2y_I = 1 - \alpha y_I - \alpha + \alpha 2y_I$$

$$1 - \alpha = 2y_I + \alpha y_I - \alpha 2y_I$$

$$1 - \alpha = y_I(2 - \alpha)(\alpha + 1)$$

$$y_I = (1 - \alpha) / (2 - \alpha)(\alpha + 1)$$

Since $x_I^* = 2y_I$:

$$x_I^* = 2 (1 - \alpha) / (2 - \alpha)(\alpha + 1)$$

$$1 - y_0 = \alpha y_I$$

$$1 - y_0 = \alpha (1 - \alpha) / (2 - \alpha)(\alpha + 1)$$

This game was one of the games solved by John von Neumann and Oskar Morganstern in their seminal work *Game Theory and Economic Behavior* (1944).

We see in this game a number of important points. First, we see that X's ability to fold or call changes the basic structure of Y's strategy. In Game #1, Y simply bet his best hands for value. However, in Game #2, if Y employs such a strategy, X simply folds all his weaker hands, and Y never makes any money. Y's response is to begin to bluff his weakest hands, forcing X to call with hands between Y's value bets and Y's bluffs in order to keep Y from bluffing.

In this game, we also see α play a prominent role. Just as in the clairvoyant game, X folds α of his hands that can beat a bluff, making Y indifferent to bluffing, and Y makes his ratio of bluffs to bets α to make X indifferent to calling with hands between y_1 and y_0 .

A third important idea here is the following. What is the value of bluffing? In the clairvoyant game, the answer was zero; if Y checked, he simply lost the pot. In this game, however, the answer is not quite so simple. At y_0 , Y is indifferent between bluffing and checking - that is, the situation is the same as the clairvoyant game (from an ex-showdown perspective). But at, for example, 1, Y's equity from bluffing is the same as his equity at y_0 (since X will only call with hands above x_1^*). So Y is not indifferent to bluffing at 1 - in fact, for all the hands he bluffs which are worse than y_0 , he increases his equity by doing so!

In this chapter, we presented two very simple half-street games and a third that was more complex. In the clairvoyant game, we saw that the clairvoyant player bluffs and value bets hi a precise ratio optimally, while the player whose hand is exposed calls and folds hi a related ratio to make the clairvoyant player indifferent to bluffing. In [0,1] Game #1, we introduced the [0,1] distribution and saw that when one player has no strategic options, the other player simply maximizes his equity. In [0,1] Game #2, we saw a similar pattern of betting best hands and bluffing worst hands, and we saw the effect that bluffs have on the calling threshold, forcing X to call with hands worse than Y's worst value bet.

Key Concepts

- Value betting and bluffing in optimal strategies are directly related to each other in half-street games by the ratio $\alpha = 1/(P+1)$. By bluffing in this ratio, the bettor makes the caller indifferent to calling.
- α is also the proper folding ratio for X in half-street games. By folding this ratio, the caller makes the bettor indifferent to bluffing.
- When one player has no strategic options, the other player should simply maximize his expectation.
- Balancing value bets with bluffs properly ensures that a player extracts value from the opponent no matter how the opponent plays; if he folds too often, the bettor gains on his bluffs, and if he calls too often, the bettor gains on his value bets.
- Optimal strategies are not always indifferent to bluffing; only the threshold hands are truly indifferent. Often the weakest bluffs have positive value (compared to checking).

Chapter 12

Headsup With High Blinds: The Jam-or-Fold Game

In Chapter 11, we discussed three half-street toy poker games. For each of these games, we were able to find the optimal strategies - full solutions to the game. Solving toy games can be useful in terms of real poker because each of these games has a particular lesson that we might attempt to apply to real-life poker games. This is even more valuable because in general, real poker games are too complex to solve directly, even using computers.

In this chapter we will work through a toy game as a preliminary stepping stone, and then discuss and provide the computed solution to a game which is of real and immediate use in no-limit holdem games being played every day. The end product of this analysis is the solution to headsup no-limit holdem *where the first player must either jam (move all-in) or fold*.

We consider these games here in the midst of our half-street games even though they have a slightly different structure because they are more like half-street games than any other type. When we look at jam-or-fold games, instead of considering ex-showdown value, we will consider the total value of the game. By doing this, we create a situation where the first player to act has similar options to Y in the half-street games. He can fold (analogous to a check), which gives him value zero. Or he can bet the amount stipulated by the game, after which the other player may call or fold. Viewed in this light, jam-or-fold games are an extension of half-street ideas.

Before we begin discussion of this game, however, we invite the reader to test his intuition by considering the following question, the solution to which will be revealed through this analysis. Some readers who know something about the jam-or-fold game may already know the answer.

Two equally skilled players contest a headsup no-limit holdem freezeout where the button must either jam or fold. Each has sixteen chips to begin the hand. The blinds are one chip on the button and two chips in the big blind. With what percentage of hands should the button, playing optimally, move all-in?

We now consider our first toy game, which is a simple game of jam-or-fold with static hand values based on the $[0,1]$ distribution.

Example 12.1 - $[0,1]$ Jam-or-Fold Game #1

Both players have equal stacks of S units.

Each player receives a number uniformly from $[0,1]$.

The player without the button (call him the *defender*, X) posts a blind of 1 unit.

The button (call him the *attacker*, Y) posts a blind of 0.5 units and acts first.

The attacker may either raise all-in to S units (*jam*), or fold, surrendering the small blind.

If the attacker goes all-in, the defender may call or fold.

If there is a showdown, then the hand with the lower number wins the pot.

The first thing to notice about this game is that the attacker has just two alternatives -jamming and folding. When a player has two strategic options, where one of the options involves putting money in the pot and one of them is folding, putting money in the pot with any hand that is worse than a hand you would fold is always dominated. Because of this, the first player's strategy will consist of just two regions - strong hands with which he will jam, and weak hands with which he will fold, in the same way, the second player will call with his strongest hands and fold his weakest hands. Unlike other situations where we bet our strong hands along with the weakest hands, it does us no good here. In fact, if our jamming frequency is x , jamming with the strongest

x hands dominates all other sets of jamming hands.

Our two players' strategies can be expressed directly as two numbers, the attacker's strategy y and the defender's strategy x . These numbers are thresholds; with hands better than y , the attacker jams, and with hands better than x , the defender calls.

Additionally, we know that y and x for the optimal strategies are *indifference points*. This means that when Y (the attacker) holds the exact hand y , he is indifferent to jamming or folding. Likewise, when X (the defender) holds the exact hand x , the defender is indifferent to calling a jam or folding. We also know that X will never call with a hand worse than y , Y's jamming threshold - if he did, he could never win the pot. Using this information, we can solve the game directly using indifference equations, as in the last chapter. For this game, we'll deviate from our usual practice and calculate the total value of the game, including the blinds.

The expectation of folding for Y is simply $-1/2$ units (he loses his small blind).

$$\langle Y, \text{fold} \rangle = -1/2$$

Since X will never call with a hand worse than y , X always wins when he calls. Then the expectation of jamming with the threshold hand y is:

$$\langle Y, \text{jam} \mid y \rangle = (\text{defender calls}) (-S) + (\text{defender folds}) (+1)$$

For a uniform distribution such as the $[0,1]$ distribution, the probability of an interval being selected is equal to the length of that interval. Then the defender calls with probability x and folds with probability $(1 - x)$.

$$\begin{aligned} \langle Y, \text{jam} \mid -y \rangle &= (x)(-S) + (1 - x)(+1) \\ \langle Y, \text{jam} \mid y \rangle &= -xS + 1 - x \end{aligned}$$

We know also that at the threshold value, Y will be indifferent between jamming and folding. Hence the value of these two strategic options will be equal. Setting this equal to the value of the folding alternative:

$$\begin{aligned} -xS + 1 - x &= -1/2 \\ x(1+S) &= 3/2 \end{aligned}$$

$$x = 3/(2 + 2S) \tag{12.1}$$

The expectation of folding to a jam is -1 unit, while the expectation of calling a jam with the threshold hand x is:

$$\langle X, \text{call} \mid x \rangle = (\text{attacker has a better hand})(-S) + (\text{attacker has a worse hand})(+S)$$

The attacker jams a total of y hands (we know that $x < y$ because X will never call with a hand worse than y .) Of these hands, $(y - x)/y$ are worse than the calling threshold x and x/y are better. For example, suppose that Y jams with 50% of hands, while X calls with 30%. ($y = 0.5, x = 0.3$) Then $(y - x)/y = 2/5$ of Y's jamming hands would be worse than X's calling threshold, and the remaining $x/y = 3/5$ of Y's jamming hands would be able to beat it.

$$\begin{aligned} \langle X, \text{call} \mid x \rangle &= p(\text{Y wins})(-S) + p(\text{X wins})(+S) \\ \langle X, \text{call} \mid x \rangle &= (x/y)(-S) + [(y - x)/y](+S) \end{aligned}$$

$$\langle X, \text{fold} \rangle = -1$$

Setting these equal so that X will be indifferent between calling and folding, we have:

$$\begin{aligned} (S)(y-2x)/y &= -1 \\ Sy - 2Sx &= -y \\ S - S2x/y &= -1 \end{aligned}$$

$$y = S2x/(S + 1) \tag{12.2}$$

Remembering Equation 12.1:

$$x = 3/(2 + 2S)$$

Substituting for x gives us:

$$y = 3S/(1 + S)^2 \tag{12.3}$$

These two equations, then, are the optimal strategies for this game: the attacker jams with $3S/(1 + S)^2$ of the hands, and the defender calls with $3/(2 + 2S)$ of the hands. We can test these solutions intuitively to see if they make sense. For example, with stacks of 5 units, the attacker jams $(3(5)/(1 + 5)^2) = 15/36$ of the time, and the defender calls $(3/(2 + 2(5))) = 1/4$ of the time. With stacks of 10 units, the attacker jams $(3(10)/(1 + 10)^2) = 30/121$ of the time, while the defender calls only $3/(2 + 2(10)) = 3/22$ of the time. We can see that at the extreme, with stacks of, for example, a million, both players' frequencies are close to zero.

Is this game good for the attacker or for the defender? We can determine the value of the game to find out. Here we use the matrix method and look at the matchup possibilities.

X →	[0, x]	[x, y]	[y, 1]
Y			
[y, 1]	-0.5	-0.5	-0.5
[x, y]	-S	1	1
[0, x]	0	1	1

When the attacker folds, the defender gains 0.5 units. This happens $1 - y$ of the time.

When the attacker jams and the defender folds, the attacker gains 1 unit. This happens $y(1 - x)$ of the time.

When the attacker jams and the defender calls, there are two cases. The first case is when the two players are both within the caller's threshold. In this case, they break even, as each will have all numbers from zero to x with equal probability, and so their distributions are equal. The second case is when the attacker is above the defender's threshold. In that case, the defender wins S units. This occurs when the attacker jams with a hand worse than the defenders calling threshold and the defender has a calling hand, which has probability $(y - x)(x)$.

So we have:

$$\begin{aligned} \langle Y \rangle &= p(\text{Y folds})(-1/2) + p(\text{Y jams, X folds})(+1) + p(\text{Y jams, X calls and } x > y)(-S) + \\ &\quad p(\text{Y jams, X calls and } x < y)(0) \\ \langle Y \rangle &= -(1 - y)/2 + (y)(1 - x)(1) + (y - x)x(-S) \end{aligned}$$

$$\langle Y \rangle = -\frac{1}{2} + y \left[\frac{1}{2} + 1 - x - Sx \right] + Sx^2$$

$$\langle Y \rangle = -\frac{1}{2} + y \left[\frac{3}{2} - (1+S)x \right] + S \left(\frac{3}{2(1+S)} \right)^2$$

$$\langle Y \rangle = -\frac{1}{2} + y \left[\frac{3}{2} - \frac{3}{2} \right] + \frac{9S}{4(1+S)^2}$$

$$\langle Y \rangle = \frac{9S}{4(1+S)^2} - \frac{1}{2}$$

Using this valuation function, we can examine the value of the game as the stack size grows from 1 to infinity. With stack sizes from 1 to 2, the game favors tire attacker, while above 2 units, the game favors the defender. As the stacks go to infinity, the advantage of the defender simply goes to the 0.5 units he wins when the attacker folds, because the attacker must fold so often. So overall, this game is a pretty poor one for the attacker.




Figure 12.2. Jam or Fold Game #1

When we play poker, however, the best hand preflop doesn't always win. In fact, this is dramatically not the case. Even if the defender calls with a fairly strong distribution, the attacker will normally not have less than 33% of equity in the pot. And so our next toy game is a variation on the previous game, but where the worst hand preflop has a 33% chance of winning at showdown.

Example 12.2 - [0,1] Jam-or-Fold Game #2

Both players have equal stacks of S units.

Each player receives a number uniformly from $[0,1]$.

The player without the button (call him the *defender* or X) posts a blind of 1 unit.

The button (call him the *attacker* or Y) posts a blind of 0.5 units and acts first.

The attacker may either raise all-in to S units (*jam*), or fold, surrendering the small blind.

If the attacker goes all-in, the defender may call or fold.

If there is a showdown, then the hand with the lower number has $\frac{2}{3}$ equity in the pot.

This game is quite similar to the previous games, and we can solve it in much the same way, except for two important differences. In the previous game, when the stacks were very small (for example, at $S = 1$), the attacker still only jammed with 75% of his hands, while the defender only called with 75% of his. This because hands very close to 1 had virtually no value. But in this game, even the very worst hand has $\frac{1}{3}$ equity in the pot because of the rules of the game. So at stack size 1 the attacker will have higher equity from putting his chips in than from folding. Secondly, in the previous game the defender could not gain from calling with a hand worse than any hand with which the attacker would jam. However, in this game, this is no longer always the case because the defender always has at least $\frac{1}{3}$ of the pot.

We divide the game into two cases.

Case 1: $S < 3$.

When the stacks are less than 3 units, we can conclude something about the defender's strategy. If the attacker jams, the defender will be getting at least two to one immediately. Since the defender always has at least $\frac{1}{3}$ of the pot, he should always call - that is, $x = 1$. This creates a situation where the defender really has no strategic options, as folding is dominated by calling. Thus, the attacker should simply maximize his equity against the defender's static strategy.

$$\langle Y, \text{jam} \rangle = p(\text{attacker has best}) (\text{pot value}) (\text{pot equity}) + p(\text{attacker has worse}) (\text{pot value}) (\text{pot equity}) - (\text{cost of bet})$$

$$\langle Y, \text{jam} \rangle = p(\text{attacker has best})(2S)(\frac{2}{3}) + p(\text{attacker has worse}) (2S) (\frac{1}{3}) - (S)$$

The defender will call with all hands. Hence, if the attacker jams with y_I , the defender will have a hand better than y_I , with $p = y_I$, and a worse hand with $p = 1 - y_I$,

$$p\{\text{attacker has best hand}\} = 1 - y_I$$

$$p(\text{attacker has worst hand}) = y_I$$

$$\langle Y, \text{jam} \mid y_I \rangle = (1 - y_I)(2S)(\frac{2}{3}) + y_I (2S)(\frac{1}{3}) - (S)$$

$$\langle Y, \text{jam} \mid y_I \rangle = S/3 - 2Sy_I/3$$

$$\langle Y, \text{fold} \mid y_I \rangle = -\frac{1}{2}$$

Then the attacker should jam whenever:

$$S/3 - 2Sy_I/3 > -\frac{1}{2}$$

$$2S - 4Sy_I > -3$$

$$y_I < (2S + 3)/4S$$

Since he will jam whenever his hand value y_I is greater than this, his threshold value y is:

$$y = (2S + 3)/4S$$

The optimal strategies for the game when the stacks are less than three, then, is for the attacker to jam whenever this equation is satisfied and fold otherwise, and for the defender to call all the time. If we examine the play at different stack sizes, we can see that this makes sense. For example, at a stack size of less than 1.5 units, the attacker is now getting two to one to jam even if he will be called 100% of the time. So he should always jam below that threshold. As we can see by substituting 1.5 for S , this satisfies the equation for all hands the attacker can hold.

Likewise, at the stack size 3, we see that the attacker should jam with $\frac{3}{4}$ of his hands. If we calculate his equity from jamming with the hand 0.75, we find that it is - 0.5 units, just the same as if he had folded.

Case 2: $S > 3$.

At stack sizes greater than 3, however, the defender is no longer getting two to one on his calls. As a result, he cannot call with all hands and do better than folding. In fact, he can no longer call with hands worse than the attacker's jamming threshold. If he does, he will simply have one-third of the equity in the pot, and since the stacks are large enough, this will have worse expectation than folding. As a result, this game is fairly similar to the first jam-or-fold game we examined (Example 12.1), except that the best hand has just $\frac{2}{3}$ of the equity in the pot.

We can solve the game using the same indifferences from that game (with the modifications regarding the distribution of the pot). One identity we will use is that the if there is a showdown (i.e., a jam and a call), the player with the best hand wins $S/3$ units, while the player with the worst hand loses $S/3$ units.

$$\langle Y, \text{jam} \mid y \rangle = (\text{defender calls}) (-S/3) + (\text{defender folds}) (+1)$$

The defender calls with probability x and folds with probability $(1 - x)$.

$$\langle Y, \text{jam} \mid y \rangle = (x)(-S/3) + (1 - x)(+1)$$

$$\langle Y, \text{jam} \mid y \rangle = (-xS/3) + 1 - x \tag{12.4}$$

$$\langle Y, \text{fold} \rangle = -1/2$$

Setting this equal to make Y indifferent between jamming and folding:

$$-xS/3 + 1 - x = -1/2$$

$$xS + 3x = 9/2$$

$$x = 9/2(S + 3) \tag{12.5}$$

For X, the expectation of folding to a jam is 1 unit, while the expectation of calling a jam with the threshold hand x is:

As in Example 12.1, the attacker jams a total of y hands (we know that $x < y$ from previously). Of these hands, $(y - x)/y$ are worse than the calling threshold x and x/y are better.

$$\langle X, \text{call} \mid x \rangle = p(\text{attacker has better}) (-S/3) + p(\text{attacker has worse}) (+S/3)$$

$$\langle X, \text{call} \mid x \rangle = (-S/3)(x/y) + (+S/3)(y - x)/y$$

$$\langle X, \text{call} \mid x \rangle = (S/3)(y - 2x)/y$$

$$\langle X, \text{fold} \rangle = -1$$

Setting these equal to make X indifferent between folding to a jam and calling it:

$$(S/3)(y - 2x)/y = -1$$

$$S(y - 2x) = -3y$$

$$y = 2xS/(S + 3)$$

$$y = 9S/(S + 3)^2 \tag{12.6}$$

So the optimal strategy is for the attacker to jam with $9S/(S+3)^2$ of his hands, and the defender to defend with $9S/(S+3)$ of his hands.

Recall that in Example 12.1, the play was comparatively tight; that is, at stacks around 10, the attacker jammed about a fourth of the time. Contrast that with this game. Here the attacker jams with $(9(10)/(10+3)^2) = 90/169$, or more than half his hands at a stack size of 10! The defender then calls with $(9/(6 + 2(10)) = 9/26$ of his hands, or a little more than a third. This might seem quite loose. After all, can it really be right to jam ten times the blind more than half the time, when one is a 2:1 underdog when called?

t is. The reason this seems intuitively wrong is because we are used to considering situations where we "risk X to win Y," putting in a fixed amount and having our chance of succeeding (by the other player folding or winning a showdown) result in a win of a different fixed amount. By that logic, it does seem rather silly to risk ten to win one if the opponent is going to call a third of the time (roughly). However, this is not the case here. Recall that even when the attacker is called, his loss is simply $S/3$, not S . So in a sense, the attacker is putting in just three and a third to win one and a half (the dead money in the pot).

If we compare the solutions to Game #2 (Example 12.2) with $S > 3$ to the solution to Game #1 (Example 12.1), we can see this relationship more strongly:

Game 1:

$$x = 3/(2 + 2S)$$

$$y = 3S/(1 + S)^2$$

Game 2:

$$x = 9/(2S + 6)$$

$$y = 9S/(S + 3)^2$$

The solution to game 2 is simply the solution to game 1, but substituting $S/3$ for S . This is because the *effective stack size* of Game 2 is just $S/3$ - this is the amount that the better hand gains when there is a showdown. We can use this concept of effective stack size to calculate the value of the game, G , by simply substituting $S/3$ into the game value formula we found in Game #1.

$$\langle Y \rangle = -1/2 + 3/4 S/((S/3+1)^2)$$

$$G = -\frac{1}{2} + \frac{27S}{4(S + 3)^2}$$

However, this is only the game value for $S > 3$. There are two other regions of stack size that have different solutions. If $S < 3/2$ then both players play all hands and the game value is zero. In the region between $3/2$ and 3, then we found in Case 1 of the last example that:

$$y = \frac{2S+3}{4S} \text{ and } x = 1, \text{ so}$$

$$\langle Y \rangle = p(\text{Y folds})(-1/2) + p(\text{Y jams})p(x_1 > y)(S/3)$$

$$\langle Y \rangle = -\frac{1}{2}(1 - y) + y(1 - y)\frac{S}{3}$$

$$\langle Y \rangle = (1 - y) \left[-\frac{1}{2} + y\frac{S}{3} \right]$$

$$\langle Y \rangle = \frac{2S - 3}{4S} \left[-\frac{1}{2} + \frac{2S + 3}{12} \right]$$

$$\langle Y \rangle = \frac{(2S - 3)^2}{48S}$$

Note that the three game value regions are continuous, as we expect, that is:

$$\langle Y \rangle = 0, \text{ for } S \leq 3/2$$

$$\langle Y \rangle = (2S - 3)^2 / 48S, \text{ for } 3/2 \leq S \leq 3$$

$$\langle Y \rangle = -1/2 + (27S)/(4(S+3)^2), \text{ for } 3 \leq S$$

However, if we look at the threshold values x and y we notice something unusual. That is, for $S < 3$, $x = 1$, but suddenly when S becomes a little larger than 3, $x = 3/4$. That seems unusual. However, it shows that strategies are not continuous with respect to stack sizes or pot sizes.

These two toy games can help us understand how the fact that the best hand does not always win the pot affects play in jam-or-fold games. When the best hand did win the whole pot, proper play was relatively right and the defender had a significant advantage for most stacks. However, when the best hand only won $2/3$ of the time in Game #2 (Example 12.2), proper play became much looser and perhaps counter to our intuition, the attacker had an advantage all the way up to a stack size of six.

As we shall see in the next section, play in no-limit holdem is similarly loose and aggressive.

Example 12.3 - Jam-or-Fold No-limit Holdem

Both players have equal stacks of S units.

Each player receives a holdem hand.

The player without the button (call him the *defender*) posts a blind of 1 unit.

The button (call him the *attacker*) posts a blind of 0.5 units and acts first.

The attacker may either raise all-in to 5 units (*jam*), or fold, surrendering the small blind.

If the attacker goes all-in, the defender may call or fold.

If there is a showdown, then a holdem board is dealt and the best poker hand wins.

Here we have what might be considered a logical extension to the jam-or-fold games we have considered previously. However, holdem hands do not conform very nicely to the $[0,1]$ structure; there are no hands which have nearly zero preflop value, and just two that have even 80% against a random hand. Also, hand values are not transitive; this can be easily shown by the famous proposition bet where the shark offers a choice from among AKo, 22, and JTs, provided that he gets to choose once the sucker has chosen. We can, however, analyze this game directly by brute force. Before we do that, however, we think it interesting to examine the situations at the extremes (when the stacks are very large, and when the stacks are very small).

In the following discussion, we will make reference to a technique called fictitious play. We have already made informal use of this technique and it is of exceptional value in solving game theory problems with computers. Much has been written in academia about this topic, but we will summarize the basic idea here. Essentially, if we have a game and a "nemesis" strategy for the game (which given an opponent's strategy returns the maximally exploitive strategy), we can use these to find the optimal strategy.

The process is fairly straightforward. We begin with a candidate strategy for each player; this can be anything, although in most cases if we have a reasonable guess at strategy for both sides,

the process converges more quickly. Suppose that we have two players, A and B. We take a candidate strategy for both players. We then calculate the maximally exploitive strategy for B given that A plays his candidate strategy. We combine this maximally exploitively strategy with the initial candidate strategy by mixing the two together using a "mixing weight" constant. This mixing weight is some sequence that converges to zero but has enough terms that strategies can move as much as they need to, such as the harmonic sequence $1/n$.

The mixing process works like this. Suppose that m is the current mixing weight. For each possible decision point, we combine the old strategy S_{old} and the new maximally exploitive strategy S_{new} .

If we apply this process of fictitious play iteratively, it has been proven mathematically that our strategies will converge on the optimal strategy given enough iterations. Often "enough" is some achievable number, which makes this a very useful method for solving games with computers.

Very Large S

When the stacks are very, very large, both sides only play aces. The attacker cannot profitably jam any other hands because the defender can simply pick him off by calling with aces only. However, as we reduce the stack size incrementally, the attacker must add additional hands to his jamming distribution or miss out on additional blind-stealing equity. It should be clear that if the attacker does not add any additional hands, there is no incentive for the defender to call with anything but aces, since the defender is taking down 1.5 unit pots each time the attacker does not hold aces ($2^{20}/2^{21}$ of the time). Since calling with any non-AA hands when the attacker is jamming with aces loses a great deal, the defender is content to play this strategy until the attacker modifies his.

Above a certain stack size, however, the attacker simply cannot make money by jamming with additional hands, as he loses too much even when he's only called by Aces. You might think that the next hand to be added (once the stacks become a little smaller) would be KK, as this hand is the second-best hand; however, it turns out that this is not the case because at these high stacks, **card removal** is more important. Hands that contain aces, even though they do substantially worse against AA than do other pairs, reduce the probability that the other player holds aces by half and thus have higher equity than KK. The hand that performs the very best in jam-or-fold against an opponent who will call with only aces is ATs. We find this by simply evaluating every hand's equity against this strategy.

If the attacker jammed with ATs (the first hand to be added, as we shall soon see) at a stack size of 2000, his equity would be:

$$\begin{aligned} \langle \text{ATs, jam 2000} \rangle &= p(\text{defender folds}) (\text{pot}) + p(\text{defender calls}) p(\text{attacker wins pot}) \\ &\quad (\text{pot size}) - \text{cost of jam} \\ \langle \text{ATs, jam 2000} \rangle &= (1^{222}/1225) 1.5 + (3/1225) ((0.13336) (4000) - 2000) \\ \langle \text{ATs, jam 2000} \rangle &= - 2.09525 \end{aligned}$$

for a loss of approximately 2 units per jam. We can use this formula, however, to find the exact stack size, x , at which ATs becomes profitable to jam against a defender who will only call with aces:

$$\langle \text{ATs, jam } x \rangle = 1.5(1^{222}/1225) + (3/1225) ((0.13336) (x) - x)$$

Setting this to zero:

$$1.5(1^{222}/1225) + (3/1225) ((0.13336) (x) - x) = 0$$

$x = 833.25$

So at a stack size of 833.25, the attacker can begin jamming with an additional hand, ATs. For stack sizes greater than this, only aces are profitable jamming hands. At the exact stack size above, ATs is actually indifferent to jamming or folding; the attacker can utilize any mixed strategy he chooses with ATs.

Now let us consider the defender's options, once the attacker begins to jam ATs as well as AA. Can he call with any additional hands? The hand (other than AA) that performs best against the distribution {AA, ATs} is actually AKs, with 41.442% equity. If the defender calls with AKs, he actually loses money against this distribution of raising hands. Once the stack size of 833.25 is reached, then the attacker's strategy goes from {jam 100% of AA} to {jam 100% of AA and 100% of ATs}. There is no gradual ramping-up where the attacker jams with just 1% of his ATs; instead, at the moment it becomes profitable to do so, he jams with every ATs.

This interesting phenomenon is analogous to phase transitions in physics. At a critical temperature, the substance suddenly transforms from solid to liquid, the substance has a drastically different configuration for a minute move in the temperature. That is because the equilibrium state may greatly change even for a small perturbation in the system. This can be seen even in the simplest systems, even in one dimension by perturbing $f(x)$ a little there are critical points where the minimum "jumps" from one place to another. In the same way, we can think of the relationship between the size of the pot and the stacks in this game as a temperature - as we raise this ratio, the game becomes warmer. Playing with blinds of 0.5 and 1 and stacks of 800 is a very low-temperature game. Of course, blinds of zero (where it would be co-optimal to fold all hands) would be analogous to absolute zero temperature - where there is no motion whatsoever.

Raising the temperature of the game just a little, we lower the stacks to approximately 833.12. This brings us to the next hand that becomes profitable to jam, which is A5s. (A5s has 13.331% equity against aces). When ATs became profitable, we simply added it full-force to the attacker's jamming range. However, as we add hands, sooner or later the defender will find it profitable to call with another hand.

For the defender, the hand which has the best equity against the range {AA, ATs, A5s} is AKs - and it is actually the favorite if the attacker jams with all three of those hands (50.9%). So if the attacker were to jam with AA, ATs and A5s, the defender could respond by calling with AKs. and the attacker would lose money. If the attacker were to simply exclude A5s from his range, however, his opponent could respond by calling with only AA and the attacker would give up equity. The attacker must figure out a mix of hands with which he can raise such that he maximizes his profit against the best response strategy his opponent can play. To do this, he will seek to make the defender indifferent to calling with AKs.

We'll assume the stacks are exactly 833, a value that is just a little warmer than the threshold of 833.12 at which A5s becomes profitable. We'll also only deal with the *change* in equity (for fee defender) that occurs between the two defender's strategies.

Attacker's hand	Equity difference ($\langle \text{AKs, call} \rangle = p(\text{call})p(\text{AKs wins pot}) - \text{cost of call}$)
AA	$(3/1225)((0.12141)(1666)-833) = -1.5447$
ATs	$(3/1225)((0.70473)(1666)-833) = 0.8353$
A5s	$(3/1225)((0.69817)(1666)-833) = 0.8085$

Given this, we can see that jamming with A5s is actually better than jamming with ATs when the opponent will call with both AA and AKs.

In order to make the opponent indifferent, we have to make it such that his equity difference from calling is zero. Hence,

$$(-1.5447) (AA\%) + (0.8353) (ATs\%) + (0.8085) (A5s\%) = 0$$

We know that AA% will be 1 because aces have far superior equity to the other hands against any hand, so:

$$0.8353 ATs\% + 0.8085 A5s\% = 1.5447$$

We add hands marginally now as they perform best - so A5s goes to 100% as well:

$$0.8353 ATs\% = 0.70935$$

$$ATs\% = 87.73\%$$

This would indicate that at a stack size of 833, if the attacker jams with the range {AA-100%, ATs - 87.73%, A5s - 100%}, the defender will be indifferent to calling with AKs, and the attacker's equity will be maximized.

Notice the intriguing effect of this - we previously jammed 100% of ATs; now at a stack size less than half a unit smaller, we jam it just a little over 25% of the time. An exhaustive analysis of this game yields any number of oddities like this - in fact, at some stack sizes, hands that were previously jams for the attacker are eliminated completely and replaced by other hands (which do better against the opponent's new range).

Very Small S

Now that we have scratched the surface of what the solution to this game looks like at very large stack sizes, let us consider what happens at very small ones.

First, consider the defender's position. For a stack size S , the defender will call with any hands that have positive equity against the opponent's range. The defender will be facing a bet of size $(S - 1)$, and a pot of size $(S + 1)$. If his equity in the pot when called is x , his equity from calling will be:

$$x(2S) - (S-1) > 0$$

$$x(2S) > S- 1$$

$$x > (S- 1)/(2S)$$

So the defender will call when his equity is greater than $(S- 1)/(2S)$.

At a stack size of 1.1, for example, the defender would call when his equity is greater than $1/22$. If the attacker jammed with all hands, then the defender would always have a call, as the equity of the worst hand (32o) is 32.3% against all hands.

Let's assume, then, that the defender will call with all hands. The attacker will then jam with all hands whose EV is better by jamming than by folding. For any hand, its EV from jamming is simply the outcome of the all-in confrontation against the range of all hands that the defender is calling against, while folding has an EV of 0.

$$\langle \text{ATT, all} \rangle (2S) - (S - 0.5) > 0$$

$$\langle \text{ATT, all} \rangle > (S - 0.5)/2S$$

So the attacker will jam with all hands whose equity is greater than $(S - 0.5)/2S$. At a stack size of 1.1, then, he will call when his equity is greater than $0.6/2.2$, or $3/11$. All hands have at least this equity, so the attacker would jam with all hands. Since neither player can unilaterally improve by changing his strategy, these strategies of jamming all hands and calling all hands are optimal for a stack size of 1.1

As we decrease the "temperature" of this very hot game, it is the attacker who first runs into hands that have insufficient equity with which to jam.

32o has 32.3% equity, so substituting in the above equity formula:

$$0.323 > (S - 0.5)/(2S)$$

$$S > 1.412$$

So when the stacks get to 1.412 units, the attacker can no longer profitably jam with 32o. Contrast this to the case of very large stacks, where in all cases, the attacker jammed with more hands than the defender called with.

We can use the same fictitious play technique to examine what the proper strategies are at different stack sizes. For example, consider the stack size 2. Using our EV formula from above, the defender will call whenever his equity hi the pot is greater than $(2 - 1)/(4)$, or $1/4$. Assume temporarily that the defender will call with all hands. The attacker's best counter-strategy then, is to jam with all hands that have more than $[2 - 0.5]/4$, or $3/8$ equity. This excludes a total of 13 hands: 32o, 42o, 52o, 62o, 72o, 82o, 43o, 53o, 63o, 73o, 83o, 32s, and 42s.

Now we look for the best counter-strategy for the defender. Against a range which excludes the above 13 hands, the hand with the lowest equity is still 32o, with 31.72%. As a result, the defender still defends with all hands. Here, both sides have again maximized their equity against then opponent's strategy and neither side can improve unilaterally, so the strategies are optimal.

This may seem a little peculiar to readers who arc experienced with no-limit. Normally it is Correct (or so the conventional wisdom goes, which we will see in a moment agrees with analysis) to call with fewer hands than one would jam. The reason the opposite is true here at very small stacks is that the additional 0.5 units that the attacker has to expend to play the hand is very significant relative to the size of the stacks. As we increase the size of the stacks, we find that this effect is diminished, and the ratio of optimal attacking hands at a specific stack to defending hands at that stack begins to rise again.

Continuing to lower the temperature of the game, we look at stack size of 3. Now the attacker's equity requirement against a defender who will always defend is $(3 - 0.5)/6 = 5/12$. Against a defender who calls all the time, this excludes 34 hands:

32o and 32s
 42o-43o and 42s-43s
 52o-54o and 52s-54s
 62o-65o and 62s-64s
 72o-75o and 72s-73s
 82o-85o and 82s-83s

92o-94o

The defender's equity to call is $(3 - 1)/6$, or $1/3$. Against the attacker's range, the following hands have insufficient equity to call against the attacker's jamming range:

32o
42o
52o
62o-63o
72o-73o
82o-83o

Now for the attacker, the defender's range has changed, and so his expected value function changes. Now he will jam whenever the following is true:

$$(\% \text{ defender calls}) (p(\text{ATT wins}) (\text{new pot}) - \text{cost of call}) + (\% \text{ defender folds}) (\text{pot}) > 0$$

This equation yields a different range of jamming hands; then the defender reacts to this new range, and so on. What we have described here is a loose variant of the fictitious play algorithm we described earlier. Applying that algorithm to the problem, we can find the optimal strategies at a stack size of 3.

Attacker jams:

{22, A2s, A2, K2s, K2, Q2s, Q2, J2s, J2, T2s, T2, 92s, 95, 84s, 85, 74s, 76, 64s, 54s}

Defender folds:

{32, 42, 52, 63-, 73-, 83-}

In fact, we can solve the jam-or-fold game using the fictitious play technique for any stack size. *The jam-or-fold tables*, which are the optimal solution to the jam-or-fold game for all stack sizes, are reproduced on the next two pages and in the reference section. The tables contain two numbers for each hand - a "jam" number, which is the stack size S below which the attacker should jam, and a "call" number, which is the stack size S below which the defender should call. It is true that (as we saw) at certain stack sizes hands do appear and disappear (in narrow bands), but this can be safely ignored as a practical matter. The upper threshold for these tables is a stack size of 50. Hands that are labeled JAM or CALL are simply calls at any stack size below 50. The few cases where hands disappear and then reappear at significantly higher stack sizes are noted in the tables.

No Limit Holdem - Jam or Fold tables

(suited hands to the upper right)

Attacker	A	K	Q	J	T	9	8	7	6				
A	JAM	JAM	JAM	JAM	JAM	JAM	JAM	JAM	JAM	JAM	JAM	JAM	47.5
K	JAM	JAM	JAM	JAM	JAM	JAM	JAM	49.2	36.2	32.2	26.4	19.9	19.3
Q	JAM	JAM	JAM	JAM	JAM	JAM	JAM	30.2	29.4	24.4	16.3	13.5	12.7
J	JAM	JAM	JAM	JAM	JAM	JAM	49.8	32.3	18.6	16.2	19.5	10.6	8.8
T	JAM	JAM	44.7	46	JAM	JAM	JAM	35.8	24.7	11.9	10.5	7.7	6.5
9	44.9	24.2	24.3	29.2	31.9	JAM	JAM	36.1	26.8	14.4	6.9	4.9	3.7
8	42.6	18.6	13	14.1	18.4	20.5	JAM	43.2	30.9	18.8	10.1	2.7	2.5
7	40.8	16.1	10.3	8.5	9.9	10.8	15.8	JAM	35.7	23.8	13.9	2.5	2.1
6	34.8	15.3	9.8	6.5	5.7	5.2	7.1	11.2	JAM	29.3	16.3	*	2
5	36.9	14.4	8.9	6	4.1	3.5	3	2.6	2.4	JAM	23.5	**	2
4	34.5	13.1	8.3	5.4	3.8	2.7	2.3	2.1	2	2.1	JAM	***	1.8
3	31.9	12.5	7.5	5	3.4	2.5	1.9	1.8	1.7	1.8	1.8	JAM	1.7
2	29.2	11.6	7	4.6	3	2.2	1.8	1.6	1.5	1.5	1.4	1.4	JAM

Starred cells indicate hands that have broken strategies:

- * 63s = 2.3, 5.1 – 7.1
- ** 53s = 2.4, 4.1 – 12.9
- *** 43s = 2.2, 4.8 – 10.0

(suited hands to the upper right)

Defender	A	K	Q	J	T	9	8	7	6	5	4	3	2
A	CALL	CALL	CALL	CALL	CALL	47.3	40.7	35.6	30.9	30.1	25.6	24.7	23.2
K	CALL	CALL	CALL	44.9	31.9	24.2	17.8	15.3	14.4	13.3	12.2	11.4	10.8
Q	CALL	46	CALL	29.1	24.3	16.2	13	10.6	10	8.9	8.5	7.8	7.2
J	CALL	27.4	19.9	CALL	18.4	13.5	10.7	8.8	7.1	6.9	6.2	5.8	5.6
T	CALL	24.1	15.6	13.2	CALL	11.6	9.3	7.4	6.3	5.2	5.2	4.8	4.5
9	39.7	17.6	11.8	9.9	6.5	CALL	8.3	7	5.8	5	4.3	4.1	3.9
8	34.5	14.1	9.8	7.7	6.7	6.1	CALL	6.5	5.6	4.8	4.1	3.6	3.5
7	29.2	12.5	8	6.4	5.5	5	4.7	CALL	5.4	4.8	4.1	3.6	3.3
6	21.9	11.1	7.4	5.4	4.7	4.2	4.1	4	CALL	4.9	4.3	3.8	3.3
5	20.6	10.3	6.8	5.1	4	3.7	3.6	3.6	3.7	42.6	4.6	4	3.6
4	15.6	9.2	6.3	4.5	3.8	3.3	3.2	3.2	3.3	3.5	31.9	3.8	3.4
3	16.9	8.5	5.9	4.5	3.6	3.1	2.8	2.8	3	3.1	3	22.4	3.3
2	16	8.3	5.6	4.2	3.5	3	2.8	2.6	2.7	2.8	2.7	2.6	15.1

For attacker, the number indicated is the stack size (in big blinds) below which he should jam.
 For defender, the number indicated is the stack size (in big blinds) below which he should call a jam.

For both players, the stacks indicated are the stack size in play before the blinds of 1/2 and 1 are posted.

Practical Applications and Additional Comments

Most situations that occur in practice do not meet the conditions of these examples - that is, we would not play jam-or-fold with 500 times the blind. But we can see that at extremely large stack sizes, it is optimal to play very few hands, and at extremely small stack sizes, it is optimal to play most, if not all, hands. As the size of the stacks decreases through the range between these extremes, it turns out to be the case that the attacker and defender strategies do in fact add hands at a more or less steady pace.

We can compare the results of jam-or-fold to the results from Jam-or-Fold Game #2 (Example 12.2), where the weaker hand had $\frac{1}{3}$ equity in the pot. In that game, at stack sizes of 10, the attacker jammed with $\frac{90}{169}$ (53.2%) of his hands, and the defender called with $\frac{9}{26}$ (34.6%) of his. In jam-or-fold no-limit holdem, at stack sizes of 10, the attacker jams with $\frac{774}{1326}$ (58.3%)

of his hands and the defender calls with $494/1326$ (37.3%) of his As you can see, these results are relatively close - the slightly more aggressive strategies in the real no-limit game come from the fact that many matchups are even closer than two to one.

It is of particular interest what types of hands are included in the jamming strategy. In most cases, the hands with the most showdown value against a random hand are included at the upper range of hands (that is, for even very large stack sizes, jamming with high pocket pairs, for example, is indicated). But as the stack size decreases through the typical ranges, a number of hands get added that actually have significantly less showdown equity against random hands. For example, 97s is a jam for stack sizes up to approximately 36 units, despite the fact that it is an underdog against a random hand. The reason for this is that for the attacker, it is irrelevant what equity a jamming hand has against a random hand; what matters is only its equity against the range of hands with which the opponent will call. And middle suited connectors do quite well (relatively speaking) against the types of hands that make up a significant percentage of calling hands, such as high pairs and big aces. The attacker tables also highlight the importance of suited cards - the additional 2.5% that suitedness adds makes it far more attractive to jam with suited cards than their unsuited brethren.

The reader may then ask, "What use is this? I do not play in games where the only options are jamming or folding." It is our contention that playing jam-or-fold for small stack sizes is either the optimal or at minimum a close approximation of the optimal strategy for heads-up no-limit holdem for chose stack sizes and blinds. The proliferation of online single-table tournaments, satellites, and other structures where heads-up high-blind situations are common makes this particularly relevant to the serious player. We concern ourselves, then, with the important question, "When can jam-or-fold reasonably be brought to apply?"

We begin with the following assumptions:

- On an information-hiding basis, it is desirable to play all hands in a consistent manner; that is, raising all hands to a small amount, all-in, calling, or folding.
- Because the distribution of hands that the attacker will play is stronger than the distribution of hands that the defender holds, the attacker should raise in preference to calling.

Call the strategy of playing jam-or-fold J . The defender maximizes his expectation against J by playing J in response. Now consider an alternate strategy J_1 , where the first player raises some set of hands (to be determined later) to 2 units.

Let us assume that the defender plays the following (definitely sub-optimal) strategy J_2 in response to the raise to 2 units:

Jam with all hands that would have called a jam.
Fold all hands that would have folded to a jam.

Against J_2 , J_1 seems to be stronger than J ; the attacker only risks 1.5 units and gets the same number of folds as he would by jamming his entire stack. If the opponent jams in response, he may be able to fold some of his hands and gain equity. At very small stacks, however, this is impossible. For example, at a stack size of 4, the opponent will jam to 4 units (after a raise to 2) with 115 of 169 different hand types, or approximately $838/1326$ (63%) of his hands (this is actually modified for card removal effects based on the attacker's hand, but we ignore this for now).

If the defender jams with 63% of his hands, the attacker clearly has higher equity from calling

than folding with any hand. Even 32o has more than 30% equity against the range of jamming hands for the defender. So the attacker raising any hand is tantamount to jamming that hand. Since raising to two units in this case is equivalent to jamming, he should only raise the hands in J . But this gives the opponent a strategic option (calling and seeing a flop). Since strategic options have non-negative value, the attacker should simply play jam-or-fold at a stack of 4.

Equally obviously, this is not true, at stacks of, say, a hundred. If the opponent played J_2 with such a large stack, this would be excellent for the attacker. He would simply raise to two units with all hands and fold to the defender's jam without a strong hand. Somewhere between these two extremes, however, there is some stack size at which the attacker can likely start to raise smaller amounts profitably. Below this stack size, we can reasonably say that jam-or-fold is essentially the optimal strategy. We're not going to attempt to prove this, or find some specific exact stack size over which this is optimal. Instead, we're looking for a practical recommendation that shows approximately where we can start playing jam-or-fold. We're looking for a stack size above which we can begin to raise a smaller amount profitably with hands outside the jam-or-fold range.

In order to find this stack size, we can look for a situation where for a hand just below the attacker's jamming cutoff, the equity from raising to 2 units and then folding to a jam from J_2 is equivalent to the equity from jamming. At a stack size of eight, the defender will reraise all-in, playing J_2 with $602/1326$, or 45.4% of his hands. Approximately the worst hand that is jammed by the attacker is Q4o. So Q3o is a hand close to the threshold. Against the defender's range of calling hands, Q3o has about 34.8% equity.

The equity of jamming eight units with Q3o is:

$$\begin{aligned} \langle \text{jam, Q3} \rangle &= p(\text{defender folds}) (\text{defender's blind}) + p(\text{defender calls}) p(\text{attacker wins}) \\ &\quad (\text{defender's stack}) + p(\text{defender calls}) (p(\text{defender wins}) (\text{loss of attacker's stack})) \\ \langle \text{jam, Q3} \rangle &= (0.546) (+1) + (0.454) (0.348) (+8) + (0.454) (0.652) (-8) \\ \langle \text{jam, Q3} \rangle &= -0.558 \text{ units} \end{aligned}$$

This value is close to what we would expect, as it is just below the cutoff for jamming (compared to the -0.5 units for folding). The difference is caused by card removal effects that impact this particular hand.

The equity of raising to two units with Q3o is (again ignoring card removal effects):

$$\begin{aligned} \langle \text{raise to 2, Q3} \rangle &= p(\text{defender rejams}) (\text{loss of raise}) + p(\text{defender folds}) (\text{gain of blind}) \\ \langle \text{raise to 2, Q3} \rangle &= (0.454) (-2) + (0.546) (+1) \\ \langle \text{raise to 2, Q3} \rangle &= -0.638 \end{aligned}$$

At a stack size of eight, then, we begin to approximate the situation we're looking for; where the attacker is indifferent to jamming S units or raising to two units and folding to a jam with hands weaker than his set of jamming hands. As a first-order approximation, then, we could use eight units as the threshold to begin playing jam-or-fold. However, there is another important effect that the above analysis neglects. The defender is not restricted to playing J_1 , in fact, it would be foolish for him to do so!

Since the defender can simply call our 2-unit raise, getting 3-1 on his call, the equity of raising two units is decreased. As a result of this, we add a couple additional units to the stack size threshold at which we begin to play jam-or-fold. Therefore, we believe it is likely optimal to play jam-or-fold in the headsup no-limit case at a stack size of 10 units or below, and potentially as

high as 12 or 13 units when the attacker will be out of position postflop. In situations where the blinds are not of the 0.5-1 structure (such as when the blinds are 2 or 3 chips) or when there are antes (such as when the full table has anted and the play has become headsup by the entire field folding), we simply allocate the antes into the blinds. Two-thirds of the money in the pot is normalized to 1 unit for the purpose of utilizing the jam-or-fold tables. This is only an approximation, but it is useful in practice. Also note that when playing jam-or-fold, only the smaller of the two stacks needs to be below the 10 blind threshold - it doesn't matter how big your stack is if the opponent's stack is small, or vice versa.

Returning to the question we posed at the beginning of this chapter:

Two equally skilled players contest a headsup jam-or-fold no-limit holdem freezeout. Each has sixteen chips to begin the hand. The blinds are one chip on the button and two chips in the big blind. With what percentage of hands should the button move all-in optimally?

Utilizing the jam-or-fold tables, we can see that the set of jamming hands is approximately 61.7% of all hands. That number startles many people on first sight. After all, who would intuitively think that you would put in sixteen chips to win three with more than three-fifths of your hands? This is a powerful idea of which many players are not aware.

These strategies are of great value in the headsup portion of tournaments and in playing in blind vs. blind situations when the blinds are relatively large. These situations occur frequently in, for example, online single-table tournaments where the blinds escalate quickly. In addition, the jam-or-fold game shows the great value of jamming preflop in no-limit holdem. By playing jam-or-fold even at high stack sizes, a weaker player can fairly effectively neutralize a strong player's postflop skill, giving up only a modest amount of equity by doing so.

Key Concepts

- In jam-or-fold where the best hand wins the entire pot, optimal strategy is fairly tight and play favors the defender for most stacks.
- When the worst hand has significant showdown equity, however, optimal strategy loosens up a significant amount.
- In jam-or-fold no-limit holdem, strategies are significantly looser than most players think.
- When selecting hands with which to jam, the showdown equity against the hands that will call is of paramount importance.
- Jam-or-fold is at least near-optimal for stack sizes up to about 10 or 11 blinds.

Chapter 13

Poker Made Simple: The AKQ Game

We continue our investigation of game theory in poker games with a game that represents perhaps the simplest form of symmetrical, non-trivial poker. Our first exposure to this game was a Card Player magazine article by Mike Caro. We suppose that he intended the game as a simplified object lesson in betting and bluffing ratios, and in that capacity it serves quite well. We have also expanded on this by studying this beyond the limit case it was first presented in, and in doing so, discovered something surprising about no-limit poker. But for now we concern ourselves with the half-street limit case for this game.

The fundamental elements of the game are these:

- There is a three-card deck, containing an Ace, a King, and a Queen.
- The players are each dealt one card without replacement and there is an initial ante.
- There is a round of betting, after which there is a showdown (if neither player folds). In the showdown, the high card wins.

The three-card deck introduces a topic that we have not so far considered in our games - *card removal*. In the [0,1] game both players had symmetric distributions, and one player's card did not affect his opponent's card. However, in this game, if you hold the ace, then the opponent holds either the king or the queen. Likewise, if you hold the queen, the opponent holds the king half the time and the ace half the time.

Example 13.1 - AKQ Game #1

One half street.

Pot size of 2 units.

Limit betting of 1 unit.

The complete ex-showdown payoff matrix (from Y's point of view) for this game as follows:

	Y	Ace		King		Queen	
X		Bet	Check	Bet	Check	Bet	Check
Ace	Call			-1	0	-1	0
	Fold			+2	0	+2	0
King	Call	+1	0			-1	0
	Fold	0	0			+2	0
Queen	Call	+1	0	+1	0		
	Fold	0	0	0	0		

Our goal is to find the optimal strategies for this game. Immediately we can reduce this game by eliminating dominated strategies. Comparing the options when X holds an ace, we can see that calling has expectation greater than or equal to that of folding in all cases. Thus folding an ace is dominated by calling with it. As a result of this, we can eliminate the strategic option of folding an ace from X's strategy options, since it is dominated by calling with an ace. In the same way, we can see that calling with a queen is dominated by folding a queen. So we can eliminate calling with a queen from the strategic options, leaving us with a simplified game:

	Y	Ace		King		Queen	
X		Bet	Check	Bet	Check	Bet	Check
Ace	Call			-1	0	-1	0
King	Call	+1	0			-1	0
	Fold	0	0			+2	0
Queen	Fold	0	0	0	0		

Now we can turn our attention to strategies that are dominated for Y. We can see that betting a king is dominated by checking it for Y, as well as checking an ace. So we can further simplify our game by eliminating these alternatives:

	Y	Ace	Queen	
X		Bet	Bet	Check
Ace	Call		-1	0
King	Call	+1	-1	0
	Fold	0	+2	0

Let's take a moment to intuitively discuss what's occurred here. First, we considered X's options. We found two strategies where X could always do at least as well: calling with his aces, and folding his queens to bets. Hopefully this is obvious - we never fold the nuts to a bet on the river, nor call when we could fold with a completely hopeless hand. Since X would never take these actions and can never gain by taking them in the future, we can simply eliminate them from the game matrix.

Next, we took the reduced game, and applied the same reasoning to Y's choices. If Y has a king, he knows he can never bluff X's aces out, and can never be called by X's queens. And if he has an ace, he can never lose by betting, and he might get a call from a king.

By eliminating the dominated strategies, we are left with a much more simplified game - in fact, from a strategic option standpoint, this game is a 2x2 matrix, just as games in Chapter 13 were. Here we see the oscillation of strategies from a pure standpoint, as we described there.

If X calls all the time with kings, Y's best response is never bluffing and only value betting. In turn, X responds to that by folding all kings. Y exploits that by bluffing all the time; and X returns to his original strategy of calling all the time with kings. This indicates that we have an equilibrium between these strategies in mixed strategies.

We can actually create a 2x2 matrix by determining the equities of the undominated strategies against each other:

The two undominated strategic options are whether Y should bluff queens, and whether X should call with Kings.

We End Y's equity from the two options by considering the possible hands that X can hold and the outcomes.

We can name the four strategic options:

Y_1 : Y bluffs with queens

Y_2 : Y checks queens.

X_1 : X calls with kings

X_2 : X folds kings.

We can find the expectation of each of these matchups:

Y_1 vs. X_1 :

When Y has an ace, he gains a bet against X's kings and nothing against X's queens. When Y has a king, no betting occurs. When Y has a queen, he bluffs and loses a bet to X's aces and kings.

Each card matchup occurs $1/6$ of the time; Y gains a bet in one, and loses a bet in two for a total of:

$$\langle Y_1, X_1 \rangle = -1/6$$

Y_2 vs. X_2 :

When Y has an ace, he gains a bet against X's kings and nothing against X's queens. When Y has a king, no betting occurs. When Y has a queen, no betting occurs.

Y gains a bet in one case, and does not lose a bet, for a total of:

$$\langle Y_2, X_1 \rangle = +1/6$$

Y_1 vs. X_2 :

When Y has an ace, he never gains a bet. (X folds.) When Y has a king, no betting occurs. When Y has a queen, he bluffs and wins the pot (of two bets) when X has a king, and loses a bet when X has an ace.

Y gains two bets and loses one, for a total of:

$$\langle Y_1, X_2 \rangle = +1/6$$

Y_2 vs. X_2 :

When Y has an ace, he never gains a bet (X folds). When Y has a king, no betting occurs. When Y has a queen, no betting occurs, and so there is no net impact.

$$\langle Y_2, X_2 \rangle = 0$$

This results in the following 2×2 matrix:

	Y bets queens	Y checks queens
X calls with kings	$-1/6$	$+1/6$
X folds kings	$+1/6$	0

Multiplying all the cells of a payoff matrix by a positive constant does not change the strategic implications of the game. Multiplying by -1 simply swaps the payouts between the two players.

If we multiply the AKQ Game #1 matrix by -6, we have:

	Y bets queens	Y checks queens
X calls with kings	+1	-1
X folds kings	-1	0

This matrix should be familiar to the reader from Chapter 10; recall that the payoff matrix of Cops and Robbers was:

	Robber robs	Robber stays home
Cop patrols	+1	-1
Cop doesn't patrol	-1	0

This game is identical, then, to Cops and Robbers. We know from the analysis of that game that the optimal strategies are $\frac{1}{3}$ and $\frac{1}{3}$, and we can compare this result to a result we will find directly.

We can also solve the game by finding for each player a strategy that makes his opponent indifferent between his two strategic options. Consider Y, who must decide whether to bluff with a queen. His ex-showdown EV from the two actions he can take with a queen (using c as X's calling frequency with a king) is:

$$\begin{aligned} \langle Y, \text{bluff} \rangle &= (\text{lose to aces}) + (\text{lose to kings}) - (\text{kings fold, win the pot}) \\ \langle Y, \text{bluff} \rangle &= p(\text{X has an ace}) (\text{lose one bet}) + p(\text{X has a king}) p(\text{X calls}) (\text{lose one bet}) + \\ &\quad p(\text{X has a king}) p(\text{X folds}) (\text{the pot}) \\ \langle Y, \text{bluff} \rangle &= (\frac{1}{2})(-1) + (\frac{1}{2})(c)(-1) + (\frac{1}{2})(1-c)(2) \\ \langle Y, \text{bluff} \rangle &= \frac{1}{2}(-1) + \frac{1}{2}[c(-1) + (1-c)(2)] \\ \langle Y, \text{bluff} \rangle &= -\frac{1}{2} - \frac{1}{2}c + 1 - c \\ \langle Y, \text{bluff} \rangle &= \frac{1}{2} - \frac{3}{2}c \\ \langle Y, \text{check} \rangle &= 0 \end{aligned}$$

Setting these two things equal:

$$\begin{aligned} \frac{1}{2} - \frac{3}{2}c &= 0 \\ c &= \frac{1}{3} \end{aligned}$$

Next, set X to be indifferent to calling with a king. Using b as Y's bluffing frequency with a queen:

$$\begin{aligned} \langle X, \text{call} \rangle &= (\text{lose to aces}) + (\text{win against queens}) \\ \langle X, \text{call} \rangle &= \frac{1}{2}(-1) + \frac{1}{2}[b(1)] \\ \langle X, \text{call} \rangle &= -\frac{1}{2} + \frac{1}{2}b \\ \langle X, \text{fold} \rangle &= \frac{1}{2}(b)(-2) \\ \langle X, \text{fold} \rangle &= -b \\ -\frac{1}{2} + \frac{1}{2}b &= -b \\ b &= \frac{1}{3} \end{aligned}$$

So we have the optimal strategy pair: X should call $\frac{1}{3}$ of the time when he has kings, and Y should bluff $\frac{1}{3}$ of the time when he has queens.

Let's again take a moment to discuss intuitively what's going on here. The idea for Y is to get value from his set of hands. In an exploitive sense, if he knew what strategy X were going to

employ with his kings, he could bet or bluff accordingly. However, let's say that he doesn't and he wants to simply get all his value no matter what the opponent does.

There are two sources of value for him: stealing pots with queens, and getting paid off with aces. So he wants to create a situation where the value he gets from each of these options is equal. Then his opponent won't be able to exploit him. What he does, then, is bluff with a frequency so that the value of his whole strategy is the same no matter what the opponent does; specifically, no matter whether the opponent calls with a king or not. It turns out for this situation that the proper ratio of aces to bluffed queens is three to one; hence he should bet $\frac{1}{3}$ of his queens, since he will bet all his aces. Recall that in Chapter 11, we called this ratio of bluffs to bets α , and showed that it had the value $\frac{1}{P+1}$. (See Equation 11.1) For the pot of two bets, we can see that this ratio indeed calculates to $\frac{1}{3}$.

Similar reasoning works for X. He wants to call often enough that the opponent is indifferent to bluffing with his queens. When his opponent bluffs successfully, he wins two bets, while bluffing unsuccessfully costs just one bet. So X must call twice as often as he folds in order to make the bluffer indifferent. This fraction is $\frac{2}{3}$ (leaving $\frac{1}{3}$ of hands to fold). But half of the time that the opponent is bluffing, X will hold an ace. So he needs to call with an additional fraction of his hands equal to $\frac{2}{3} - \frac{1}{2}$, or $\frac{1}{6}$. Since kings make up $\frac{1}{2}$ of his hands when the opponent has a queen, he needs to call with $\frac{2}{6}$ or $\frac{1}{3}$ of kings. Again recall in Chapter 11, we showed that X would fold α of the time with hands that could beat a bluff. Here again, he folds $\frac{1}{3}$ of the total hands he holds that can beat a bluff.

Y gains value in this game, as you would expect; after all, he could just check behind all the time to get zero value, and the betting and bluffing is valuable to him. We can calculate the value of the game to Y by evaluating the strategies and outcomes. We could calculate all the fractional possibilities for each action, but instead we can use a shortcut. Recall that when the players play optimal strategies, X is indifferent to calling; that is, he does equally well calling or folding. So we can simply pretend that he folds all the time. Now if he actually did fold all the time, he would be exploitable. But since we're only calculating the value of the game and not the strategies, we can simply hold Y's strategy constant.

The value of the game to Y is actually $\frac{1}{18}$ of a bet. We invite the reader to verify that this is the value of the game for all strategies that X may employ as far as calling with kings.

We have now been able to solve this half street game, finding the optimal strategies, for a single pot size (2 bets) - however, we can use the same methodology as we employed to solve it for any arbitrary pot size. If we restate the game with an arbitrary pot size P, we have the easily solvable AKQ Game #2.

Example 13.2 - AKQ Game #2

One half street.

Pot size of P units.

Limit betting of 1 unit

Again, we have the complete game payoff matrix:

Y		Ace		King		Queen	
X		Bet	Check	Bet	Check	Bet	Check
Ace	Call			-1	0	-1	0
	Fold			+P	0	+P	0
King	Call	+1	0			-1	0
	Fold	0	0			+P	0
Queen	Call	+1	0	+1	0		
	Fold	0	0	0	0		

This matrix can be reduced by eliminating dominated strategies, just as in AKQ Game #1 (Example 13.1), yielding:

	Y	Ace	Queen	
X		Bet	Bet	Check
Ace	Call		-1	0
King	Call	+1	-1	0
	Fold	0	+P	0

Our two indifference equations are modified as follows (assume b is Y's bluffing frequency with queens, and c is X's calling frequency with kings):

$$\begin{aligned} \langle Y, \text{ bluff } \rangle &= (\text{lose to aces}) + (\text{lose to kings}) + (\text{kings fold, win the pot}) \\ \langle Y, \text{ bluff } \rangle &= (\frac{1}{2})(-1) + (\frac{1}{2})(c)(-1) + (\frac{1}{2})(1-c)(P) \\ \langle Y, \text{ bluff } \rangle &= \frac{1}{2}(-1) + \frac{1}{2}[c(-1) + (1-c)(P)] \\ \langle Y, \text{ bluff } \rangle &= -\frac{1}{2} - \frac{1}{2}c + \frac{1}{2}P - \frac{1}{2}Pc \\ \langle Y, \text{ check } \rangle &= 0 \end{aligned}$$

Setting these two things equal:

$$\begin{aligned} -\frac{1}{2} - \frac{1}{2}c + \frac{1}{2}P - \frac{1}{2}Pc &= 0 \\ 1 + c - P + Pc &= 0 \\ (P+1)c &= P-1 \\ c &= (P-1)/(P+1) \end{aligned}$$

This value of c is X's calling frequency with kings. When Y has a queen, X holds half aces and half kings. So the total fraction of his hands (aces and kings) with which he calls is $\frac{1}{2} + (\frac{1}{2})(P-1)/(P+1)$. This in turn means that he folds $1/(P+1)$, or the now-familiar constant α , of his hands, to keep Y indifferent from bluffing.

In a similar manner, we have:

$$\begin{aligned} \langle X, \text{ call } \rangle &= (\text{lose to aces}) + (\text{win against queens}) \\ \langle X, \text{ call } \rangle &= \frac{1}{2}(-1) + \frac{1}{2}[b(1)] \\ \langle X, \text{ call } \rangle &= -\frac{1}{2} + \frac{1}{2}b \\ \langle X, \text{ fold } \rangle &= \frac{1}{2}(b)(-P) \end{aligned}$$

Again, setting these two equal:

$$\begin{aligned}
 -\frac{1}{2} + \frac{1}{2}b &= \frac{1}{2}(b)(-P) \\
 -1 + b &= -bP \\
 b(1+P) &= 1 \\
 b &= \frac{1}{1+P} \\
 b &= \alpha
 \end{aligned}$$

And again, we see that the bluffing to betting ratio for Y here should be α as well.

We can also quickly calculate the value of the game to Y:

Hand Matchup	Value to second player	$p(\text{matchup})$	Weighted Value
(A, K)	0	$\frac{1}{6}$	0
(A, Q)	$-\frac{1}{P+1}$	$\frac{1}{6}$	$-\frac{1}{6}(P+1)$
(K, A)	0	$\frac{1}{6}$	0
(K, Q)	$+\frac{P}{P+1}$	$\frac{1}{6}$	$\frac{P}{6}(P+1)$
(Q, A)	0	$\frac{1}{6}$	0
(Q, K)	0	$\frac{1}{6}$	0
Total			$\frac{P-1}{6(P+1)}$

So the overall value of the game to Y is $\frac{P-1}{6(P+1)}$. We can see that as the pot size goes to infinity, the value of this game converges on $\frac{1}{6}$. If the pot is extremely large, then X will simply call with just about all his kings, and Y will almost never bluff, making the case where Y holds an ace and X a king worth a full bet to Y. This case occurs $\frac{1}{6}$ of the time.

One additional comment must be made - the above solution is only valid if the pot size is greater than 1. The reason is that when the pot is less than 1, Y loses money by bluffing queens - X calls with aces, which make up half his hands. Since Y gains the pot when X folds a king, and loses a full bet when X calls with an ace, bluffing always loses money, even when X folds all his kings. The entire game changes at that point. As the pot approaches one from the positive side, the bluffing ratio rises - at 2, it's $\frac{1}{3}$, at 1.5, it's $\frac{2}{5}$, at 1.1 it's $\frac{10}{11}$, and then at precisely 1, Y stops bluffing.

We can also use a graph to see the value of the game as the size of the pot changes. Notice that the graph suddenly changes at $P=1$; Y goes from never bluffing to bluffing half of his hands. This happens because calling suddenly becomes a mixed strategy, as does bluffing. Below this pot size, neither side bluffs nor calls; this cannot be exploited because of the frequency of aces, in the opponent's distribution. We can also see at large pot sizes that X has to call frequently while Y can bluff infrequently, as a result, the game favors Y as the pot size increases.

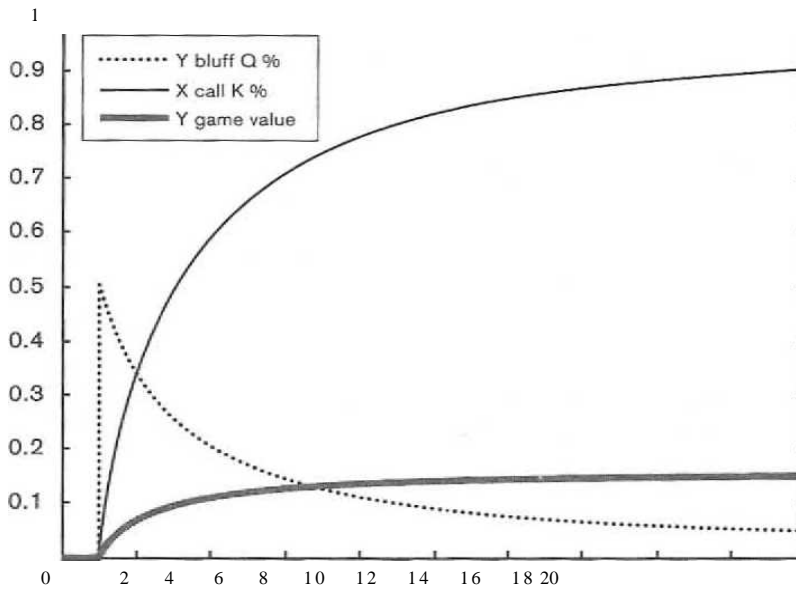


Figure 13.1. Optimal strategies for the limit AKQ game by pot size.

The important lessons here are that changing the size of the pot impacts the players' strategies and their equities; also solving games with specific pot sizes is often just as easy as solving them for arbitrary pot sizes. In the future, we will generally forego solving games with specific pot sizes unless it is necessary for the discussion, preferring to simply treat games as having pot sizes of P .

Key Concepts

- Betting and bluffing ratios are the key to optimal play when both sides have clearly defined hand values (i.e., they have the nuts, a hand in the middle, or a dead hand).
- In limit poker, it is virtually always wrong to bet hands in second position on the river that can beat only a bluff but will be called by only better hands.
- Card removal can have a significant impact on the shape of the opponent's distribution.

Chapter 14

You Don't Have To Guess: No-Limit Bet Sizing

No-limit bet sizing is an important topic, but a topic that has generally been neglected from an analytical standpoint. We find recommendations in the existing literature ranging from pot-sized bets, to small underbets, to betting different amounts according to hand strength, and so on. These recommendations are usually based on a combination of experience and consensus. In some cases, however, we can find the analytical answer.

Let us first consider the half-street no-limit clairvoyant game. When we previously looked at the clairvoyant game, it was limit betting, and the clairvoyant player bet all his whirling hands and a fraction of the losers such that the other player was indifferent to calling. But in no-limit, there are a multitude of bet sizes to choose from. What makes one bet size better than another? How can we find the optimal strategies for such a game?

Example 14.1 - The Half-Street No-Limit Clairvoyance Game

One half street.

Pot size of 1 bet.

No-limit betting. (Stacks are n bets, where n is large)

Y is clairvoyant and his hand is randomly drawn from a distribution that is half winners and half dead hands.

Note: In limit games, we use the convention that the bet size is 1 and the pot is P units. However, in no-limit games with variable-sized bets, we instead use the convention that the pot is 1 unit, and the bet is s . Thus a bet of s units is a bet of s times the pot.

Let us first describe the structure of the strategy. After X checks dark, Y will clearly want to bet his winning hands for value, and protect those winning hands with bluffs such that X is indifferent to calling. Y will choose a bet size s that he will bet with both types of hands. If Y were to bet two different amounts, it would be necessary for such a strategy to be optimal that he have equal expectation against the nemesis with those two bet sizes. Then X will call in a proportion that makes Y indifferent to bluffing.

We can create a function $f(s)$, which is Y 's expectation from betting an amount. Then if we find the place where this function is maximum, we can find Y 's proper bet size and have the solution to the game. When Y bets and bluffs appropriately, then X can fold or call at his discretion - both have equal expectation. Hence for finding Y 's expectation, we can assume that X calls any bet as long as we hold Y 's strategy constant. We will use this idea much more when we look at the $[0,1]$ games later in Part III.

We know from the games we solved in Chapter 11 that Y 's ratio of bluffs to bets will be α , and also that X will fold α of his hands that beat a bluff. Recall from Equation 11.2 that for no-limit games, we have an equivalent definition of α :

$$\alpha = s/(1 + s)$$

For any bet value s , Y will bluff $s/(1 + s)$ as many hands as he value bets. Suppose y is the fraction of Y 's total hands that are winning hands.

$$f(s) = p(Y \text{ has nuts}) (\text{one bet}) + p(Y \text{ bluffs}) (\text{lose one bet})$$

$$f(s) = (y)(s) + (ys/(1 + s))(-s)$$

$$f(s) = (y)(s) - (ys^2)/(1+s)$$

$$f(s) = ys/(1 + s)$$

This function, as seen in the following graph, lacks any local maxima or minima:

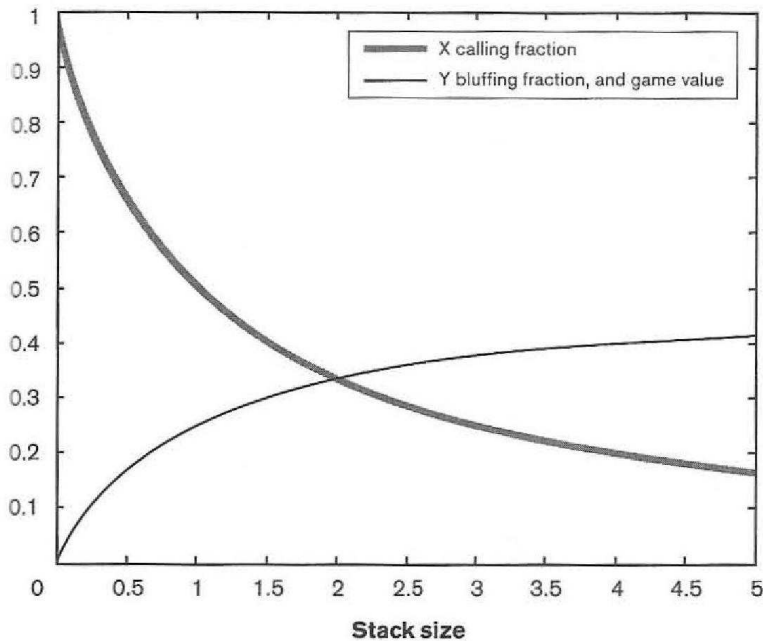


Figure 14.1. The half street clairvoyance game, $y=0.5$

In fact, this function continues to increase as s goes to infinity. Hence Y maximizes his value by betting his whole stack and bluffing with the appropriate corresponding function.

We assumed in the above that Y 's distribution was 50% nut hands and 50% dead hands. But suppose that y is $3/5$, meaning that $3/5$ Y 's hands are winners and $2/5$ are bluff, and the stacks are just five pot-sized units. Now Y bets his whole stack. How should X respond? What if Y simply bets the pot or $s = 1$, how should X respond?

We consider the first question first. Well, the immediate answer based on the reasoning above would be that X should fold $\alpha = 5/(1 + 5) = 5/6$ of his hands and call with the rest. This will make Y indifferent to bluffing. But to see why this is wrong, let's consider the strategy-pair we're discussing:

Y : Bet 5 with all winning hands and with all dead hands.

X : Call with $1/6$ of hands.

With a bluff, Y gains the pot of 1 for $5/6$ of the time for a total win of $5/6$, while losing the bet of $s = 5$ for $1/6$ of the time, for a loss of $5/6$ of a unit, thus breaking even. With value bets, however, Y wins 5 units $1/6$ of the time, for a gain of $5/6$ unit per value bet. Since value bets make up $3/5$ of his hands overall, this is a total win of $1/2$ unit.

But what if X folded all the time? Then Y would win the entire pot with his bluffs, which make up $2/5$ of total hands, and gain nothing otherwise. So Y 's equity would just be $2/5$ of a unit. So these cannot be the optimal strategies! X can unilaterally improve by simply folding all hands. Y could exploit this by bluffing all his hands. Suppose he does this. Then he would gain nothing with his value hands and the entire pot with this dead hands, for a total win of $2/5$ unit, Against

this Y strategy, X's best response is to simply fold.

This is an important result; X's optimal strategy in this situation is to fold all his hands, even though Y is bluffing all of his dead hands. This is because Y's distribution is so strong that he can simply bet and "claim" the pot with all hands, because X loses money by calling at all. We will have more to say about these situations in the chapter on multi-street play. We have the game and equity graphs as follows:




Figure 14.2. Half street no-limit clairvoyance game, $y=0.6$

Suppose Y doesn't bet all his chips, but simply bets the pot ($s = 1$)? It might seem as though we can simply look at the above graph for $s = 1$ and see that we should call with half our hands. This will make him indifferent to bluffing a value $s = 1$, and so on. Unfortunately, it isn't that simple. The right answer is that X should *still* fold all hands. This is even true if Y had bet $s = 1\%$ of the pot.

To see why, consider how Y might exploit X's call. We know that X will fold all his hands if Y bets all his chips. Now suppose that X will respond to a smaller bet by calling in some mixed strategy. Y can exploit this by adopting a strategy of betting the smaller amount with his winning hands and betting all his chips with his bluffs. X could exploit that in turn, but since Y can increase his equity against X's strategy, it cannot be optimal. From another perspective, Y can simply push all-in with every hand and claim the pot—that is Y's optimal play. X does no better playing against Y's optimal strategy than he does by throwing his hand away in the dark. Hence X can play optimally by always throwing his hand away!

In Chapter 13, we considered the limit half-street AKQ game. There, we found that with a fixed pot in relation to the bet size, we could solve the game and find the proper betting and bluffing frequencies for both players. We continue our no-limit bet sizing investigations by considering the no-limit version of this game.

Example 14.2 - AKQ Game #3

One half street.
Pot size of 1 unit.

No-limit betting. (Stacks are n bets, where n is large)

In this game, we have the same situation as previously. X checks dark, and Y then seeks to extract as much value from his hand distribution as possible. However, instead of a binary decision of "bet or check," Y is now faced with a more complex decision. If he does decide to bet, he must choose a bet amount as well. In fact, we can simply consider a "check" as betting zero units.

At a glance, we can see that it would be foolish for Y to bet more than the pot. If he does, then X picks off his bluffs with aces only and never pay off with kings. If he does this, he never loses money when Y has an ace, and gains more than the pot when Y bluffs a queen. We know from solving limit games previously that Y can profit by betting amounts smaller than the pot. Hence, we can discard all bet amounts greater than the size of the pot.

Next, consider what would happen if we were to value bet a different amount than we bluffed. This would immediately be exploitable by X: X could simply fold kings to the value betting amounts and call the bluff amounts. For each value betting amount we utilize, we must have corresponding bluffs so that the ratio of bluffs to value bets (α) is preserved for that bet amount. One additional idea that is important here is that if we have a bet amount s such that we make more profit than betting some alternate amount t , then our strategies cannot be optimal. We could simply switch all the hands that bet t over to betting s and improve our equity. This means that either: 1) several bet amounts will have the same equity; or 2) there will be just one bet amount in Y's optimal strategy.

With all those ideas in mind, we can calculate the optimal strategies as follows. First we will find the optimal bluffing and calling frequencies for any bet amount s . These are the action frequencies which are X's best response to Y betting j and Y's proper bluffing frequency when betting s . Using these frequencies, we can solve for the expectation of the game, as a function of s . Once we do this, we can find the value of j for which Y's equity is maximized. Y will choose to bet and bluff at that particular value: X can do nothing to prevent this.

This is a common technique we use when considering games where different strategy choices can be made. Sometimes, we can solve the game or portions thereof in terms of a variable that only one player controls. When this is the case, we can find the strategies, calculate expectation, and then the player who controls that variable will choose its value such that this expectation is maximized.

Again we know from the games we solved in Chapter 11 that Y's ratio of bluffs to bets will be α , and also that X will fold α of his hands that beat a bluff, where $\alpha = s/(I + s)$. (See Equation 11.2)

So for any bet value s , Y will bluff $s/(I + s)$ of the time when he holds a queen. Likewise, X will fold a total fraction of his hands equal to $s/(I + s)$.

Based on these frequencies, we can calculate the expectation of betting and bluffing any bet size s . We will create a **value function** $f(s)$ that represents this expectation for any s .

We know from previous experience that Y will break even on his bluffs, because X will call often enough to make Y indifferent between bluffing and checking, and checking has an EV of zero. On his value bets, then, Y gains value when the first player calls with a king.

X must call with an overall frequency of $1 - \alpha$, or $1/(1 + s)$. Subtracting $1/2$ from this expression for the aces with which X always calls:

$$c_{kings} = 2(1/(1 + s) - 1/2)$$

$$c_{kings} = 2/(1 + s) - (1 + s)/(1 + s)$$

$$c_{kings} = (1 - s)/(1 + s)$$

We find that the percentage of kings with which X will call is $(1 - s)/(1 + s)$. When he does call, Y's gain will be s units. So in the case where Y holds an ace and X holds a king, he gains $s(1 - s)/(1 + s)$ units. This case represents $1/6$ of the total hands, so the total expectation of the game to Y is:

$$f(s) = s(1 - s)/6(1 + s)$$

Y now chooses s such that the above expression is maximized. To do this requires a little calculus, but before we do that, you can see approximately where the peak will be by looking at the following graph:

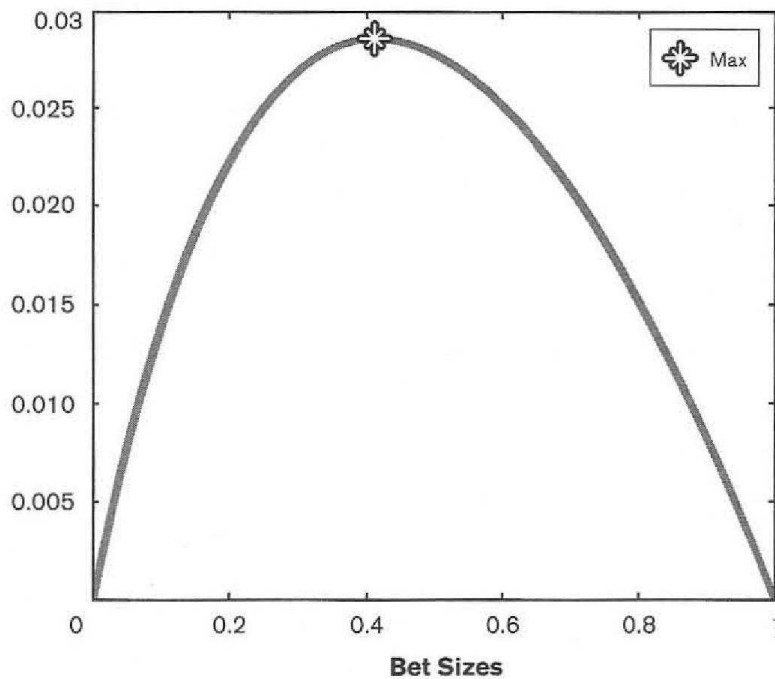


Figure 13.3. Equity of the NL clairvoyance game

Note: Remember that more complex derivations are marked off in boxes so that our less mathematical readers can just "take our word for it" if they choose.

We set the derivative of this expression equal to 0:

$$f(s) = s(1 - s)/(1 + s)$$

$$f(s) = -(s^2 + 2s - 1)/(1 + s)^2$$

$$s^2 + 2s - 1 = 0$$

$$(1 + s)^2 = 2$$

$$s = -1 \pm \sqrt{2}$$

$-1 - \sqrt{2}$ must be an extraneous solution, because we can't bet a negative amount.

So the value $\sqrt{2} - 1$ or approximately 0.114 is the proper amount to bet.

In his book *The Theory of Poker*, David Sklansky coined the phrase "Fundamental Theorem of Poker," a tongue-in-cheek reference to the Fundamental Theorem of Algebra and Fundamental

Theorem of Calculus from introductory texts on those two subjects. The constant $\sqrt{2} - 1$ appears so often in poker analysis that we will in the same vein go so far as to call it *the golden mean of poker*, and we call it r for short. We will see this value in a number of important results throughout this book.

$$r = \sqrt{2} - 1 \tag{14.1}$$

The solution to this game is the first example we have seen of finding the proper bet amount in no-limit. This important concept revolves around the idea that we can find the expectation of the game for all different bet amounts. Once we have done this, we can find the bet amount value so that the bettor's expectation is maximized.

So in a game where the following conditions hold:

- You hold the nuts sometimes and a bluff sometimes.
- When you hold a bluff, your opponent holds the nuts half the time and the other half the time a hand that can only beat a bluff.

The optimal strategy is for you to bet $\sqrt{2} - 1$ times the pot (a little more than 41%) with all of your nut hands, and to bluff with the corresponding frequency of bluffing hands

$(1 - 1/\sqrt{2})$. Your opponent in turn calls with all his nut hands and enough of his in-between hands such that his total calling frequency is $1/\sqrt{2}$.

In addition to discrete games such as the AKQ game, we can use a similar idea to solve more complex no-limit games. Even continuous games such as the [0,1] game can be played no-limit. In the AKO game, we found that there was exactly one optimal bet size that Y used. In the half-street [0,1] game, however, we find that we occasionally bet all different bet sizes.

Example 14.3 - [0,1] Game #3

One half street.

Betting is no-limit; both players have arbitrarily large stacks.

The pot is 1 unit.

Folding is allowed.

Here, we have a similar game, where X checks dark, and Y tries to maximize his equity. However, we have a situation now where there are two variables that must be related - strength of hand and size of bet. In the limit AKQ game, if Y held a queen, he had a simple binary decision - bet or check. In the no-limit case, Y had to additionally choose what amount to bet.

In the half-street [0,1] no-limit game, Y must choose an amount to bet with each hand; it's possible that he may bet one particular amount, or vary his bet based on the strength of his hand. Just as in the previous game, we will attempt to solve the game in terms of the variables, and then have each player choose the variables he controls in order to maximize his equity.

We can first appeal to our previous investigations to conclude that Y will play a two-pronged strategy of betting and bluffing, and that his betting and bluffing regions will be related to each other by a ratio like α . So when X is faced with a bet of t from Y, he will need to call to make Y indifferent to bluffing at the edge of Y's bluffing threshold.

We can express X's strategy as a function of t . Let $x(t)$ be X's calling threshold when faced with a

bet of t . So $x(t)$ will only consist of hands that can beat Y's bluffs - X will never call with hands that will lose even to Y's bluffs. The threshold where X will never call is $x(0)$; We will also call this value x_0 .

X will choose his function $x(t)$ such that if Y's hand is weaker than x_0 , then Y is indifferent to bluffing an amount s . In other words, X will call a bet of size t with exactly enough hands that Y is indifferent between betting t or checking for any value of t . So t_1, t_2, t_3 and so on, all have equity equal to checking. Hence Y is indifferent to betting any s value.

This leads to the following indifference for Y at y :

X's hand	Y bet's	Y checks
$[0, x(s)]$	$-s$	0
$[x(s), y]$	1	0

$$y - x(s) - sx(s) = k \quad (\text{where } k \text{ is constant})$$

$$x(s) = (y - k)/(1+s)$$

We know that $x(0) = x_0$, so $y - k = x_0$, and therefore

$$x(s) = x_0/(1+s) \tag{14.2}$$

To recap, the above equation means that after X checks, if Y bets some amount s , X will call with hands better than $x_0/(1+s)$. At this point, we don't know what the value of x_0 is, only that X will never call with hands worse than that, no matter how little Y bets. We will solve for this value later.

Next we turn to Y's strategy.

In the same way, Y's value betting strategy will be a function of s , the bet amount that returns the hand value at which Y should bet a particular s .

If Y bets s for value at a particular hand value y , he will lose s when X holds a hand better than y (this occurs with probability y) and he will win s when X holds a hand better than x 's calling threshold $x(s)$ but worse than y . This occurs with probability $(x(s) - y)$. Then his equity is:

$$\langle Y \rangle = (\text{chance that Y gets called and loses}) (-s) + (\text{chance that Y gets called and wins})(s)$$

$$\langle Y \rangle = (y)(-s) + (x(s) - y)s$$

$$\langle Y \rangle = sx(s) - 2sy.$$

Again using calculus, we can maximize this value:

$$sx^2(s) + x(s) - 2y = 0$$

$$s(-x_0/(1+s)^2) + x_0/(1+s) - 2y = 0$$

$$x_0/(1+s)^2 - 2y = 0$$

$$y(s) = x_0/2(1+s)^2$$

This is Y's value betting strategy. Y's strategy bets different amounts at different Y values - in fact, bets a different amount at every Y value! The principle of information hiding would usually indicate that this is a terrible mistake because it can telegraph Y's hand. However, it turns out that in this particular game, X cannot make use of this information in an exploitive way, Y will

match every value bet with a corresponding bluff in the proper ratio. Then, when faced with a bet of s , X can either call with confidence (when he knows that he has the best hand), call with uncertainty (knowing that Y either has a bluff which he can beat or a value bet which he cannot), or fold. Because X cannot raise, Y gives up no useful information by limiting the strength of his hand; X would always call with hands that beat Y.

Now the next part - matching up the bluffing region with the value betting region - is a little tricky; it again involves calculus. If you've understood what's happened up to this point, skipping the following section is fine. What we've done so far is show that the strategies for this game are as follows:

On $[0, x(0)]$, X calls by his function $x(s)$. On $[x(0), 1]$, X folds.

On $[0, y(0)]$, Y value bets by his function $y(s)$. On $[y(0), y'(0)]$. Y checks behind. And with worse hands, Y bluffs according to $y'(s)$.

In the following section, we will solve for the value of x_0 and show the relationship of the betting and bluffing region for Y.

We have a known value betting function:

$$y(s) = x_0/2(1 + s)^2$$

We can then construct a value betting differential; effectively the size of one infinitesimal movement along that function. Intuitively, this might be thought of as the size of the region in which we bet one particular amount. The "dy" is the movement of Y per every "ds" movement in s.

$$dy = x_0/(1 + s)^2 ds$$

It turns out that our familiar bluffing ratio of $s/(1 + s)$ applies to this differential.

So we have a bluffing differential:

$$sx_0/(1+s)^4 ds$$

The total size of our bluffing region is then:

$$\int_0^{\infty} \frac{sx_0}{(1 + s)^4} ds$$

Making the substitution $t = 1 + s$, we have:

$$\begin{aligned} & x_0 \int_{t=1}^{\infty} \frac{t-1}{t^4} dt \\ &= x_0 \int_{t=1}^{\infty} \frac{1}{t^3} - \frac{1}{t^4} dt \\ &= x_0 [(-1/2 t^2) + (1/3 t^3)] | 1^{\infty} \\ &= 1/6 x_0 \end{aligned}$$

The total size of the bluffing region is thus $\frac{1}{6}x_0$. Since the bluffing region begins immediately after x_0 , we have:

$$x_0 + \frac{1}{6}x_0 = 1$$

So:

$$x_0 = \frac{6}{7}$$

Using the technique above, we can likewise map the bluffing region to the value betting region for Y; we leave this as an exercise for determined mathematicians. Instead, we return and state our strategy functions (Y's bluffing strategy as a function of s is omitted, but he bluffs some amount s with a total of $\frac{1}{7}$ of hands):

$$x(s) = \frac{6}{7}(1+s)$$

$$y(s) = \frac{3}{7}(1+s)^2$$

The important thing here, however, is not the particular strategies themselves, but the technique that we use for finding optimal bet sizes. In each of these two half-street games the AKQ and the [0,1]), we solve the game using variables to stand in for the unilateral strategy choices - in this case, the size of the bet that Y will make. Then the player who controls the variable selects its value such that his expectation is maximized.

We can learn many useful lessons from half-street games; over the past four chapters we have investigated a range of games:

- The clairvoyance game, which illustrated betting and bluffing ratios.
- The [0,1] Game #1, which introduced the concept of strategy regions and thresholds.
- The [0,1] Game #2, which showed the application of betting and bluffing ratios when the distributions are continuous and symmetrical.
- The [0,1] Jam-or-Fold games, which showed how uncertainty in the outcome of the showdown makes preflop jam-or-fold play much looser.
- The Jam-or-Fold no-limit game, which is a real-life solution to a game of some importance.
- The AKQ limit game, which again showed betting and bluffing ratios as well as introducing the idea of card removal.
- The AKQ no-limit game, which demonstrated the idea of calculating optimal bet size.
- The [0,1] game, which showed the much more complicated types of betting functions that arise from continuous distributions.

Key Concepts

- Bet sizing can be attacked analytically. To do this, we can first solve games for the strategies in terms of the bet size under consideration. Then the player who has control of the bet size maximizes his equity through his choice of a bet size.
- When one player has a very strong distribution or clairvoyance, situations can arise where the other player should simply give up no matter how small a bet the clairvoyant player makes.
- Information hiding is especially important when different bet sizes are available.

Chapter 15

Player X Strikes Back: Full-Street Games

In the first part of Part III, we discussed half-street games. In these games, Player X acted first, but was compelled to check dark. After this Player Y could either check and take a showdown, or value bet and bluff appropriately. In Chapter 15, we begin our investigation of full-street games. In these games, Player X again acts first, but now can either bet or check. If he bets, Player Y will sometimes have the option to raise or fold, and can always call. If he checks, then Player Y can check or bet.

It is important to recognize that some concepts will carry over from half-street games, while others will not. For example, betting and bluffing ratios are just as valid for Player X as they were for Y. But in addition to this, the option that X has to bet ***changes the distribution of hands he has when he checks***, so Y's play behind his checks will frequently change. This is a feature of virtually all expanded toy games - the play on earlier decisions significantly impacts the play on the later ones.

A second important idea is that ***Player X cannot lose value moving from a half-street game to a full-street one***, assuming that the rest of the game is the same. If Player X's strategies are such that he has less expectation in the full-street game than the half-street one, he should simply adopt the strategy previously forced upon him of checking in the dark. This is a natural extension of the idea that strategic options have non-negative value - X's ability to bet right off the bat can at worst be worth zero, since he can always decline the option and take his half-street equity.

The third and most important idea that applies to full-street games, and to any game of more complexity than a half-street game, is the concept and the way in which a strategy is expressed. In half-street games, particularly the [0,1] game, each player had only one action. Player Y could bet or check, and if Player Y bet, Player X could call or fold. Since each of their strategies was fundamentally just a single action in the game, there was a direct relationship between the ***strategy*** "fold" and the ***strategic option*** "fold."

In the half street games, each player's strategy was made up of regions - sometimes these regions were a range of hands with which the player would use a particular strategic option. This occurred in games like the [0,1] game. In other cases, the regions were mixing frequencies for the discrete hands - such as in the AKQ game. But in either case, the player selected a distribution of hands with which to take each strategic option, and the total of those distributions comprised the strategy.

In full-street games, though, the situation is slightly different, than in half-street games. Player X's initial check was required by the rules of the game. But now he has what might be called two different decisions - to check or bet initially, and to call or fold to a bet if he is faced with that scenario. In our view, these are not really two separate decisions. Instead, when X chooses a strategy, we create a new set of more complex options that reflect his choices for the game. Instead of two options on the first decision and two on the second, X simply has three decisions - to bet, to check-call, or to check-fold. When we specify decisions for full-street games, we will create payoff matrices that reflect the expanded strategic options and not simply the immediately available ones.

With these ideas in mind, we introduce our first full-street game.

Example 15.1 - AKQ Game #4

One full street.

Pot size of P units. (P > 1 for the same reason as in AKQ Game #1 - Example 13.1.)

Limit betting of 1 unit.

Raising is not allowed.

The complete ex-showdown payoff matrix for this game as follows:

	Y	Ace				King				Queen			
		X bets		X checks		X bets		X checks		X bets		X checks	
X		Call	Fold	Bet	Chk	Call	Fold	Bet	Chk	Call	Fold	Bet	Chk
Ace	Bet					-1	0			-1	0		
	Chk/Call							-1	0			-1	0
	Chk/Fold							+P	0			+P	0
King	Bet	+1	-P							-1	0		
	Chk/Call			+1	0							-1	0
	Chk/Fold			0	0							+P	0
Queen	Bet	+1	-P			+1	-P						
	Chk/Call			+1	0			+1	0				
	Chk/Fold			0	0			0	0				

We can again immediately remove dominated options from this game to create a simpler game. For Player X, we can see that checking and folding an ace is dominated by checking and calling with it, and checking and calling with a queen is dominated by checking and folding it. For Player Y, calling a bet with a queen or folding to a bet with an ace are dominated. Additionally, checking an ace behind a check is dominated by betting it. Removing these leaves the following game:

	Y	Ace		King				Queen		
		Call	Bet	Call	Fold	Bet	Chk	Fold	Bet	Chk
Ace	Bet			-1	0			0		
	Chk/Call					-1	0		-1	0
King	Bet	+1						0		
	Chk/Call		+1						-1	0
	Chk/Fold		0						+P	0
Queen	Bet	+1		+1	-P					
	Chk/Fold		0			0	0			

In this simpler game, we can see that betting kings is dominated for both players, and so we have:

	Y	Ace	King		Queen	
		Call/Bet	Call/Chk	Fold/Chk	Fold/Bet	Fold/Chk
Ace	Bet		-1	0	0	0
	Chk/Call		0	0	-1	0
King	Chk/Call	+1			-1	0
	Chk/Fold	0			-P	0
Queen	Bet	+1	+1	-P		
	Chk/Fold	0	0	0		

Let us consider the strategy choices for X. X will check all his kings, and call and fold with them to make Y indifferent to bluffing a queen. In addition to this, he must choose a strategy for his aces and for his queens. When he has an ace, he gains by check-calling when his opponent has a queen by *including a bluff*. However, he gains by betting when his opponent has a king and his opponent will check-call. With queens, he gains by bluffing out his opponent's kings and loses

when his opponent holds an ace or a king that will call.

For Y, the picture is a little simpler. He has only to choose strategies for his kings and for his queens. With a king, he will call often enough to make X indifferent to bluffing with a queen. With queens, he will bluff often enough to make X indifferent to calling with a king.

The strategy for X, then, will consist of three values:

A betting frequency with aces, which we will call x .

A calling frequency with kings, which we will call c .

(Note that X checks with all his kings, so this is his check-calling frequency)

A bluffing frequency with queens, which we will call b .

The strategy for Y will consist of just two values:

A calling frequency with kings, which we will call c_y .

A bluffing frequency with queens, which we will call b_y .

Consider the situation when X bets. We know that X will only bet when he holds aces or queens. Y must make him indifferent to bluffing by calling with a particular frequency of kings, and this is true whether the first player bets all his aces or not. Just as in the half-street games, when a player bets a bifurcated set of aces and queens, his opponent must fold precisely a of his hands that beat bluffs. Using the same logic from AXQ Game #2 (Example 13.2), we can end that when X bets, Y must call with $(P-1)/(P+1)$ of his kings so that X is indifferent to bluffing queens. This ensures that Y will be folding α of his hands that can beat a bluff.

X, however, must make his ratio of bluffing queens to value betting aces such that Y is indifferent to calling with kings. What we do here is consider the difference in ex-showdown value between Y's two options. When Y calls X's bet with a king, X can hold two different hands. When X holds an ace, X gains one bet when Y calls compared to gaining zero bets if Y folded. When X holds a queen, Y gains not only the bet he called, but the entire pot - since the alternative for Y was to fold and allow X to steal the pot. Since Y wins $(P+1)$ bets by calling a queen and loses one bet by calling an ace. X must bet one queen for every $P+1$ aces.

This can be written as the following equation:

$$b/x = 1/(P+1) = \alpha \tag{15.1}$$

Again, we see that the fundamental a bluffing-to-betting relationship holds. This allows us to ignore the value b in the rest of our analysis, since we can easily find b by simply using the above equation.

Next, let us consider the situation where X checks. Here we have a similar situation to the one in Game #2, except that instead of X having all his aces, he might have bet x of them. Y will bet all his aces and α of his queens, using the same reasoning as we have used throughout. X has some remaining distribution of hands. To make Y indifferent to bluffing, X must fold α of the hands that beat a bluff that remain in his distribution.

If X checked all his aces (suppose $x = 0$), then he would hold equal parts kings and aces. He would then fold α of his remaining hands, selecting kings to fold first. We showed in AXQ Game #2 (Example 13.2) that the fraction of his kings he would fold is $(P-1)/(P+1)$. Here, however, his total number of hands has been reduced by x of his aces, so he will need to call with more kings.

The total number of hands that X holds at this point is:

Total hands = (all his aces) (fraction of aces he checked) + (all kings)

$$\text{Total hands} = \left(\frac{1}{2}\right) (1 - x) + \left(\frac{1}{2}\right)$$

$$\text{Total hands} = \frac{1}{2} - \frac{1}{2}x + \frac{1}{2}$$

Fraction to fold = (α) (Total Hands)

$$\frac{1}{2} (1 - c) = \alpha \left(\frac{1}{2} - \frac{1}{2}x + \frac{1}{2}\right)$$

$$1 - c = \alpha (2 - x)$$

$$c = 1 - 2\alpha + \alpha x \tag{15.2}$$

We can check this formula by looking at two specific values of x . First consider $x = 0\%$, in which c is:

$$c = 1 - 2\alpha + \alpha x$$

$$c = 1 - 2\alpha$$

$$c = 1 - 2\alpha \left[\frac{1}{(P+1)}\right]$$

$$1 - 2\alpha = (P + 1 - 2)/(P + 1)$$

$$\alpha = (P - 1)/(P + 1)$$

as we expected. Alternatively, suppose $x = 100\%$. If X bets all his aces, he needs to fold α of his kings to make Y indifferent to bluffing a queen. Our formula shows that $c = 1 - \alpha$. again matching our expectation.

With this equation, then, we know how both sides will play in all situations. However, we still don't know what the value of x is.

To recap the solution so far:

X bets x of his aces and αx of his queens.

X checks and folds the remaining queens and checks and calls with the remaining aces.

X checks and calls with $(1 - 2\alpha + \alpha x)$ of his kings and check-folds the rest.

Y calls with $1 - 2\alpha$ of his kings and all of his aces if X bets.

Y folds the remaining kings and all queens if X bets.

If X checks, Y bets all his aces and α of his queens, and checks all other hands.

This is a common occurrence when we solve games; often we will obtain optimal solutions except for some particular variable. Then the side that controls that variable will simply select it such that that side's expectation is maximized.

To find x , we must find the expectation of the game. We have six cases.

Case 1:

X: A, Y: K.

In this case, X will bet x of the time, and Y will call $1 - 2\alpha$ of the time, for a total value of $-x(1 - 2\alpha)$. When X checks, the value is zero.

Case 2:

X: A, Y: Q.

In this case, when X bets, the value of the game is zero. X will check and induce a bluff $1 - x$ of the time, and Y will bluff α of the time. The value is thus $-\alpha(1 - x)$.

Case 3:

X: K, Y: A.

Here X checks all the time. When Y bets, X calls $1 - 2\alpha + \alpha x$ of the time. So the value of this case is $1 - 2\alpha + \alpha x$ to Y.

Case 4:

X:K,Y:Q

Here X checks all the time. Y will bluff α of the time, and when X calls, Y will lose a bet. When X folds, however, Y wins the entire pot. We know that Y's total value from all his bluff must be zero, since X's optimal strategy makes him indifferent between bluffing and checking and the value of checking is zero. So the value of Case 4 plus the value of Case 2 must equal zero. That makes Case 4's value $\alpha(1 - x)$.

Case 5:

X:Q, Y: A

X bluffs αx of the time. Y also calls for a value of αx .

Case 6:

X: Q, Y: K

Again, we know that the total value to the bluffer must be zero from bluffing, since the optimal strategy makes him indifferent to bluffing and checking and the value of checking is zero. So the value of Case 6 must be the negative of Case 5, or $-\alpha x$.

We can view these values as a value matrix:

Hand Matchup	Value to second player
(A, K)	$-x(1 - 2\alpha)$
(A, Q)	$-\alpha(1 - x)$
(K, A)	$1 - 2\alpha + \alpha x$
(K, Q)	$\alpha(1 - x)$
(Q, A)	αx
(Q, K)	$-\alpha x$
Total	$(x(-1 + 3\alpha) + 1 - 2\alpha) / 6$

This last expression gives us the value of the game in terms of x . The following graph shows the

game value for values of P and x :

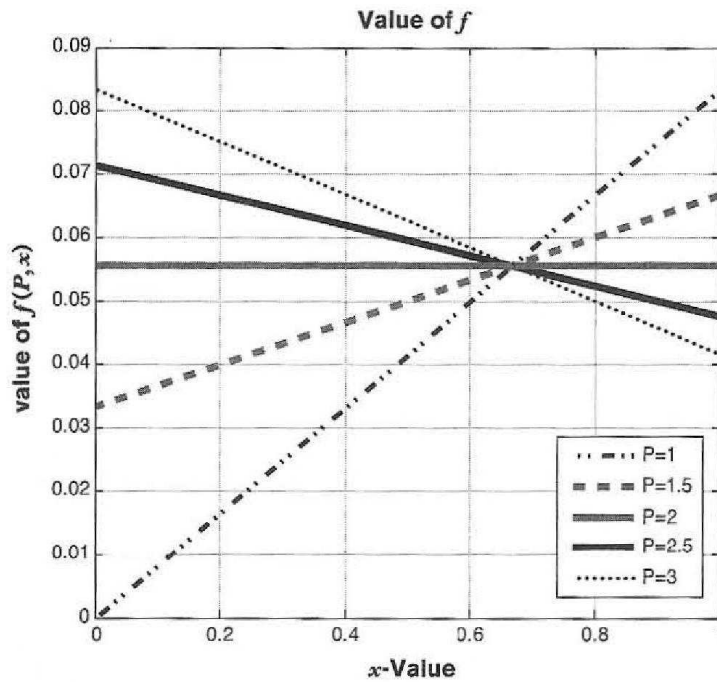


Figure 15.1. Equity of AKQ Game #4 for values of x .

The value expression for this game is minimized at $x = 1$ for values of P larger than 2 and at $x = 0$ for values of P between 1 and 2.

So we have the following optimal strategies:

Solution:

If $P < 2$:

Player X:

Checks all hands.

Calls with aces and $(P - 1)/(P + 1)$ of his kings.

Player Y:

Bets all aces and $1/(P + 1)$ of his queens.

If $P > 2$:

Player X:

Bets all aces and $1/(P + 1)$ of his queens.

Calls with $P/(P + 1)$ of his kings.

Player Y:

If the first player checks, bets all aces and $1/(P + 1)$ of his queens.

If the first player bets, calls with aces and $(P - 1)/(P + 1)$ of his kings.

In stating this game, we excluded raising as an option. However, we can verify based on the above solution that if raising is allowed, neither player can improve by raising any of his opponents' bets (except in the cases where he holds an ace and his opponent will fold a queen). Hence this is the solution to the more complete game where raising is allowed as well.

This simple game introduces full-street games. Essentially the differences between these and then half-street counterparts lie in the number of different strategic options available to the players. We saw the recurrence of the same betting to bluffing ratios (α) and the same folding ratio (α again) as we saw in the half-street game. The additional complexity came from X's new decision about whether to induce bluffs or to bet out with his aces.

We turn now back to the no-limit case. Recall that AKQ Game #3 (Example 14.1) was solved in a relatively similar way; we solved the game for strategies in terms of Y's bet size s and then maximized Y's equity through his selection of s . Having solved the half-street game, it would seem that the jump from the half-street no-limit game to the full-street game should involve only incremental complexity. At the end of this chapter, we proceed to doing just that.

However, before we go on and attempt to solve the AKQ for a full no-limit street, consider the following game.

Example 15.2 - AKQ Game #5

One full street.

Pot size of 4 units.

A special form of spread-limit betting. Either side may check (if no action yet), or bet or raise either 1 or 2 units at their turn.

AKQ Game #5 is the same as AKQ Game #4 (Example 15.1) with pot = 2, except that we've added one additional option for both sides; to bet $\frac{1}{4}$ of the pot instead half the pot. When we solved Game #4, we found that at $P = 2$, X was indifferent to checking or betting aces; hence we know that the value of the game is equal to the value of the half-street game, where X was forced to check dark, since X can check dark in this game without losing value.

In Example 15.1, we found that the value of the game to Y was:

$$f(1, P) = (x(-1 + 3\alpha) + 1 - 2\alpha) / 6$$

Here x is zero (as X can check dark) and $\alpha = 1/(2 + 1) = \frac{1}{3}$.

$$f(1, 2) = ((0)(-1 + 3\alpha) + 1 - 2(\frac{1}{3})) / 6$$

$$f(1, 2) = (1 - \frac{2}{3}) / 6$$

$$f(1, 2) = \frac{1}{18}$$

This result, however, is normalized to a bet size of 1. To correspond it to a bet size of two, we double it:

$$f(s = 2, P = 4) = \frac{1}{9}$$

This is the equity of the game if both players simply follow the strategy of playing as though the only allowed bet size is two units. We can now check to see if either player can benefit from betting 1 unit instead. Suppose that X bets 1 unit with his aces and bluffs appropriately with queens. Then Y will respond by making him indifferent to bluffing, which means calling with

$\frac{4}{5}$ of his hands overall, or $\frac{3}{5}$ of kings in addition to all his aces. By doing this, X will win 1 unit $\frac{3}{10}$ of the time, for a total win of $\frac{3}{10}$ of a unit. Contrast this to his equity when checking. If he checks, Y will bluff $\frac{1}{3}$ of his queens (which he holds $\frac{1}{2}$ of the time) for a win for the ace of two units, for a total win of $\frac{1}{3}$ of a unit. Since $\frac{1}{3}$ is larger than $\frac{3}{10}$, X unilaterally loses equity by betting 1 unit instead of two: hence this is not optimal.

If X checks dark, then Y's equity from changing his bet amount to one unit instead of two:

$$f(1, P) = (x(-1 + 3\alpha) + 1 - 2\alpha) / 6$$

So for $P = 4$:

$$f(1, 4) = (1 - \frac{2}{5}) / 6$$

$$f(s = 1, P = 4) = \frac{1}{10}$$

Hence Y cannot gain from betting one unit with aces and bluffing appropriately with queens. We also know that Y may raise in this game. However, since X would only bet with aces and queens, neither of which are hands that Y can profitably raise against, Y will just raise with his aces (because calling with them is dominated) and fold queens. The same reasoning applies to X. It seems, then, that this type of game is just a mild extension of the principles we have already studied.

Or is it?

Based on the previous analysis, we have the following strategies:

X:

Check all hands. (This is co-optimal with other strategies.)

Call a 2-unit bet with all aces and $\frac{1}{3}$ of kings. Call a 1 unit bet with $\frac{3}{5}$ of kings.

Y:

If X bets, call a 2-unit bet with aces and $\frac{1}{3}$ of kings and a 1 unit bet with aces and $\frac{3}{5}$ of kings.

If X checks, bet two units with all aces and $\frac{1}{3}$ of queens.

Recall now the important principle of optimal strategies that *neither player can improve their expectation unilaterally by changing strategies*. This means that if we can improve either player's expectation with any hand by playing it differently, the strategies aren't optimal.

Currently, X loses $\frac{2}{3}$ of a unit each time he has a king. Half the time, his opponent has a queen and bluffs $\frac{1}{3}$ of the time. X calls $\frac{1}{3}$ of the time, so $\frac{1}{3}$ of the time he gains two units. The other $\frac{2}{3}$ of the time, he loses the entire pot of four units, for a total loss when the opponent has a queen of $\frac{2}{3}$ of a unit. The other half of the time, the opponent has an ace. In these cases, he loses the bet of two units $\frac{1}{3}$ of the time (because he calls). So his total loss when the opponent has an ace is $\frac{2}{3}$ of a unit. The average of these, of course, is $\frac{2}{3}$.

If we changed his strategy to call more often with kings - for example, to call $\frac{1}{2}$ of the time - he would gain value against queens by not losing the pot when the queens bluffed, but lose some value by paying off aces. Since the bluffing to betting ratio for the second player is 3 to 1, the

total amount he loses with a king remains unchanged.

But what if we bet a king?

In the games we have considered so far, betting a king was dominated once we removed the first set of dominated strategies. We would never bet a king because to do so guarantees that we will lose a bet against the opponent's aces while gaining nothing against the opponent's queens. But consider now betting one unit with kings. Our strategy for Y indicates that Y should respond by calling with aces and $\frac{3}{5}$ of kings.

If Y does this, then X's total loss with a king is reduced to losing one unit half the time (when he gets called by an ace), for a total loss of $\frac{1}{2}$. This loss is smaller than the $\frac{2}{3}$ of a unit that X was incurring by checking all his kings. Therefore the strategies presented before cannot be optimal - X can unilaterally improve his equity by betting one unit with kings instead of checking them.

How can Y respond to this bet? We previously excluded raise-bluffing from the strategies, because there were no instances where a player would bet with a king. However, now we have such an occurrence, and so Y must consider raise-bluffing as a non-dominated option. When X bets and Y holds an ace, clearly raising dominates folding. But if he only raises with aces, he cannot gain value because X can simply fold his kings to a raise. So he must also protect his value raises with a corresponding ratio of bluffs.

However, the familiar ratio of s to $(1 + s)$ no longer applies, because the pot is now bigger (X made it bigger by betting, and the second player must match his bet before raising). There are also two bluff amounts that the second player can choose; a raise of two units (putting in a total of three units) or a raise of one unit (putting in two units total). At first glance, we might think that the ratio of bets to bluffs must make the first player indifferent to calling with a king after betting one unit. However, as we shall see shortly, bluffing that often is more costly for Y. Instead, Y will bluff often enough that X is indifferent to betting one unit with a king and folding to a raise or checking it.

X gained $\frac{2}{3} - \frac{1}{2} = \frac{1}{6}$ of a unit by betting one unit with his kings instead of checking them. So Y must bluff often enough to make up this difference. Let q be the bluff-raising frequency for the second player (for raising one unit more). Then the second player gains five units each time he bluffs a king successfully. Let q be Y's bluffing frequency.

$$\begin{aligned} (\text{pot})p(\text{Y has a queen}) (\text{bluffing frequency}) &= \frac{1}{6} \\ (5)(\frac{1}{2})q &= \frac{1}{6} \\ q &= \frac{1}{15} \end{aligned}$$

If the raise amount is two units, Y bluff-raises $\frac{1}{15}$ of the time with queens. This makes the first player indifferent to betting one unit or checking kings. The first player should still fold all his kings when the second player raises, because $\frac{15}{16}$ of the time, the second player will hold an ace. Notice that it doesn't matter for Y whether he raises one unit or two units when the first player holds a king because the king will fold anyway. However, if X holds an ace, Y would clearly prefer to have raised the smaller amount with a queen.

Now a new strategic option becomes profitable for X. He is indifferent to checking kings or betting them, but his expectation has changed for taking various actions with aces. He had been gaining $\frac{1}{3}$ of a unit with aces by checking and picking off Y's bluffs. (Y bluffs $\frac{1}{3}$ of his queens, two units, and has a queen $\frac{1}{2}$ the time). Now he has a new option, which is to bet one

unit with aces and an appropriate ratio of queens. This gains him some value from the second player calling with kings. He also gets to pick off some raise-bluffs from the second player.

His expectation from betting one unit with aces is:

$$\langle X, \text{bet } 1 \rangle = p(Y \text{ calls})(\text{bet}) + p(Y \text{ bluff-raises})(\text{raise amount})$$

$$\langle X, \text{bet } 1 \rangle = \left(\frac{3}{5}\right)\left(\frac{1}{2}\right)(1) + \left(\frac{1}{15}\right)\left(\frac{1}{2}\right)(2)$$

$$\langle X, \text{bet } 1 \rangle = \frac{11}{30}$$

His expectation from checking aces is:

$$\langle X, \text{check} \rangle = p(Y \text{ bluffs})(\text{bet size})$$

$$\langle X, \text{check} \rangle = \left(\frac{1}{2}\right)\left(\frac{1}{3}\right)(2)$$

$$\langle X, \text{check} \rangle = \frac{1}{3}$$

So he now gains $\frac{1}{30}$ of a unit by betting one unit with aces.

Let's review what's happened so far.

We set out to find the optimal strategies for this game with two different bet sizes. Initially, we planned to solve it in the same manner as the previous games, where kings were never bet and the rest was simply maximizing including betting to bluffing ratios. We found that both players ignored the smaller bet and bet aces and queens in ratios as though this were simply Game #4 with a pot of 2. But then a problem arose when we reviewed the strategy. We were able to increase X's equity by having him bet the smaller amount with kings, an option that was previously dominated (when only one bet size was allowed).

We then had to find a value-raising response that would enable Y to avoid being exploited in this way. We found a proper value-raise and corresponding bluffs that made X indifferent to playing this way with kings. But then a new strategic option became profitable for X. He can now bet the smaller amount with his aces and some queens, pick up most of the equity he would have made from checking by getting called by Y's kings, and additionally gain some equity from picking off raise-bluffs.

In all of these strategies, X will only bet queens enough to protect his aces (such that they get called by kings), and so that he will break even on his bluffs with queens.

We know that X has equal equity checking or betting kings (since the second player raise-bluffs enough to ensure this). We gain $\frac{1}{30}$ of a unit each time by betting aces, so X will want to bet all his aces. We must then find a ratio of kings to aces which forces Y to continue to raise-bluff. Too many kings, and Y profits by raise-bluffing more often; too few, and Y profits by raise-bluffing less often.

The raise bluff gains 5 units when it works (when X has a king), and loses two units when it fails (when X has an ace). So to make Y indifferent to raise-bluffing, X must bet aces and kings in a ratio of 5 to 2. Since X is betting 100% of his aces, he must bet 40% of his kings.

This leads us, then, to a new strategy:

Solution:

X:

Bet one unit with all aces. Bet one unit with 40% of kings. Bet one unit with 20% of queens.
 If Y raises one unit, reraise with aces and fold kings.
 If Y raises two units, reraise with aces and fold kings.
 With the remaining kings and queens, check.
 Call one unit bets $\frac{4}{5}$ of the time with kings and two unit bets $\frac{2}{3}$ of the time with kings.

Y:

If X bets two units, raise two units with aces and call with $\frac{1}{3}$ of kings. Fold queens.
 If X bets one unit, raise one unit with aces and $\frac{1}{15}$ of queens. Call with $\frac{3}{5}$ of kings.
 If X checks, bet two units with aces and $\frac{1}{3}$ of queens. Check kings.

The total value of this game is as follows:

P1 Card	P2 Card	Action	Value	Probability	EV Weight
A	K	bl - call	1	3/30	3/30
A	K	bl - fold	0	2/30	0
A	Q	bl - r1 - r - f	2	1/90	1/45
A	Q	bl - fold	0	14/90	0
K	A	bl - r1 - f	-1	1/15	-1/15
K	A	chk - b2 - c	-2	1/15	-2/15
K	A	chk - b2 - f	0	1/30	0
K	Q	bl - r1 - f	-5	1/225	-1/45
K	Q	bl - f	0	14/225	0
K	Q	chk - b2 - c	2	1/45	2/45
K	Q	chk - b2 - f	4	1/90	-2/45
K	Q	chk - chk	0	1/15	0
Q	A	bl - r2 - f	-1	1/30	-1/30
Q	A	chk - b2 - f	0	2/15	0
Q	K	bl - call	1	1/50	-1/50
Q	K	bl - fold	1	1/75	+1/75
Q	K	chk - chk	0	2/15	0
Total				1	-1/10

The value of the naive strategy (not involving betting a king) was $\frac{1}{6}$. The value of this strategy is $-\frac{1}{10}$, a difference of $\frac{1}{90}$.

It is interesting to note that this is a clear example of information hiding - if X knew his opponent's strategy was fixed, he would prefer to bet two units with his aces instead of one unit. However, with his kings, he prefers to bet the smaller amount. Since he has to choose a single bet size or be exploited, he chooses the smaller amount because it is more valuable (in this case) to make preemptive bets than to bet larger with aces.

Key Concepts

- It is important not to stop short of the solution when solving games. In this case, we went through the usual process, utilizing results from past games and so on, in order to solve this expanded game. At one point, it seemed we had found what looked like a reasonable solution. Only by going deeper and checking the solutions by comparing the equities of each option did we find an entirely new idea.
- Strategies are only optimal when both sides cannot improve! When we first discovered the above problem (while attempting to solve the full no-limit case), we were quite convinced that we had made a mistake; after all, betting a hand that has no chance of being called by a worse hand and no chance of folding a better hand seems foolish.
- This play with a king is quite out of character with what we see in many other circumstances; bets on the river are "supposed" to be value bets, intended to have positive expectation when they are called, or bluffs, intended to sometimes fold better hands. But this bet is neither. We call it a *preemptive bet*. Basically what X is doing is betting and preempting Y's ability to bluff with queens as effectively as he'd like to. Y can respond by raise-bluffing, but now walks into difficulty because X gets more value with his aces. This type of play exists when there are different bet amounts possible and the first player can bet a smaller amount than the second player wants to.
- Betting preemptively is a perfect example of a play that has negative expectation, but the expectation of the play is less negative than the alternative of checking. The important feature is that the value of the strategy is higher than the value of the alternative strategy. It is likely that no-limit holdem on the river contains preemptive betting as part of optimal strategy.
- One other important feature of this game is that after X makes a preemptive bet with a king, he folds all the time to a raise, even though the opponent is bluffing sometimes. This is because he has an excess of strong hands to pick off bluffs with and does not need to call raise bluffs as often. In addition, the raise-bluffer does not attempt to make the original bettor indifferent to calling with a king; instead the raise-bluffer only attempts to make X indifferent to betting a king (vs. checking) in the first place.

With these concepts in mind, we turn our attention now to the full no-limit AKQ game. Because of the complexity and difficulty of this game, we have created an Appendix to Chapter 15, which contains the solution and justification. *If you understood how the spread-limit solution was derived, then you have understood the most important lesson of the no-limit game.* At the beginning of the appendix, there is a section that summarizes the steps to find the solution and presents the solution without any justification. The remainder of the appendix contains the mathematics behind the solution, which will likely be of interest primarily to those with a mathematical background.

Appendix to Chapter 15

The No-Limit AKQ Game

Summary

The solution to the full no-limit AKQ game is surprisingly close to the solution to the spread-limit game from Chapter 17. The steps in deriving the strategy are:

- We can first solve a full-street game where X may not bet kings (or is known not to bet them). We find that X bets r (as was the case for Y in the half-street game) with some aces and checks the remaining aces, while Y bets a different, larger amount s_0 if checked to. Both players balance their strategies with aces and queens appropriately.
- X can then improve on his strategy if we relax the constraint on betting kings by betting a smaller amount t with some kings.
- Y responds to X betting t with kings by raising to another value $s(t)$ and raise-bluffing with queens with some frequency β . By doing this he makes X at most indifferent to checking or betting an amount t with kings. β is related to s_0 , the amount that Y bets when X checks.
- X responds by betting t with some aces in addition to kings and gaining value by picking off raise-bluffs and getting called by Y's kings. Y cannot prevent X from gaining value in one of these two ways.
- However, in the no-limit case. X cannot simply bet all his aces (as he did in the spread-limit case). If he does, and leaves himself with only kings in his checking distribution, Y can exploit him by betting a large amount with aces and some queens. Since he has only kings, he must call sometimes, which is more costly than the amount he can gain by betting aces.
- We can construct a function $A(t)$, which is the value of betting t with aces. This is the sum of the values of being called by kings and picking off raise-bluffs when the opponent holds queens.
- For any particular s_0 value, we can find the optimal t value, by maximizing $A(t)$. We call this value τ . This is the amount that X should bet with some aces and some kings.
- Since he will be playing a mixed strategy of betting some aces and checking some aces, he must have the same expectation from checking and betting τ ; that is $A(0) = A(\tau)$.
- We can then find a raise amount t_2 that Y would like to raise to. Unfortunately, the rules force Y to raise to at minimum $2t$.
- We then balance the strategies as appropriate with bluffs, calls, and raise-bluffs such that each player is indifferent to changing various parts of his strategy.

The final strategies are: (all decimal values are approximate, and all percentages are "of the hands he has played this far in this way")

Solution:

X:

Checks 61.2% of his aces, and if Y bets behind, X raises all-in.

Bets $\tau = 0.2521$ with the rest of his aces, and reraises all-in if Y raises.

Bets $\tau = 0.2521$ with 15.62% of his kings and folds to any raise.

Checks the remainder of his kings and calls a bet of s with an appropriate

fraction of kings such that his total calling or raising frequency with hands he checked is $1/(1 + s)$.

Bets $\tau = 0.2521$ with 7.81% of his queens and folds to any raise.

Checks and folds with the remaining 92.19% of his queens.

Y:

If X checks, Y bets $s_0 = 0.5422$ with all his aces and 35.16% of his queens.

If X check-raises, Y reraises all-in with his aces and folds his queens.

If X bets an amount t , Y responds as follows:

Define the following:

s_0 is Y's bet amount if X checks = 0.5422.

$\alpha_0 = s_0/(1+s_0)$

$t_2 = \alpha_0/(1 - \beta)$ (if this is greater than $2t$), otherwise $2t$

$\beta = (\alpha_0 - t)/(1 + t)$ if $\alpha_0 > t$, 0 otherwise.

Y raises to t_2 with all aces and with β of his queens. If X raises again, Y raises all-in with aces and folds queens.

Y calls with enough kings to make his total calling frequency with hands that can beat a queen $1/(1 + t)$.

Derivation

First we consider the game with the constraint that X may not bet kings. We can see immediately that if he bets with all aces, he can be exploited by Y betting a very large amount with all aces and an appropriate fraction of queens. Since he has bet all aces, he is left with only kings to defend against the opponent's bluffs. He must call with an appropriate fraction of kings, but gives up too much value to the opponents' aces to make this worthwhile.

To see this clearly, presume that X bets all his aces and checks all his kings. Then suppose Y bets 100 units (where pot = 1) with all his aces and 100/101 of his queens. If he does this X must call with $1/101$ of his kings to prevent Y from bluffing. This means he is giving up $100/101$ of a unit when he holds a king and X holds an ace. But he could improve by checking a few ace; and calling the bets which will all be bluffs). So betting all aces initially must not be optimal.

Since he must check all kings (by the rules of the game), and he must check some aces to provide a credible threat against the opponent's bluffing large amounts, he may either check all aces, or simply some fraction of aces. Call the fraction of aces he checks A_k . X will then bet some amount t with the aces he does not check and an appropriate fraction of queens. If he bets an amount t , then Y will respond by (at least) calling with a total fraction of his hands equal to $1/(1 + t)$. This will make X indifferent to bluffing queens. If X checks, then Y will bet an amount s . Thus we have three unknown elements of the strategy, in addition to the various easy-to-calculate betting and bluffing ratios.

For Y, the fraction of all calling hands that are kings is $1/(1 + t) - 1/2$ unless X bets an amount greater than 1, in which case he will call with all aces and no kings. X, then, has the following equity from betting t when he holds an ace and Y holds a king:

$$\langle A \text{ vs. K, bet } t \rangle = (t)(2/(1+t) - 1) = (t-t)^2/(t+1)$$

Maximizing this with respect to t :

$$\begin{aligned} (t+1)(1-2t) - (t-t^2)(1)/(t+1)^2 &= 0 \\ 1-t^2-2t &= 0 \\ t &= \sqrt{2} - 1 \end{aligned}$$

Recalling Equation 14.1:

$$t = r$$

Y will then call with $(t-1)/(t+1)$ of his kings as described above.

When X checks, we have a similar situation, Y will bet some amount s with his aces and appropriate queens. X will call with the aces that remain (remember, he bet some aces), and then enough kings to make Y indifferent to bluffing with queens. We know that X must be indifferent to betting r or checking with his aces because he will play a mixed strategy of these two options. The expectations of the two possible plays with an ace are:

$$\begin{aligned} \langle A, \text{check} \rangle &= (s(1+s))(s)(1/2) = s^2/2(1+s) \\ \langle A, \text{bet } r \rangle &= (r-r^2)/2(1+r) = 3/2 - \sqrt{2} \end{aligned}$$

Setting these equal:

$$\begin{aligned} 3/2 - \sqrt{2} &= s^2/2(1+s) \\ 3 - 2\sqrt{2} + 3s - 2\sqrt{2}s &= s^2 \\ s &\approx 0.50879 \end{aligned}$$

The third remaining unknown of the final strategy is the frequency A_k with which X checks. As we saw earlier, if A_k is not chosen properly, then Y can exploit X by bluffing larger or smaller amounts. Hence X will select A_k such that Y's equity is maximized by betting s .

When X holds a king and Y an ace, Y gains one bet (call this bet q) when X calls. X's calling frequency, in turn, depends on A_k . Recall that this value is the fraction of aces that remain. The total hands that X has remaining is $A_k(1/2) + 1/2$, or $1/2(1+A_k)$. If Y bets s , then X will call with a total fraction of hands (aces and kings) equal to:

$$(1/2)(1+A_k)(1/(1+q)) = (1+A_k)/2(1+q)$$

Of these, X will call with aces $1/2 A_k$ of the time, and then with remaining kings as necessary. The number of kings that X will call with is:

$$\begin{aligned} (1+A_k)/2(1+q) - 1/2 A_k \\ = (1+A_k - (A_k+qA_k))/2(1+q) \\ = (1-qA_k)/2(1+q) \end{aligned}$$

and Y's equity from being called by kings is this value times q , or

$$\langle K/A, \text{bet } q \rangle = (q - q^2 A_k)/2(1+q).$$

We can maximize this expression for Y:

$$(1 + q)(1 - 2qA_k) - (q - q^2A_k)/(1 + q)^2 = 0$$

$$L = 2qA_k + q^2A_k$$

Completing the square, we have:

$$(q + 1)^2 = 1/A_k + 1$$

$$q + 1 = \pm\sqrt{(A_k + 1)}/A_k$$

$$q = \sqrt{(A_k + 1)}/A_k - 1$$

Thus, in response to a choice of A_k by X, Y maximizes his equity by betting an amount q equal to this value. By setting $q = s$, we can find the optimal choice of A_k for X.

$$0.50879 = \sqrt{(A_k + 1)}/A_k - 1$$

$$A_k \approx 0.78342$$

So the final strategy for this simplified game (where X must check kings) is:

X checks aces about 78.3% of the time, bets r (about 41.2% of the pot) with the remainder, bets r times the pot with $r/(1 + r)$ times as many queens, or about 6.3% of his queens overall, and is forced to check all kings.

Y responds to a bet of t ($t > 0$) by calling with (at minimum) all aces, and if the bet size is small enough, with enough kings to make his total calling frequency (if the opponent is bluffing) equal to $1/(1 + s)$

If X checks, Y bets s (about 50.8% of the pot) with all aces and about 33.7% of his queens. X then calls with enough kings to make his total calling frequency (with the aces he checked and kings) equal to $1/(1 + s)$. For the optimal bet sizes, these calling frequencies are: Y calls an r sized bet with r of his kings, and X calls an s size bet with about 39.8% of his kings.

The value of this simplified game to Y is approximately 0.02761 units.

If we relax the "must check kings" restriction on the game, however, X can now improve. When X checks a king, we can use a shortcut to find his equity. Since we know that Y will make him indifferent to calling with a king by bluffing and betting in the proper ratio, we can assume he folds all the time. He loses one unit whenever a queen bluffs. Y will bluff an amount $s_0/(1 + s_0)$ of the time and have a queen half the time. We define $\alpha_0 = s_0/(1 + s_0)$. Thus his equity is:

$$\langle P1, \text{check K} \rangle = (1/2) [-s/(1 + s)] = -\alpha_0/2$$

Obviously if Y continues to respond to a bet of t by calling with aces and kings, X can bet a tiny amount and prevent Y from bluffing. So let us take as our candidate strategy that Y will respond to a bet of t with a raise to either $2t$ (the minimum allowed) or some other value t_2 , whichever is greater, with all aces and with some fraction of queens $\beta(t)$.

Then X's equity of betting t with a king is:

$$\langle P1, \text{bet } t/\text{K} \rangle = (1/2)(-t) + (1/2)(\beta) [-(1 + t)]$$

β must be large enough that this expression is smaller than the first expression; that is, that X's equity from betting t is less than or equal to his equity from checking. The cases where his equity from betting t is less than his equity occur when he bets too much with a king. For now we will concern ourselves only with reasonable preemptive bet sizes. (Later we will see that p can be zero if the bet size is large enough.) So we have:

$$\begin{aligned} (1/2)(-t) + (1/2)(\beta) [-(1+t)] &= -\alpha_0/2 \\ t + \beta t + \beta &= \alpha_0 \\ \beta &= (\alpha_0 - t)/(1+t) \quad (\alpha_0 > t) \end{aligned}$$

Recall that when X bets, Y calls with $(1-t)/(1+t)$ of his kings. If t_2 is the amount that Y raises to, then the equity of betting t with an ace is:

$$A(t) = (1/2)(t(1-t)/(1+t)) + (1/2)(\beta)(t_2)$$

If we substitute a minimum raise amount $2t$ for t_2 and the expression for β , we obtain:

$$A(t) \geq (1/2)(t(1-t)/(1+t)) + (1/2)((\alpha_0 - t)/(1+t))(2t) \quad (\alpha_0 > t)$$

Note that we changed the equality to an inequality because in some cases the raise t_2 might be larger than $2t$, which would make $A(t)$ larger. Simplifying this expression:

$$A(t) \geq \frac{t - t^2 + 2t\alpha_0 - 2t^2}{2(1+t)} \rightarrow \frac{t}{1+t} \left(\frac{1 - 3t + 2\alpha_0}{2} \right)$$

Maximizing this expression, we have:

$$A'(t) = (1/(1+t)^2)(1 - 3t + 2\alpha_0)/2 + (t/(1+t))(-3/2)$$

$$A'(t) = \frac{1 - 3t + 2\alpha_0 - 3(t + t^2)}{2(1+t)^2}$$

$$0 = 1 - 3t + 2\alpha_0 - 3t - t^2$$

$$3t^2 + 6t + 2\alpha_0 - 1 = 0$$

$$t^2 + 2t + 2/3\alpha_0 - 1/3 = 0$$

$$(t+1)2 = 2/3\alpha_0 + 4/3$$

$$t = \sqrt{4/3 + 2/3\alpha_0} - 1$$

This t value is therefore a special value; if Y will respond to a check with a bet of s_0 , then the optimal amount for X to bet with his kings is this value of t (related to s_0 , of course, because α_0 is a function of s_0). We call this value τ .

We know from previous discussion that X will be indifferent to checking or betting τ for a particular s -value because he will play a mixed strategy with aces. (If he does not, Y can exploit him by betting large amounts as we pointed out earlier). We can use this fact to find the proper s_0 value.

$$A(0) = A(\tau)$$

$$(1/2)(\alpha_0)(s_0) = (1/2)(\tau(1-\tau)/(1+\tau)) + (1/2)((\alpha_0 - \tau)/(1+\tau))(2\tau)$$

τ is a function of α_0 and in turn of s_0 , so we can solve for s_0 . Solving numerically, we obtain $s_0 \approx 0.54224$, $\alpha_0 \approx 0.35159$, and $\tau \approx 0.25209$.

Recall the previous formula for $\beta(t)$.

$$\beta = (\alpha_0 - t)(1+t) \quad (\alpha_0 > t)$$

It's clear that we cannot raise-bluff a negative percentage of the time. Hence β is zero for values of t greater than α_0 . Using this formula for β , we can also find the proper raise-bluffing amount; Y should raise an amount such that X is indifferent to calling with kings (if this is possible - otherwise he should raise-bluff the minimum). If the bet is t , and the raise-bluff is to t_2 , then we have the following indifference for X betting t with a king:

Second player's hand:	P1 calls with a king	P1 folds
Ace	$-(t_2 - t)$	0
Queen	t_2	$-(1 + t)$

$$\begin{aligned} t_2 - t &= \beta(t_2 + 1 + t) \\ (1 - \beta)t_2 &= \beta(1 + t) + t \\ (1 - \beta)t_2 &= \alpha_0 - t + t \\ t_2 &= \alpha_0 / (1 - \beta) \end{aligned}$$

This value t_2 is what we might call the *theoretical* value t_2 - the value to which Y would raise if there were no rule limitations on the game requiring that the raise be at minimum twice the original bet. We can immediately see that at bet amounts above α_0 , t_2 is lower than the bet amount; this means that Y need not raise and raise-bluff at all in order to ensure his equity in the game if X bets this value or higher. Also, at t values that are very small, Y cannot simply raise to $2t$; by doing so, he loses value against the opposing kings. Instead he raises to a value close to s_0 . There is a point close to 0.2001 where the theoretical value is equal to the minimum raise value.

$$\begin{aligned} s(t) &= \alpha_0 / (1 - \beta) & (t < 0.2001) \\ s(t) &= 2t & (t > 0.2001) \end{aligned}$$

We now have the outline of the optimal strategies. The only major variables left to find are the checking and betting frequencies of X with aces and kings. What happens here is that X checks aces and kings in a ratio to make Y bet s_0 when X checks, and bets aces and kings in a ratio to make Y indifferent to raise-bluffing the optimal amount. We can then combine these two ratios to find the precise frequencies of each action.

When X checks, he will call $1/(1 + s_0)$ of the time with all his hands to make the opponent indifferent to bluffing with queens. Call k the fraction of checking hands that are kings. If Y bets q with an ace when X checks, then X will call with $1/(1 + q)$ of his hands total.

$1 - k$ of those hands will be aces, leaving $1/(1 + q) - 1 + k$ kings out of a total of k kings to start.

This amount is $(k - q + kq)/k(1 + q)$.

Y's equity from betting q with an ace when X checks is then:

$$(kq - q^2 + kq^2)/k(1 + q).$$

Maximizing this expression:

$$k(l + q)(k - 2q + 2kq) - k(kq - q^2 + kq^2)/k^2(l + q)^2 = 0$$

$$k - 2q + 2kq + kq - 2q^2 + 2kq^2 - kq + q^2 - kq^2 = 0$$

$$k - 2q + 2kq - q^2 + kq^2 = 0$$

$$k(1 + 2q + q^2) = 2q + q^2$$

Substituting the known value of s_0 , we get:

$$k \approx 0.57957$$

In a like manner, X has a ratio of aces and kings with which he bets t ; this ratio is designed to make Y's equity maximum by raise-bluffing the optimal amount. Y gains $1 + \tau$ from successfully bluffing, but loses his raise amount q when called. X, however, will not call with his kings! He can simply fold all his kings to the raise; Y can no longer exploit him by raising a smaller amount because the raise amount must be at least $2t$. He will call enough to make Y indifferent to raise-bluffing, or $(1 + \tau)/(1 + \tau + q)$ of the time.

So his ratio of kings to aces for betting is $q/(1 + \tau + q) \approx 0.28707$.

Utilizing these two ratios, we can then find the full strategy for aces and kings. Letting c be the total amount of hands checked and b the total amount of hands bet, we have:

$$0.28707b + 0.57957c = 1$$

$$0.71293b + 0.42043c = 1$$

$$b \approx 1.45595$$

$$c \approx 0.55405$$

$$A_k = 61.213\%$$

$$A_b = 38.787\%$$

$$K_k = 84.382\%$$

$$K_b = 15.618\%$$

And with that, we have the optimal strategy to the game.

The value of this game to Y is approximately 0.02682 units per play.

Chapter 16

Small Bets, Big Pots: No-Fold [0,1] Games

In Chapter 15, we solved our first full-street game, where both players have the option to bet. There, we returned to consider the AKQ game. In this chapter, we return to a game that we considered when we began looking at half-street games; the [0,1] game. The [0,1] game is among the most important games we consider, because it is the game that is most closely related to poker, primarily because of the many hand types and the necessity of utilizing betting regions.

Throughout this chapter, all of the games we consider will have a special rule that we first introduced in Chapter 11, which is that neither player may fold. The "neither player may fold" rule approximates play when the pot is extremely large. What this means for us here is that none of the strategies that follow contain bluffing; such strategies would be of no use, since there is no opportunity for folding.

No-fold games are exercises in extracting the proper amount of value from strong hands.

The first and simplest [0,1] full street game is where there is only one bet remaining (i.e., no raising is allowed). Hence, there is only one decision for each player. X must decide what hands to bet, and Y must decide what hands to bet if X checks. If either player bets, the other player must call.

Before going on to the solution, see if you can figure out the optimal strategies; this is a good method of improving intuition.

Example 16.1 - [0,1] Game #4

One full street.

Neither player may fold.

Only one bet is allowed.

In this game, there is only one bet available, so the number of possible betting sequences is three:

X bets-Y calls.

X checks - Y bets - X calls.

X checks - Y checks.

Solving the problem intuitively, X is going to bet at least some of the time. If he bets, he will never bet a hand that is worse than a hand he would check, since Y can never fold. If X checks, however, Y can bet all the hands that X would have bet, since he is sure to win (given that X checked). In addition, Y can bet the strongest half of the remaining hands. If X's hand is in that same region, then Y will break even on X's calls (overall), but if X's hand is the lower half of that region, then Y will win a bet.

To recap Y's strategy:

If X bets, then Y's action is forced (he must call).

If X checks, then Y responds by betting all hands that X would have bet, and the best half of the hands that he would check.

Given this strategy for Y, then, what is X's strategy?

Y gains value when he bets a hand after X checks. Otherwise, Y's strategy is forced. So to

prevent Y from exploiting him, X can simply bet all hands. If he does this, he guarantees that the game will have equal equity for both players (one bet will go in).

So the optimal strategy for this game is for X to bet every hand, and Y to call every hand.

Analytically, recall that we referred to thresholds earlier as the key components of the solutions to our strategies. In fact, even in this simple game we solve for the thresholds of each strategy.

We parameterize the optimal strategies for both sides as follows:

There is a threshold x_1 such that below this threshold, X will bet, and above this threshold, X will check.

There is a threshold y_1 such that below this threshold, Y will bet if checked to, and above this threshold, Y will check.

To make X's strategy optimal, X must make it so that Y cannot unilaterally improve his equity by changing y_1 , and vice versa. Each player, then, wants to play so that the other is indifferent between the two actions on either side of each of his threshold points. This indifference means that the equity of the two actions is equal. We showed earlier how we can solve games by application of this concept. We write indifference equations for each threshold point.

Recall that we are referring to ex-showdown expectation here; only the change in equity that occurs as a result of the postflop betting is considered.

The indifference at y_1 (where Y is indifferent between checking and betting):

X's hand	$p(\text{X's hand})$	<Y, bet>	Product	<Y, check>	Product
$[0, x_1]$	x_1	N/A	N/A	N/A	N/A
$[x_1, y_1]$	$y_1 - x_1$	-1	$(x_1 - y_1)$	0	0
$[y_1, 1]$	$1 - y_1$	+1	$(1 - y_1)$	0	0
Total			$x_1 - 2y_1 + 1$		0

As we did in Example 11.3 in Chapter 11, we sum the total of the "Product" columns, and set the weighted expectations equal:

$$0 = x_1 - 2y_1 + 1$$

$$2y_1 = x_1 + 1$$

We next consider the indifference at x_1 .

At x_1 (where X is indifferent between checking and betting):

X's hand	$p(\text{X's hand})$	<Y, bet>	Product	<Y, check>	Product
$[0, x_1]$	x_1	-1	x_1	-1	x_1
$[x_1, y_1]$	$y_1 - x_1$	+1	$(y_1 - x_1)$	+1	$y_1 - x_1$
$[y_1, 1]$	$1 - y_1$	+1	$(1 - y_1)$	0	0
Total			$1 - 2x_1$		$y_1 - 2x_1$

Setting the expected values equal and solving for y_1 .

$$1 - 2x_1 = y_1 - 2x_1$$

$$y_1 = 1$$

Substitution:

$$\begin{aligned} 2 y_1 &= x_1 + 1 \\ 2 &= x_1 + 1 \\ x_1 &= 1 \end{aligned}$$

This is the same answer we found intuitively; both x_1 and y_1 are located at 1. So X bets all hands, and Y never gets a chance to bet; he must simply call.

Example 16.2 - [0,1] Game #5

The next simplest [0,1] full-street game introduces the concept of raising for value and has the following parameters:


- One full street.
- Two bets left.
- Check-raising is not allowed.
- Neither player may fold.

In analyzing this game we will use a notation that we have alluded to and which should seem intuitive; the threshold between betting and raising, y_2 . There is no x_2 in this game, as due to the impermissibility of check-raise, X can never put in the second bet. However, when X bets, Y will raise with some fraction of the hands X bets.

In Game #4 (Example 16.1), we saw that X can prevent Y from exploiting him by simply betting all his hands. However, this strategy will not do here, as Y has the option of raising and X can improve on his equity with his very worst hands by checking and calling instead of betting.

The first step is to parameterize the solution:

$$0 < y_2 < x_1 < y_1 < 1$$



16.1 Strategy structure for [0,1] Game #5

What this parameterization means is that the optimal strategies will have the following form: X will bet some amount of hands. Y will raise with fewer hands than that. If X checks, Y will bet more hands than X would have.

We can then write indifference equations at each threshold value:

In our table, we can simplify the process by identifying only the difference between one EV and the other for each row. Then we can make all the differences add up to zero.

At y_2 (the indifference between Y calling and Y raising):

X's hand	Y calls	Y raises	Y difference	Product
$[0, y_2]$	-1	-2	1	y_2
$[y_2, x_1]$	+1	+2	-1	$-1(x_1 - y_2)$
Total				$2y_2 - x_1$

Recall again that Y's equity from calling or raising must be the same due to our principle of indifference. Likewise, the lengths of the intervals $[0, y_2]$ and $[y_2, x_1]$ represent probabilities because the hands were originally drawn from a uniform distribution.

$$2y_2 - x_1 = 0$$

$$y_2 = x_1 / 2$$

This indicates that Y will raise with half the hands that X bets. The reasoning for this is similar to the reasoning that we saw in the half-street game. In a no-fold game, if X holds some number of hands, Y can put in one additional bet. X must call and may not fold. Whether the bet Y puts in is the first or second bet, Y puts it in with half the hands that X can hold.

Even in cases where X can reraise, however, we find that in these no-fold games that there is a ratio between the hands that X can hold and the hands that Y will put in the next bet. We will use the variable R to represent the ratio between the two:

$$y_2 = Rx_1$$

This may seem like an unnecessary complication, but it will be prove useful. R is an important concept, representing the fraction of X's value betting hands that Y should raise. Here R is $1/2$, because X cannot reraise. In fact, this is always the case on the last bet - when a player cannot be reraised and the other player cannot fold, R is $1/2$.

At x_1 , we see the effect that Y's ability to raise has on X's betting thresholds. As betting exposes X to being raised, he is forced to bet a narrower distribution of hands.

For X at x_1 (the indifference between X checking and betting):

Y's hand	X check-calls	X bets	Difference	Product
$[0, y_2]$	-1	-2	+1	y_2
$[y_2, x_1]$	-1	-1	0	0
$[x_1, y_1]$	-1	-1	0	0
$[y_1, 1]$	0	+1	-1	$-(1 - y_1)$
Total				$y_2 - 1 + y_1$

Writing the indifference equation here results in an equation that lacks any x values at all; what it

states is that the size of the raising region of Y's optimal strategy must be equal to the size of the region he checks behind when X checks.

$$y_2 = 1 - y_1$$

At y_1 , we see the familiar pattern of Y betting behind X's check with $1/2$ of the hands that X would check there:

At y_1 (the indifference between Y checking and Y betting):

X's hand	Y check	Y bets	Y difference	Product
[0, x_1]	N/A	N/A	N/A	N/A
[x_1 , y_1]	0	-1	+1	$y_1 - x_1$
[y_1 , 1]	0	+1	-1	$-1(1 - y_1)$
Total				$2y_2 - x_1 - 1$

$$2y_2 - x_1 - 1 = 0$$

$$y_1 = (1 + x_1)/2$$

We now have three equations and three thresholds:

$$y_2 = Rx_1$$

$$y_2 = 1 - y_1$$

$$y_1 = (1 + x_1)/2$$

To solve, we set the first two expressions for y_2 equal and substitute the third expression for y_1 :

$$Rx_1 = 1 - (1 + x_1)/2$$

$$x_1 = \frac{1}{1 + 2R}$$

We know from earlier that $R = 1/2$. Substituting, we have:

Solution:

$$x_1 = 1/2$$

$$y_2 = 1/4$$

$$y_1 = 3/4$$

To verify that this solution is indeed optimal, we can examine the play and see if either player can improve by playing differently against the opponent's optimal strategy.

Consider X's play when this hand is x between 0 and $1/4$. His choices are to check and call or to bet. If he bets, then the following outcomes occur:

When Y is better than x , X loses two bets. $(p = x)$

When Y is worse than x but better than $1/4$, X gains two bets. $(p = (1/4 - x))$

When Y is worse than $\frac{1}{4}$ X gains one bet.

$$(p = \frac{3}{4})$$

$$\langle y, x \text{ bets} \rangle = (x)(2) + (\frac{1}{4} - x)(-2) + (\frac{3}{4})(-1)$$

$$\langle y, x \text{ bets} \rangle = 4x - \frac{5}{4}$$

If X checks a hand that is between 0 and $\frac{1}{4}$.

When Y is better than x, X loses one bet.

$$(p = x)$$

When Y is worse than x but better than $\frac{1}{4}$, X gains one bet.

$$(p = \frac{1}{4} - x)$$

When Y is worse than $\frac{1}{4}$ but better than $\frac{3}{4}$, X gains one bet.

$$(p = \frac{1}{2})$$

When Y is worse than $\frac{3}{4}$, X gains 0.

$$(p = \frac{1}{4})$$

$$\langle x, \text{check} \rangle = (x)(1) + (\frac{1}{4} - x)(-1) + (\frac{1}{2})(-1)$$

$$\langle x, \text{check} \rangle = 2x - \frac{3}{4}$$

This second value is greater than the first (Y's equity is greater) when:

$$2x - \frac{3}{4} > 4x - \frac{5}{4}$$

$$x < \frac{1}{4}$$

So we can see that for X's hands between 0 and $\frac{1}{4}$, betting yields higher equity than checking - hence, X cannot improve by switching strategies. We could go on and check all the regions, looking for hands where we can unilaterally improve one player's equity. However, since the strategies we found here are optimal, no such opportunities exist.

In this last section, we also touched on an important concept - *the game value*. Playing this game has value to Y, because of his positional advantage. In Chapter 18 we will look at game values in depth, especially comparing game values from different types of games to each other.

Example 16.3 - [0,1] Game #6 - The Raising Game

One full street.

Players are allowed to raise and reraise without limit and have arbitrarily large stacks.

Check-raise is not allowed.

Neither player may fold.

This game is similar to the previous game, except that instead of just two bets, we now allow the number of bets to grow forever. Of course, in our optimal strategies, one or the other player will eventually stop raising and call.

Our thresholds can be described as follows:

We have x_1 , which is the threshold where X is indifferent to checking and betting. Since check-raise is not allowed, when X checks, there is just one remaining threshold y_1 . We know from our discussion in the previous game that y_1 will be half the distance from x_1 to 1, because this is the last bet and X cannot raise. We also know from the previous discussion that X will never "bluff" by putting in the nth bet with a hand weaker than a hand that would not put in that bet.

When X does bet, however, we have an infinite series of thresholds, beginning with y_2 , the threshold between hands that will raise and hands that will simply call. To the left of y_2 , we have a threshold where X reraises x_3 , and to the left of that, a threshold where Y puts in the fourth bet,

y_4 , and so on.

So our parameterization looks like this:

$$0 < y_{2n} < x_{2n-1} < \dots < y_2 < x_1 < y_1 < 1$$

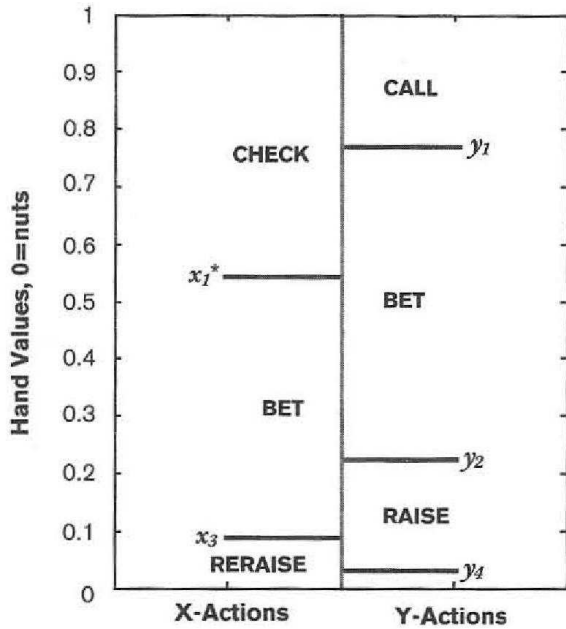


Figure 16.2 Strategy structure for [0,1] Game #6

This game has a sort of symmetry to it. First, consider the case where X bets. Y will raise with some fraction of his hands. Call this fraction (as before) R.

$$y_2 = Rx_1,$$

(just as in the previous game)

However, in this game, R is not $1/2$; X can reraise. Consider the situation once Y has raised. Y holds a hand in the range $[0, Rx_1]$. Now X wants to put in a raise with a fraction of the hands Y raised with. The situation for X here is the same as the situation for Y was on the previous bet - Y has a range, X wants to raise; his raising standard will be the same as Y's - he will raise with R of the hands that Y has.

$$x_3 = Ry_2$$

And now Y faces a bet from X, and will raise with R of the hands X has, and so on. Generally, then, for $n > 1$:

$$y_n = Rx_{n-1} \text{ (note that } y_3, y_5, \text{ etc., do not exist)}$$

$$x_n = Ry_{n-1} \text{ (note that } x_2, x_4, \text{ etc., do not exist)}$$

We have called [0,1] Game #6 "the raising game," because it contains the answer to a frequently asked question about poker - "how often should I raise my opponent if he can reraise me and no one can fold?" If we solve for R, we will know the proper frequency and be able to answer this question.

Next, consider the situation where X checks. We know that Y will bet with all the hands better

than x_1 and half the hands worse than x_1 :

$$y_1 = (1+x_1) / 2$$

Additionally, the same indifference at y_1 holds from the previous game (as there is no check-raising), leading to the size of Y's raising region y_2 being the same size as his checking region:

$$y_2 = 1 - y_1$$

Substituting:

$$y_2 = 1 - (1 + x_1)/2$$

$$2Rx_1 = 2 - 1 - x_1$$

$$x_1 = 1/(2R + 1)$$

Using this value, we can now find the values of all the thresholds y_2 , x_3 , y_4 and so on, in terms of R by multiplying by R the appropriate number of times. Additionally, we can find y_1 in terms of R .

But what is the value of R ? For this we need one more indifference equation:

For Y at y_2 (indifference between raising and calling):

X's hand	Y raises	Y calls	Y difference	Product
[0, x_3]	-3	-1	+2	$2x_3$
[x_3 , y_2]	-2	-1	+1	$y_2 - x_3$
[y_2 , x_1]	+2	+1	-1	$-1(x_1 - y_1)$
Total				$x_3+2y_2 - x_1$

$$x_3+2y_2 - x_1 = 0$$

This is an indifference equation that relates three thresholds. However, we have all these things in terms of R and x_1 (for convenience):

$$R^2x_1 + 2Rx_1 - x_1 = 0$$

$$R^2 + 2R - 1 = 0$$

$$R = \sqrt{2} - 1$$

Recall that we christened $\sqrt{2} - 1$ as "the golden mean of poker." (See Equation 14.1.) Here it appears again as the answer to our question: "How often should I raise an opponent who cannot fold and can raise me back?"

So the solution to our game is as follows:

Solution:

$$x_1 = 1/(1+2r) \approx 0.5469$$

$$y_1 = (1/(1 + 2r) + 1)/2 = (1 + r)/(1 + 2r) \approx 0.7735$$

$$y_2 = r/(1+2r) \approx 0.2265$$

$$x_3 = r^2/(1 + 2r) \approx 0.0938$$

$$y_n = r^{n-1}/(1 + 2r) \quad (\text{for even } n > 1)$$

$$x_n = r^{n-1}/(1 + 2r) \quad (\text{for odd } n > 2)$$

In our next game, we give Player X yet another option - check-raise.

Example 16.4 - [0,1] Game #7

One full street.

Two bets left.

Check-raising is allowed.

Folding is not allowed.

In Game #5 (Example 16.2), the two-bet case without check-raising, recall that there was no threshold x_2 . When check-raise was not allowed, x_2 was meaningless, as there was no sequence of betting that would lead to X putting in the second bet. However, when we give X the option to check-raise, we have the same game, but with one additional threshold for which we must solve.

The key to this game is to identify hands with which X wishes to check-raise. Clearly, X wants to put in the second bet with hands that are better than y_1 - obviously hands that are worse than this would lose value by check-raising. Clearly, X will not put the second bet in with a hand weaker than he would put in the first - we've shown that this is dominated already. Intuitively, it would seem that X would want to be the one to put in two bets with his strongest hands; although, it turns out that X's strategies are co-optimal as long as he chooses the proper region size for check-raising with hands better than y_2 . However, for simplicity, and because it aids us when more bets are available, we parameterize the game as follows:

$$0 < x_2 < y_2 < x_1 < y_1 < 1$$




Figure 16.3 Strategy structure for [0,1] Game # 7

This translates to the following strategies.

For X:

X checks and raises if Y bets on $[0, x_2]$.

X bets for value on $[x_2, x_1]$.

X checks and calls on $[x_1, 1]$.

For Y:

If X bets:

Y raises for value on $[0, y_2]$.

Y calls on $[y_2, 1]$.

If X checks:

Y bets for value on $[0, y_1]$.

Y checks on $[y_1, 1]$.

Creating indifference equations for the four thresholds:

For X at x_2 (indifference between check-raising and betting):

Y's hand	X check-raises	X bets	Difference	Product
$[0, y_2]$	-2	-2	0	0
$[y_2, y_1]$	+2	+1	+1	$y_1 - y_2$
$[y_1, 1]$	0	+1	-1	$-(1 - y_1)$
Total				$2y_1 - y_2 - 1$

$$2y_1 - y_2 - 1 = 0$$

$$y_2 = 2y_1 - 1$$

For Y at y_2 (indifference between raising and calling):

X's hand	Y raises	Y calls	Difference	Product
$[0, x_2]$	N/A	N/A	N/A	N/A
$[x_2, y_2]$	-2	-1	+1	$y_2 - x_2$
$[y_2, x_1]$	+2	+1	-1	$-1(x_1 - y_2)$
$[x_1, 1]$	N/A	N/A	N/A	N/A
Total				$2y_2 - x_1 - x_2$

$$2y_2 - x_1 - x_2 = 0$$

$$y_2 = (x_1 + x_2)/2$$

For X at x_1 (indifference between check-calling and betting):

Y's hand	X check-calls	X bets	Difference	Product
$[0, y_2]$	-1	-2	+1	y_2
$[y_2, y_1]$	-1	-1	0	0
$[y_1, 1]$	0	+1	-1	$-(1 - y_1)$
Total				$y_2 - 1 + y_1$

$$y_2 - 1 + y_1 = 0$$

$$y_2 = 1 - y_1$$

For Y at y_1 (indifference between betting and checking):

X's hand	Y raises	Y calls	Difference	Product
$[0, x_2]$	0	-2	+2	$+2x_2$
$[x_2, x_1]$	N/A	N/A	N/A	N/A
$[x_1, y_1]$	0	-1	-1	$-1(y_1 - x_1)$
$[y_1, 1]$	0	+1	+1	$1 - y_1$
Total				$2x_2 + 2y_2 - x_1 + 1$

$$2x_2 + 2y_2 - x_1 + 1 = 0$$

$$2y_1 = 1 + x_1 - 2x_2$$

We now have four indifference equations, and we can solve for the four unknowns:

$$2y_1 = 1 + x_1 - 2x_2$$

$$y_2 = (x_1 + x_2)/2$$

$$y_2 = 2y_1 - 1$$

$$y_2 = 1 - y_1$$

$$2y_1 - 1 = 1 - y_1$$

$$y_1 = 2/3$$

$$y_2 = 1/3$$

$$4/3 = 1 + x_1 - 2x_2$$

$$1/3 = (x_1 + x_2)/2$$

$$x_1 = 2/3 - x_2$$

$$4/3 = 1 + 2/3 - 3x_2$$

$$3x_2 = 1/3$$

$$x_2 = 1/9$$

$$x_1 = 5/9$$

Solution:

$$x_1 = 5/9$$

$$y_1 = 2/3$$

$$x_2 = 1/9$$

$$y_2 = 1/3$$

We can compare this game to Game #5 (Example 16.2), the similar game without check-raise. What would we expect to be different? First, we would expect Y to bet fewer hands when X checks. This is because by betting in Game #5, Y was only exposing one bet. However, in Game #7, Y exposes a second bet, because sometimes X will check-raise. So Y must bet fewer hands behind X's check.

But this, in turn, causes X a problem. In Game #5, X was checking and calling with hands higher than but near $1/2$. These hands were gaining check-call value because Y was betting hands as weak as $3/4$. In the new game, however, these hands of X's lose some of that value because Y

will no longer bet those mediocre hands for value. So X will likely bet additional hands worse than $1/2$. This in turn will allow Y to loosen up his raising threshold, since two things have occurred. First, X is checking some of his best hands, so Y does not have to fear X having those hands. Secondly, X is betting additional mediocre hands near $1/2$. Y can therefore raise for value with somewhat weaker hands in this game than in Game #5.

Inspecting the thresholds, we see that these changes are in fact reflected in the solution to the game:

Threshold	Game #5	Game #7
x_1	$1/2$	$5/9$
y_1	$3/4$	$2/3$
y_2	$1/4$	$1/3$

Our last full-street no-fold game is the analogue to Game #6 with check-raise: this game might be thought of as the full no-fold one-street $[0,1]$ game.

Example 16.5 - $[0,1]$ Game #8: The Raising Game with Check-Raise

One full street.

Players are allowed to raise and reraise without limit and have arbitrarily large stacks.

Check-raise is allowed.

Neither player may fold.

The final no-fold game in the one-street cases is the game where both players can put in as many raises as they think their hands merit and check-raise is allowed. Like Game #6 (Example 16.3), in this game we have an infinite number of thresholds, and so we cannot simply write all the indifference equations for this game. However, we can create a recursion for higher-valued indifference equations and solve that to find the values for each threshold and the solution to the game.

Based on the principles of the last few sections, we can parameterize the game this way:

For any positive integer n , $0 < x_{n+1} < y_{n+1} < x_n < y_n < 1$




Figure 16.4. Strategy structure for $[0,1]$ Game #8

We can begin by writing indifference equations at the various x thresholds:

At x_1 :

$$y_2 = 1 - y_1 \quad (\text{This is unchanged from Game \#6})$$

For X at x_2 (indifference between check-raising and betting):

Y's hand	X check-raises	X bets	Difference	Product
[0, y_3]	-3	-2	-1	$-y_3$
[y_2 , y_1]	+2	+1	+1	$y_1 - y_2$
[y_1 , 1]	0	+1	-1	$-(1 - y_1)$
Total				$-y_3 + 1 - 2y_1 - y_2$

$$-y_3 + 1 - 2y_1 - y_2 = 0$$

$$y_3 + y_2 + 1 = 2y_1$$

Substituting:

$$y_3 + 1 - y_1 + 1 = 2y_1$$

$$y_3 + 2 = 2y_1$$

For X at x_3 (indifference between bet-re-raising and check-raising):

Y's hand	X bet-re-raises	X check-raises	Difference	Product
[0, y_4]	-4	-3	-1	$-y_4$
[y_3 , y_2]	+3	+2	+1	$y_2 - y_3$
[y_2 , y_1]	+1	+2	-1	$-1(y_1 - y_2)$
[y_1 , 1]	+1	0	+1	$1 - y_1$
Total				$-y_4 - y_3 + 2y_2 - 2y_1 + 1$

$$-y_4 - y_3 + 2y_2 - 2y_1 + 1 = 0$$

$$y_4 + (y_1 - y_2) = (y_2 - y_3) + (1 - y_1)$$

For X at x_4 (indifference between check-raise-re-raising and bet-re-raising):

Y's hand	X bet-re-raises	X check-raises	Difference	Product
[0, y_5]	-5	-4	-1	$-y_5$
[y_4 , y_3]	+4	+3	+1	$y_3 - y_4$
[y_3 , y_2]	+2	+3	-1	$-(y_2 - y_3)$
[y_2 , y_1]	+2	+1	+1	$y_1 - y_2$
[y_1 , 1]	0	+1	-1	$-(1 - y_1)$
Total				$-y_5 + 2y_3 - y_4 - 2y_2 + 2y_1 + 1$

$$-y_5 + 2y_3 - y_4 - 2y_2 + 2y_1 + 1 = 0$$

$$(y_3 - y_4) + (y_1 - y_2) = y_5 + (y_2 - y_3) + (1 - y_1)$$

More generally, at x_n we have the following indifference equation when n is odd:

$$y_{n+1} + (y_{n-2} - y_{n-1}) + (y_{n-4} - y_{n-3}) [\dots] + [y_1 - y_2] = (y_{n-1} - y_n) + (y_{n-3} - y_{n-2}) + [\dots] + (1 - y_1)$$

And when n is even:

$$(y_{n-1} - y_n) + (y_{n-3} - y_{n-2}) [\dots] + [y_1 - y_2] = y_{n+1} + (y_{n-2} - y_{n-1}) + (y_{n-4} - y_{n-3}) + [\dots] + (1 - y_1)$$

Let's look at two actual thresholds, x_6 and x_7 .

Writing the indifference equations for these yields:

$$(y_5 - y_6) + (y_3 - y_4) + (y_1 - y_2) = y_7 + (y_4 - y_5) + (y_2 - y_3) + (1 - y_1)$$

$$y_8 + (y_5 - y_6) + (y_3 - y_4) + (y_1 - y_2) = (y_6 - y_7) + (y_4 - y_5) + (y_2 - y_3) + (1 - y_1)$$

Subtracting the second of these equations from the first:

$$y_8 = -2y_7 + y_6$$

It's not too difficult to see that we can do the same thing for any value of n , such as $k + 1$ and k :

$$y_{k+2} + (y_{k-1} - y_k) + (y_{k-3} - y_{k-2}) [\dots] + (y_1 - y_2) = (y_k - y_{k+1}) + (y_{k-2} - y_{k-1}) + (y_{k-4} - y_{k-3}) [\dots] + (1 - y_1)$$

$$(y_{k-1} - y_k) + (y_{k-3} - y_{k-2}) [\dots] + (y_1 - y_2) = y_{k+1} + (y_{k-2} - y_{k-1}) + (y_{k-4} - y_{k-3}) [\dots] + (1 - y_1)$$

Again subtracting:

$$y_{k+2} = -2y_{k+1} + y_k$$

$$y_{k+2} + 2y_{k+1} - y_k = 0 \tag{16.1}$$

Equations of this type are called *difference equations* and solving them requires a slightly more advanced mathematical technique. We present the derivation of the solution to Equation 16.1 in the Appendix to this chapter. The solution is that:

$$y_n = r^n / (1 - r) \tag{for } y > 0$$

This is Y's entire strategy, then, captured in a single equation, y_1 is $1/(1 - r)$, and Y puts in the $(n + 1)$ th bet with a fraction of the hands with which he put in the n th bet equal to r . Compare this with Y's strategy in the game without check-raise. There, the factors of r separated each alternating threshold, from X to Y. So y_n was rx_{n-1} , and so on. Here, instead, the factors of r separate each of Y's thresholds from the next Y thresholds, so we have:

$$y_n = ry_{n-1}$$

With Y's strategy solved, we move on to X's strategy, which is solved in a similar manner but contains a little more complexity. Just as Y's strategy involved a pattern for specifying y_n , we can specify x_n . However, Y's strategy was unaffected by regions in thresholds with n much lower than the current one. X, on the other hand, alternates his regions between check-raising and betting initially, and this is the case for thresholds as low as y_1 .

For Y at y_1 :

X's hand	Y check	Y bets
$[y_1, 1]$	0	+1
$[x_1, y_1]$	0	-1
$[x_3, x_2]$	0	-2
$[x_5, x_4]$	0	-2
$[x_7, x_6]$	0	-2
[...]		
$[x_{2k+1}, x_{2k}]$	0	-2

$$1 - y_1 = (y_1 - x_1) + 2(x_2 - x_3) + 2(x_4 - x_5) + 2(x_6 - x_7) + [...] + 2(x_{2k} - x_{2k+1}) \quad (k \rightarrow \infty)$$

For this threshold, when the decision is between betting and checking, Y gains a bet when X has checked out of weakness, but loses two bets whenever X is planning to check-raise.

For Y at y_2 :

X's hand	Y raises	Y calls
$[y_2, x_1]$	+2	+1
$[x_2, y_2]$	-2	-1
$[x_4, x_3]$	0	-2
$[x_6, x_5]$	0	-2
$[x_8, x_7]$	0	-2
[...]		
$[x_{2n}, x_{2n+1}]$	0	-2

$$x_1 - y_2 = (y_2 - x_2) + 2(x_3 - x_4) + 2(x_5 - x_6) + 2(x_7 - x_8) + [...] + 2(x_{2k} - x_{2k+1}) \quad (k \rightarrow \infty)$$

And more generally, for Y at y_n :

X's hand	Y raises	Y calls
$[y_n, x_{n-1}]$	+2	+1
$[x_n, y_n]$	-2	-1
$[x_{n+2}, x_{n+1}]$	0	-2
$[x_{n+4}, x_{n+3}]$	0	-2
[...]		
$[x_{2n}, x_{2n+1}]$	0	-2

$$x_{n-1} - y_n = (y_n - x_n) + 2(x_{n+1} - x_{n+2}) + 2(x_{n+3} - x_{n+4}) + [...] + 2(x_{2k} - x_{2k+1}) \quad (k \rightarrow \infty)$$

Returning briefly to y_1 we see that this equation is the general form of y_n for all thresholds, if we define $x_0 = 1$ (which makes intuitive sense, as X will never bluff in a game where Y cannot fold.)

Consider this equation for a particular k and also for $k + 1$:

$$x_{k-1} - y_k = (y_k - x_k) + 2(x_{k+1} - x_{k+2}) + 2(x_{k+3} - x_{k+4}) + [...]$$

$$x_k - y_{k+1} = (y_{k+1} - x_{k+1}) + 2(x_{k+2} - x_{k+3}) + 2(x_{k+4} - x_{k+5}) + [...]$$

Adding these two equations to each other, we find that the higher k terms cancel out neatly because of the alternating signs:

$$\begin{aligned} x_{k+1} - y_k + x_k - y_{k+1} &= y_k - x_k + y_{k+1} + x_{k+1} \\ x_{k-1} + 2x_k - x_{k+1} &= 2y_k + 2y_{k+1} \\ x_k + 2x_{k+1} - x_{k+2} &= 2y_{k+1} + 2y_{k+2} \end{aligned}$$

Substituting for y_n :

$$\begin{aligned} x_k + 2x_{k+1} - x_{k+2} &= 2(r^{k+1})/(1-r) + 2(r^{k+2})/(1-r) \\ r/(1-r) &= (\sqrt{2}-1)/(2-\sqrt{2}) = 1/\sqrt{2} \\ x_k + 2x_{k+1} - x_{k+2} &= \sqrt{2}(r^k) + \sqrt{2}(r^{k+1}) \\ x_k + 2x_{k+1} - x_{k+2} &= \sqrt{2}(r^k)(1+r) \end{aligned}$$

$$x_k + 2x_{k+1} - x_{k+2} = 2r^k \tag{16.2}$$

This equation, too, can be solved using a more advanced mathematical technique, and the steps of the derivation are presented in the Appendix to this chapter. The solution to this equation yields X's strategy:

$$x_n = r^{n-1}[(2r-1)(-1)^n + 1]/2$$

And so we have a recursive relationship that completely describes X's strategy.

Solution:

$$x_n = r^{n-1}[(2r-1)(-1)^n + 1]/2$$

$$y_n = r^n/(1-r)$$

The contrast between these two strategies is quite interesting: Y's strategy is relatively straightforward, with a simple exponential relationship between each threshold. Y puts in the first bet if checked to with hands below $r/(1-r)$, and is willing to put in one additional bet with each multiple of r of his hands if X keeps raising.

Each X threshold lies between each of Y's thresholds, beginning with x_1 at $1-r$. However, these thresholds are not evenly spaced between Y's thresholds. X's check-raising regions are much narrower than his betting regions. This occurs in the formulas as the term in x_n with the $(-1)^n$ in it. The intuitive explanation for this is that when X has a very strong hand, he often wants to bet, because he gets a very large amount of value from the sequence "bet-call." As a result, X's thresholds oscillate in their relationship to Y's thresholds.

A partial table of the thresholds in the game

Threshold (n)	y_n	x_n	y_n/x_n
0	1	1	1
1	0.707107	0.585786	$r/(1-r) + 1/2$
2	0.292893	0.171573	$r/(1-r) + 1$
3	0.12132	0.100505	$r/(1-r) + 1/2$
4	0.050253	0.029437	$r/(1-r) + 1$
5	0.020815	0.017244	$r/(1-r) + 1/2$

$$x_n = r^{n-1}[(2r - 1)(-1)^n + 1]/2$$

The portion in brackets evaluates to $2r$ when n is even, and to $2(1 - r)$ when n is odd.

So $x_n = r^n$ when n is even and $r^{n-1}(1 - r)$ when n is odd.

We can find something interesting if we look at the size of x 's checking regions:

$$\begin{aligned} x_0 \rightarrow x_1 &: 1 - (1 - r) = r \\ x_2 \rightarrow x_3 &: r^2 - r^2(1 - r) = r^2(1 - (1 - r)) = r^3 \\ x_4 \rightarrow x_5 &: r^4 - r^4(1 - r) = r^4(1 - (1 - r)) = r^5 \\ x_{2n} \rightarrow x_{2n+1} &: r^{2n} - r^{2n}(1 - r) = r^{2n+1} \end{aligned}$$

Summing the size of all these regions, we get the following expression

$$r \sum_{n=1}^{\infty} (r^2)^n$$

r^2 is less than 1, so the sum of this infinite series is $r(1 / (1 - r^2))$.

$$\begin{aligned} r^2 &= 2 \cdot 2\sqrt{2} + 1 = 3 - 2\sqrt{2} \\ (\sqrt{2} - 1)/(1 - (3 - 2\sqrt{2})) &= 1/2 \end{aligned}$$

This intriguing result indicates that X checks half his hands overall. Most of these hands, of course, will reluctantly call (in fact, r of them). The remainder will check-raise and be ready-to put in additional bets as appropriate.

Key Concepts

- In full-street games with no folding, there is an important relationship between the hands with which the opponent could have taken an action, and the hands with which you will take your actions.
- If you cannot be raised and your opponent cannot fold, you should bet or raise with $1/2$ of the hands your opponent has bet or checked so far.
- If your opponent *can* raise you back, this ratio shrinks. In the case where an infinite number of raises are allowed, it becomes r .
- No-fold games are exercises in extracting the proper amount of value from strong hands.

Appendix to Chapter 16

Solving the Difference Equations

Equation 16.1:

$$y_{k+2} + 2y_{k+1} - y_k = 0 \quad (16.1)$$

This equation is called a *second-order homogeneous difference equation* and describes the relationship between three consecutive y thresholds.

All equations of the form

$$y_{n+2} + cy_{n+1} + dy_n = 0$$

where c and d are constants have a solution of the form:

$$y_n = ap^n + bq^n \quad (16.1a)$$

where a and b are constants and p and q are the real distinct roots of the "characteristic equation" of the original difference equation, which is $x^2 + cx + d = 0$. If the roots are the same, there is a different solution, but this is irrelevant to our solution.

For this difference equation, the characteristic equation is

$$x^2 + 2x - 1 = 0$$

or $x = -1 \pm \sqrt{2}$

We recall that $\sqrt{2} - 1$ is the "Golden Mean of Poker" we discussed in Chapter 14 (see Equation 14.1). We substitute the constant r for that value.

Returning to Equation 16.1a, we substitute the two roots for p and q , and have:

$$y_n = ar^n + b(-1 - \sqrt{2})q^n$$

Using a bit of intuition, we know that as $n \rightarrow \infty$, y_n converges on 0, as Y will need a hand closer and closer to zero to put in additional bets as n becomes large.

The second term of this equation, however, does not converge on 0 for any value of b except 0. Therefore $b = 0$ and the equation simplifies to:

$$y_n = ar^n$$

To find a , we simply plug back in to the equation at the threshold x_1 :

$$y_2 = 1 - y_1$$

$$ar^2 = 1 - ar$$

$$ar^2 + ar - 1 = 0$$

We could solve this as a quadratic in r , but instead let us make use of the following identity:

$$r^2 = (\sqrt{2} - 1)^2$$

$$r^2 = 2 - 2\sqrt{2} + 1$$

$$r^2 = -2(\sqrt{2} - 1) + 1$$

$$r^2 = 1 - 2r$$

$$a(1 - 2r) + ar - 1 = 0$$

$$a - ar - 1 = 0$$

$$a(1 - r) - 1$$

$$a = 1/(1 - r)$$

Hence:

$$y_n = r^n/(1 - r)$$

Equation 16.2:

$$x_k + 2x_{k+1} - x_{k+2} = 2r^k \tag{16.2}$$

This equation, like Equation 20.1, is a second-order difference equation. However, because the x_k terms of this equation sum to a value other than zero, this type of equation is called **non-homogeneous** and this adds an additional step to solving it. If Z is a "particular solution" - that is, a solution for some k value, then all solutions to the general equation will be the "general solution," which is the solution to the homogeneous version of the equation, summed with Z.

$$x_n = ap^n + bq^n + Z \tag{16.2a}$$

We now solve the homogeneous form of the difference equation to find the non-Z part of the equation:

$$1 + 2x - x^2 = 0$$

$$x = 1 \pm \sqrt{2}$$

$$p = 1 - \sqrt{2} = -r$$

$$q = 1 + \sqrt{2}$$

$$x_n = a(-r)^n + b(1 + \sqrt{2})^n + Z$$

We can now guess at a form for the particular solution. We conjecture that Z has the form:

$$Z = cx^n$$

for some constant c and some constant x.

If we substitute our particular solution Z back into equation 16.2, we obtain:

$$x_k + 2x_{k+1} - x_{k+2} = 2r^k \tag{16.2}$$

$$cx^n + 2cx^{n+1} - cx^{n+2} = 2r^n$$

$$cx^n(1 + 2x - x^2) = 2r^n$$

Here $(1 + 2x - x^2)$ is a constant, so we find that for a specific value of c, $x = r$ is a solution:

$$cr^n(1 + 2r - r^2) = 2r^n$$

$$c(1 + 2r - (1 - 2r)) = 2$$

$$c(4r) = 2$$

$$c = r/2$$

So our particular solution Z will be:

$$Z = (r/2)r^n$$

We can see that as n goes to infinity, Z converges on zero. Likewise, by the same reasoning we used in solving 16.1, we know that the second term of the general solution only converges on z for $b = 0$.

$$x_n = a(-r)^n + (r/2) r^n$$

To solve for a , we simply use our definition $x_0 = 1$.

$$x_0 = a(-r)^0 + (1/2r) r^0$$

$$1 = a + 1/2 r$$

$$a = (2r - 1)/2r$$

$$x_n = ((2r - 1)/2r)(-1)^n r^n + (1/2r)r^n$$

$$x_n = r^{n-1}[(2r - 1)(-1)^n + 1]/2$$

Chapter 17

Mixing in Bluffs: Finite Pot [0,1] Games

For the games in Chapter 16, we made a special rule that does not normally apply to poker situations - that no folding was allowed. In practice, we can sometimes encounter situations analogous to this when the pot is very large. One of the important features of these games is that there is no bluffing - the betting regions for both players are defined only in terms of the strongest hands and are designed only to extract value with those strongest hands.

However, bluffing is a key element of poker, and with its introduction to the game we obtain much more complexity, and additionally better approximations and insight into real poker. In this chapter, we will consider full-street games with finite pots; the primary modification that we make to these games compared to no-fold games is the introduction of bluffing regions that correspond to the value betting and raising regions.

Example 17.1 - [0,1] Game #9

One full street.

One bet remaining.

The pot size is P and the bet size 1 unit.

Recall that when we solved the first finite pot [0,1] game in Chapter 11 (see Example 11.2), in addition to the y_n notation we used throughout Chapter 16, we used two other notations, y_0 for the threshold between checking and bluffing, and x_1^* for the threshold between X check-calling and X check-folding to a bet. In the same way that Y bluffed his very worst hands and value bet his best hands, X will likewise bluff his worst hands and value bet his best.

In the full-street game, X can additionally take the action of betting (producing a threshold x_1 and of bluffing (threshold x_0). Y has to add to his options a threshold between calling or folding if X bets, y_0^* . Here, because the options are between calling and folding, it is dominated for Y to call with hands weaker than hands he would fold. Instead, he could just call with the stronger hands and fold the weaker.

We parameterize the solution this way:

X's strategy consists of the following thresholds:

- x_0 between check-folding and bluffing.
- x_1^* between check-folding and check-calling
- x_1 between value betting and check-calling.

Y's strategy likewise consists of three thresholds:

- y_0 between checking and bluffing if X checks.
- y_1^* between calling X's bet and folding to X's bet (if X bets)
- y_1 between value-betting and checking if X checks.

As far as ordering these thresholds, we can conjecture from our previous games that Y will bet more hands behind X's check than X would have bet - hence $y_1 > x_1$. We will also find that α , the bluffing to betting ratio (see Equation 11.1), for this game is the same as α in [0.1] Game #2 (see Example 11.2) for both players. Since the size of the bluffing region will be αx_1 for X and αy_1 for Y, we can place y_0 to the left of x_0 .

The calling thresholds for both players will be somewhere between the opponent's value betting region and his bluffing region (to make him indifferent to bluffing). Hence we know that $x_1 < y_1^* < x_0$ and $y_1 < x_1^* < y_0$. The relationship between y_1^* and x_1^* actually turns out to be irrelevant - when we write the indifference equations for the various

thresholds, we only consider one or the other of these thresholds.

So our overall parameterization looks like this:

$$0 < x_1 < y_1 < x_1^* , y_1^* < y_0 < x_1 < 1$$

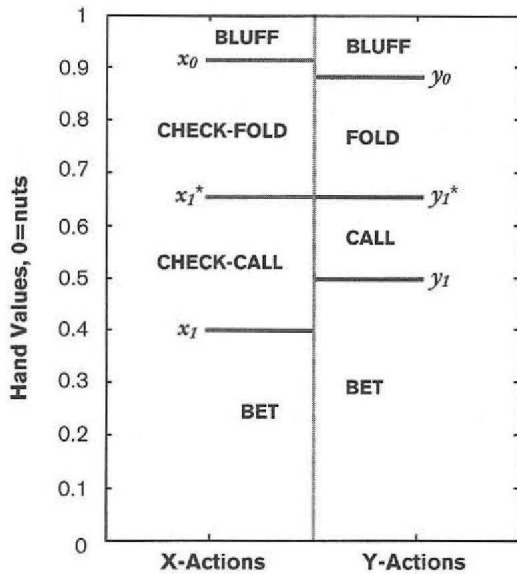


Figure 17.1, Strategy structure for [0,1] Game #9

We can now write indifference equations for each of these thresholds. We have purposely left out regions where the given action cannot occur: i.e., at y_1 , Y's decision is between checking and betting. X's hands in regions $[0, x_1]$ and $[x_0, 1]$ do not allow for Y to make that decision because X has already bet; hence they are omitted.

For Y at y_1 (indifference between checking and betting):

X's hand	Y bets	Y checks	Difference	Product
$[x_1, y_1]$	-1	0	-1	$-(y_1-x_1)$
$[y_1, x_1^*]$	+1	0	+1	$x_1^*-y_1$
$[x_1^*, x_0]$	0	0	0	0
Total				$-2y_1 + x_1 + x_1^*$

$$-2y_1 + x_1 + x_1^* = 0$$

$$y_1 = (x_1 + x_1^*)/2$$

The intuitive meaning of this equation is that Y will clearly bet all the hands that X would have bet for value. Additionally, he will bet with half the hands with which X will check call. So y_1 will be the midpoint of $[x_1, x_1^*]$.

For Y at y_1^* (indifference between checking and betting):

X's hand	Y calls	Y folds	Difference	Product
[0, x_1]	-1	0	-1	$-x_1$
[x_0 , 1]	$+(P+1)$	0	$+(P+1)$	$(P+1)(1-x_0)$
Total				$(P+1)(1-x_0) - x_1$

$$(P+1)(1-x_0) - x_1 = 0$$

$$1-x_0 = \alpha x_1$$

This equation expresses the relationship between the size of the betting region and the size of the bluffing region. $1-x_0$ is the size of the bluffing region, and x_1 is the size of the value betting region. The multiplier α (recall that $\alpha = 1/(P + 1)$) is the ratio between the two regions.

For X at x_1^* (indifference between check-calling and check-folding):

Y's hand	X check-calls	X check-folds	Difference	Product
[0, y_1]	-1	0	-1	$-y_1$
[y_0 , 1]	$P + 1$	0	$+(P + 1)$	$(P + 1)(1 - y_0)$
Total				$(P + 1)(1 - y_0) - y_1$

$$(P + 1)(1 - y_0) - y_1 = 0$$

$$1 - y_0 = \alpha y_1$$

For Y at y_0 (indifference between check-folding and bluffing):

X's hand	Y bets	Y checks	Difference	Product
[x_1 , x_1^*]	-1	0	-1	$-(x_1^* - x_1)$
[x_1^* , y_0]	$+ P$	0	$+ P$	$(P)(y_0 - x_1^*)$
[y_0 , x_0]	0	0	0	0
Total				$-(x_1^* - x_1) + (P)(y_0 - x_1^*)$

$$-(x_1^* - x_1) + (P)(y_0 - x_1^*) = 0$$

$$x_1^* - x_1 = P(y_0 - x_1^*)$$

$$x_1^* = \alpha(py_0 + x_1)$$

$$x_1^* = (1 - \alpha)(y_0) + \alpha x_1 + (x_1 - x_1)$$

$$x_1^* = (1 - \alpha)(y_0) - (1 - \alpha)x_1 + x_1$$

$$x_1^* = (1 - \alpha)(y_0 - x_1) + x_1$$

We saw in [0,1] Game #2 (Example 11.3) that just as the bluffing to betting ratio is α , the player who has been bet into calls with $1 - \alpha$ of his hands that can beat a bluff. In this case, however, when X is deciding whether to call, he has already bet with some hands. The hands he has checked are from x_1 to x_0 - of those, only the hands better than y_0 can beat a bluff.

Then x_1^* is equal to $(1 - \alpha)$ times the size of X's region of hands that can beat a bluff plus x_1 .

For X at x_0 . (indifference between check-folding and bluffing):

Y's hand	X check-calls	X bets	Difference	Product
$[0, y_1^*]$	0	-1	+1	y_1^*
$[y_1^*, x_0]$	0	+P	-P	$(-P)(x_0 - y_1^*)$
$[y_0, 1]$	-P	0	-P	$-P(1 - x_0)$
Total				$(P + I) y_1^* - P$

$$(P + I) y_1^* - P = 0$$

$$y_1^* = P/(P+I)$$

$$y_1^* = 1 - \alpha$$

Here we have the same situation, except that when X bets, Y has not yet limited his distribution (as X did by checking initially at the last threshold). Hence, Y calls with $1 - \alpha$ of his hands. We omit the qualifier "that can beat a bluff" here, because for X the decision is between check-folding and bluffing; if he check-folds, he would lose to Y's weakest hands anyway, because Y would bluff.

For X at x_1 (indifference between check-calling and betting):

Y's hand	X bets	X check-calls	Difference	Product
$[0, x_1]$	-1	-1	0	x_1
$[x_1, y_1]$	+1	+1	0	0
$[y_1, y_1^*]$	+1	0	+1	$y_1^* - y_1$
$[y_1^*, y_0]$	0	0	0	0
$[y_0, 1]$	0	+1	-1	$(1 - y_0)$
Total				$y_1^* - y_1 - (1 - y_0)$

$$y_1^* - y_1 - (1 - y_0) = 0$$

$$y_1^* - y_1 = 1 - y_0$$

Here we can see that X has two reasonable options with a hand on the cusp of initially betting or checking. When he checks, he induces bluffs and gains a bet; when he bets, he picks up a bet from Y's hands that are too weak to bet but will call a bet. Therefore, the size of these two regions is equal.

We now have six equations, each of which represents a particular relationship that must be satisfied by the optimal strategy. We also have six unknown values (the thresholds). We can solve this system of equations using algebra, to find the values of the thresholds.

$$y_1 = (x_1^* + x_1)/2$$

$$1 - x_0 = \alpha x_1$$

$$1 - y_0 = \alpha y_1$$

$$x_1^* = (1 - \alpha)(y_0 - x_1) + x_1$$

$$y_1^* = 1 - \alpha$$

$$y_1^* - y_1 = 1 - y_0$$

Solution:

$$y_1 = (1 - \alpha)/(1 + \alpha)$$

$$x_1 = (1 - \alpha)^2/(1 + \alpha)$$

$$y_0 = (1 + \alpha)^2/(1 + \alpha)$$

$$x_0 = 1 - \alpha (1 - \alpha)^2/(1 + \alpha)$$

$$y_1^* = 1 - \alpha$$

$$x_1^* = 1 - \alpha$$

Lets look at these thresholds in a specific example. Suppose the pot is 4 bets. Then, using Equation 11.1, $\alpha = 1/5$. We know that X will value bet some fraction of his hands. Y, of course, calls with hands equal to the range with which X bluffs, breaking even on that betting, but he must also call with additional hands to make X indifferent to bluffing. When X bluffs and Y calls, he loses a single bet; when Y folds, he gains 4 bets. In order to make X indifferent to bluffing, Y must call four times as often as he folds. Since Y has the entire interval [0,1] to choose from, he calls with 4/5 of his hands overall, or $1 - \alpha$.

However, X, too, must balance his value bets and bluffs so that Y cannot exploit him by calling more or less often. For example, if X never bluffed, then Y could simply call with fewer hands than X value bet, gaining value on each one. X wants to extract value with his value bets, so he must bluff in order to force Y to call with some hands between X's value betting range and X's bluffing range.

X must bluff enough so that Y is indifferent to calling with hands between x_1 and x_0 . Assuming the pot is again 4, X gains 4 bets when Y folds but loses one bet when Y calls. To make Y indifferent to calling, X must value bet four times as often as he bluffs', he must bluff, then, $1/5$ as often as he value bets, or αx_1 . It should also be clear by this logic that whatever hands Y value bets after X checks must be equally supported by a ratio of bluffs in the region $[y_0, 1]$.

We then consider how often X should value bet. At the threshold point between betting and check-calling, X has what are essentially two options; he can bet, in which case he gains a bet when Y's hand is between y_2 and y_1^* (when Y calls his bet with a worse hand but would not bet if checked to), or he can check, in which case he gains a bet when he induces a bluff when Y is between y_0 and 1. Since we know that X will be indifferent between these two actions, the size of these two regions $[y_1, y_1^*]$ and $[y_0, 1]$ must be equal.

Since we know that the second region's size is αy_1 and the first region's size is $1 - \alpha - y_1$ we can easily solve for y_1 the number of hands that Y will bet if checked to. As we saw in the solution above, this value is $(1 - \alpha)/(1 + \alpha)$. Returning to the example with pot size 4, we know that y_1^* is $4/5$. We also know that $1 - y_0$ is $1/5$ of y_1 . So $4/5 - y_1 = 1/5 y_1$, or $y_1 = 2/3$.

Now we have intuitively shown the derivations for y_1 , y_1^* , y_0 , and the relationship between x_1 and $1 - x_0$. Solving for x_1^* and x_1 is best done simultaneously. We find two relationships between these two quantities and then reduce them to numbers.

In the no-fold games, recall that we had a principle that when X checks, Y bets half the hands that X would check (provided that X cannot raise). In the finite pot game, however, when Y is deciding to check or bet at y_1 he gains a bet when X is between y_1 and x_1^* and loses a bet when X is between x_1 and y_1 . These two regions must be of equal size, so y_1 must be the midpoint between x_1 and x_1^* .

We can restate this as:

If your opponent could have bet but checks, you cannot be raised, and your opponent may call or fold, the proper strategy is to value bet all hands that your opponent would have value bet and half the hands with which he will call.

We can see that this dovetails with the earlier principle - in the no-fold case, the opponent must always call. Hence, value betting half the hands is correct. However, if your opponent may fold, the idea is to bet half the hands he will call.

Additionally, X must choose x_1^* such that he defends against Y's bluffs in the same way that Y defended against X's bluffs. However, X does not have the entire $[0,1]$ distribution for defense; instead, he has only the subset $[x_1, x_0]$, as he would have bet hands outside of this range, either as a bluff or as a value bet. Still, though, he must defend against Y's bluffs. So he must call with $(1 - \alpha)$ of the region $[x_1, x_0]$. This implies that x_1^* must be $1 - \alpha$ of the way from x_1 to x_0 .

With pot size 4, $y_2 = 2/3$ is the midpoint of x_1 and x_1^* while x_1^* is $4/5$ of the way from x_1 to $1 - (1/5)x_1$. If we call c the distance from x_1 to x_0 , a the distance from x_1 to y_1 and b the distance from 0 to x_1 , we get the following relationships:

$$\begin{aligned} b &= 2/3 - \alpha \\ 1 &= 4/5 b + 5/2 \alpha \\ b &= 4/5 (1 - 5/2 \alpha) = 4/5 - 2\alpha \\ 4/5 - 2\alpha &= 2/3 - \alpha \\ \alpha &= 2/15 \\ b &= 8/15 \\ c &= 1/3 \end{aligned}$$

.. .which matches the results we get from plugging into the formulas ($b = x_1$).

Understanding this game is essential for moving on to more complex finite pot games. We spent a fair amount of time explaining intuitively the meanings of the equations. As we move forward, many things that were explained here will be taken for granted, so we encourage you to take some time to try to work back through how the intuitive explanations mirror the algebra above. The next game is the counterpart to $[0,1]$ Game #6 (Example 16.3) - the introduction to raising.

Example 17.2 - $[0,1]$ Game #10

One full street.

Two bets left.

Check-raising is not allowed.

The pot size is P and the bet is 1 unit.

You may have already guessed that in this game we will again utilize the y_2 notation. Additionally, we will introduce the last piece of notation necessary to generalize the $[0,1]$ game - the raise-bluff notation.

Consider the situation of Player Y after X has bet. What should his strategy look like? It should be intuitively clear by now that in α certain region (his best hands), he will raise for value; with a different, weaker region of hands, he will call to force his opponent to bluff correctly, and with yet another region he will fold because his hands lack the necessary value to continue.

Suppose we set up a strategy with these three regions. Now consider the situation after value-raises. Given a strategy of the type above, how should X respond? X knows that Y has a strong hand for a raise, so he can fold hands beneath Y's raising threshold and equalize. This is analogous to what X would do if there were only one bet left and Y never bluffed; X would simply throw away his mediocre hands. So Y cannot just value raise with his best hands without protecting those value raises with some type of raise-bluffs.

The next question is: From where should those bluffs come?

To answer this question, we must consider the nature of a bluff. In cases where there is only one bet remaining, we bluff with our very worst hands. We choose the very worst hands because they are the hands that have the least showdown value if the pot is checked through. These hands stand to gain the most from being played as a bluff. When the opponent calls our bluffs, we should never win the pot if the opponent is playing optimally (or even just rationally).

But now consider the situation when the opponent has already bet. We will still want to raise with our best hands and call with our next group of hands because we want to use the showdown value of our hands most efficiently. Using the principle that raise-bluffs should never win the pot if called, we can see that the raise-bluffs will come from the remaining fraction of hands that are neither raising hands nor calling hands. Of that fraction, some will be raise-bluffs, in a proportion that protects our value raises properly with respect to the size of the pot and the size of our value-raising region, and the remainder will be folded.

Perhaps the reader's intuition now leads to a situation similar to bluffing in one-bet scenarios, where the bluffs are the worst hands in the region. But this is incorrect. Technically speaking, for Y the raise-bluffs can come from anywhere in the group of hands worse than y_1^* and be at least co-optimal. However, since we are going to fold the non-raise-bluff hands and there is no hope of a showdown with hands better than a bluff, using the strongest hands in the (fold, raise-bluff) region dominates any other strategy.

This gives rise to a new type of threshold and also forces us to redefine a previous threshold for games with more than one bet.

- y_n^* is redefined as the threshold between calling the n^{th} bet and, if the $(n + 1)^{\text{th}}$ bet and folding are allowed, putting that bet in as a bluff, or otherwise folding.
- $y_n^\#$ is the threshold between the region which puts the n^{th} bet in as a bluff and the region which folds to the $(n - 1)^{\text{th}}$ bet.

As a result of these definitions, y_n^* will be paired and closely related to $y_{n+1}^\#$.

When we perform parameterizations, different game trees are sometimes possible based on previous actions. As a result, sometimes we may draw thresholds above or below other ones when the other is not necessary, because the two are never material at the same time.

For example, x_1^* only matters if X checks. If X does check, then y_1^* and $y_2^\#$ are irrelevant. So whether x_1^* lies above or below y_1^* is arbitrary for the parameterization (although the solution will still tell the location).

With the above caveats, we parameterize [0,1] Game #6 this way:

$0, y_2, x_2^*, x_1, y_1, x_1^*, y_1^*, y_2^\#, y_0, x_0, 1$




Figure 17.2. Strategy structure for [0,1] Game #10

Translating this parameterization into strategies:

For X:

- X bets for value and calls a raise on $[0, x_2^*)$.
- X bets for value and folds to a raise on $[x_2^*, x_1]$,
- X checks and calls a bet on $[x_1, x_1^*]$.
- X checks and folds to a bet on $[x_1^*, x_0]$.
- X bluffs on $[x_0, 1]$,

For Y:

If X bets:

- Y raises for value on $[0, y_2]$.
- Y calls on $[y_2, y_1^*]$
- Y raise-bluffs on $[y_1^*, y_2\#]$
- Y folds on $[y_2\#, 1]$,

If X checks:

- Y bets for value on $[0, y_1]$.
- Y checks on $[y_1, y_0]$
- Y bluffs on $[y_0, 1]$.

Again, we advise you to carefully review this notation. Understanding what the threshold values mean is of great significance to understanding this and more complex games. In later games, we will see values such as, $x_4\#$, which is the threshold between X check-raising and putting in the fourth bet as a bluff and X folding to the third bet after check-raising, and it is easy to become lost as a result of not mastering the notation.

In the one-bet case, we introduced a variable a to represent the fundamental ratio $1/(P + 1)$ that governs in one way or another the relationships between betting and bluffing, as well as calling. In two-bet cases, however, we have situations where there is a larger pot already, and the bluff itself costs two bets. Hence, we will have a different ratio.

$$\alpha_2 = 1/(P + 3)$$

In other games, we will use this notation for larger number of bets than two, so more generally:

$$\alpha_n = 1/(P + (2_n - 1)) \tag{17.1}$$

The α from [0.1] Game #2 (Example 11.3), then, is more strictly α_1 . However, we will use the convention that an unmarked α is α_1 .

To solve this game, we now have nine unknowns. At each of these unknowns we will have an indifference equation, and solving the system of indifference equations will yield our solution.

For Y at y_2 :

X's hand	Y calls	Y raises
$[0, y_2]$	-1	-2
$[y_2, x_2^*]$	+1	+2
$[x_2^*, x_1]$	+1	+1
$[x_0, 1]$	+1	+1

Y loses an extra bet by raising X's best hands, and gains a bet from raising hands worse than y_2 that will call. Since X cannot raise, we see the familiar idea of Y raising with half the hands that X will call.

$$y_2 = x_2^* - y_2$$

$$y_2 = x_2^*/2$$

For Y at y_1 , we have exactly the same situation as in [0,1] Game #9 (Example 17.1), since X cannot raise. Hence, our indifference equation here is the same.

$$y_1 = (x_1^* + x_1)/2$$

y_1 is still the midpoint of X's betting thresholds and X's calling threshold.

For Y at y_1^* :

X's hand	Y calls	Y raise-bluffs
$[0, x_2^*]$	-1	-2
$[x_2^*, x_1]$	-1	+(P+1)
$[x_0, 1]$	+1	+1

$$x_2^* = (P + 2)(x_1 - x_2^*)$$

$$x_2^* = x_1(P + 2)/(P + 3)$$

$$x_2^* = (1 - \alpha_2)x_1$$

This is the first appearance of the idea of α_2 . Here we see that the size of X's betting region x_1 is related to the size of the region where he bets and calls a raise. In fact, he bets and folds to a raise α_2 of the time. This is because he must make his opponent indifferent to raising as a bluff. α_2 is used instead of α because the pot is now larger - X's bet and Y's raise are both now in the pot.

For Y at y_2 :

X's hand	Y raise-bluffs	Y folds
$[0, x_2^*]$	-2	0
$[x_2^*, x_1]$	$+(P + 1)$	0
$[x_0, 1]$	+1	- P

First, a bit of algebra:

$$2x_2^* = (P + 1)(x_1 - x_2) + (P + 1)(1 - x_0)$$

$$2x_2^* = (P + 1)(x_1 - x_2^* + 1 - x_0)$$

$$(P + 3)x_2^* = (P + 1)(x_1 + 1 - x_0)$$

$$\alpha x_2 = \alpha_2(x_1 + 1 - x_0)$$

$$\alpha(1 - \alpha_2)x_1 = \alpha_2(x_1 + 1 - x_0)$$

$$(\alpha/\alpha_2 - \alpha)x_1 = x_1 + 1 - x_0$$

We can also create an identity:

$$\alpha/\alpha_2 = (1/(P + 1))/(1/(P + 3))$$

$$\alpha/\alpha_2 = (P + 3)/(P + 1)$$

$$\alpha/\alpha_2 = 1 + 2\alpha$$

Substituting for α/α_2 :

$$(1 + \alpha)x_1 = x_1 + 1 - x_0$$

$$\alpha x_1 = 1 - x_0$$

Remember that previously we commented that Y's raise-bluffing region could really be anywhere in his folding region. We insist on playing non-dominated strategies, so as a result, we get an indifference between raise-bluffing and folding. However, we could also solve the game for a different threshold between calling and folding; then the above equation would simply appear instead of forcing us to go through algebraic manipulations.

This is the familiar equation (e.g., Equation 11.4, among others) that relates the size of X's betting region to the size of his bluffing region by a factor of α .

For Y at y_0 , we have the same equation as in [0,1] Game #9 (Example 17.1), because X still cannot check-raise.

$$x_1^* = (1 - \alpha)(y_0 - x_1) + x_1$$

For X at x_2^* :

Y's hand	X bet-calls	X bet-folds
$[0, y_2]$	-1	-2
$[y_1^*, y_2^\#]$	+2	-(P+1)

Notice that X's strategies are more complicated in this game. When we examine the indifference between two actions, it is not enough to state "X bets" or "X checks" - instead we must specify the entire strategy : "X bet-calls" or "X bet-folds."

x_2^* is one of the new thresholds, between X calling a raise and folding to the raise. X loses a bet by calling against Y's value raises, but picks off Y's raise bluffs. As we shall see, because it costs Y two bets instead of one to make a raise-bluff, the calling ratio for X is differently shaped.

$$y_2 = (P + 3)(y_2^\# - y_1^*)$$

$$\alpha_2 y_2 = (y_2^\# - y_1^*)$$

This equation is basically the raise-bluff equation analogous to Equation 11.4, but for two bets instead of one from earlier games. The size of the raise-bluffing region $y_2^* - y_1^\#$ is the size of the value raising region (y_2) multiplied by the appropriate ratio (α_2). This is a general rule that appears throughout all the [0,1] finite pot games. For every value region, there is a corresponding bluffing region. The size of the bluffing region is related to the size of the value region by some α value.

For X at x_1 :

Y's hand	X bet-folds	X check-calls
$[0, x_1]$	-1	-1
$[x_1, y_1]$	+1	+1
$[y_1, y_1^*]$	+1	0
$[y_1^*, y_2^\#]$	-(P + 1)	0
$[y_2^\#, y_0]$	0	0
$[y_0, 1]$	0	+1

$$y_1^* - y_1 = (P+1)(y_2^\# - y_1^*) + (1 - y_0)$$

At x_1 , we have an equation that relates the size of the region where Y calls a bet but wouldn't have bet himself ($y_1^* - y_1$), and the size of the raise-bluffing and bluffing regions for Y. This is X's choice between trying to induce a bluff with his borderline value betting hand or trying to extract value. If he bets, he gains a bet when Y calls but wouldn't have bet, but loses P+ 1 bets when Y raise-bluffs him out of the pot. If he checks, he sometimes induces a bluff from Y. The value of checking and the value of betting must be equal; this equation reflects that.

For X at x_1^* , we again have the familiar relationship:

$$1 - y_0 = \alpha y_1$$

For X at x_0 :

Y's hand	X check-folds	X bluffs
$[0, y_2\#]$	0	-1
$[y_2\#, x_0]$	0	+P
$[x_0, 1]$	-P	0

$$y_2\# = (x_0 - y_2\#)(P) + (1-x_0)(P)$$

$$(P + 1)y_2\# = P$$

$$y_2\# = 1 - \alpha$$

This result is somewhat analogous to the equation governing $y_1\#$ in $[0,1]$ Game #9 (Example 17.1); the total hands that Y folds to X's bet with is α . In $[0,1]$ Game#9 we characterized this relationship as "Y calls with $(1 - \alpha)$ of his hands," but a more accurate characterization in light of this result is "Y folds α of his hands to a bet."

We now have nine equations. Many of them are the same as the previous game.

$$y_1 = (x_1^* + x_1)/2 \tag{17.2}$$

$$y_2 = x_2^*/2 \tag{17.3}$$

$$x_2^* = (1 - \alpha_2)x_1 \tag{17.4}$$

$$\alpha x_1 = 1 - x_0 \tag{17.5}$$

$$x_1^* = (1 - \alpha)(y_0 - x_1) + x_1 \tag{17.6}$$

$$\alpha_2 y_2 = (y_2\# - y_1^*) \tag{17.7}$$

$$y_1^* - y_1 = (P + 1)(y_2\# - y_1^*) + (1 - y_0) \tag{17.8}$$

$$1 - y_0 = \alpha y_1 \tag{17.9}$$

$$y_2\# = 1 - \alpha \tag{17.10}$$

We can solve this system of equations and obtain values for all the thresholds. We will (hopefully to your pleasure) omit much of the algebra, but there is at least one point worth a bit of additional scrutiny.

Combining (17.3) and (17.4), we have the important relationship:

$$y_2 = x_1(1 - \alpha_2)/2 \tag{17.11}$$

As in $[0,1]$ Game #6 (Example 16.3), we can rewrite this to capture its identity as the ratio between the hands that X bets and the hands that Y raises:

$$y_2 = R x_1$$

Here, $R = (1 - \alpha_2)/2$. We can see that as P goes to infinity, this value converges on $1/2$, which was the value that it took in $[0,1]$ Game #6.

The idea behind solving this system of equations is to convert everything to be in terms of y_1 and x_1 . We can solve Equations 17.2 and 17.6 for x_1^* , set them equal, and obtain one equation that contains only x_1 and y_1 . A second equation can be obtained by working on Equation 17.8. This equation contains y_1^* , y_1 , $y_2^\#$, and y_0 .

y_0 can be put in terms of y_1 with 17.9. $y_2^\#$ is a constant by 17.10. We can put y_1^* in terms of y_1 by using 17.7. y_1 can be put in terms of x_1 , as $y_2 = Rx_1$. Since $y_2^\#$ is a constant, y_1^* is now in terms of x_1 . Substituting for all these variables can make 17.8 a second equation in terms of only x_1 and y_1 .

We can reduce these two equations to the following:

$$x_1 = (1 - \alpha)^2 / [1 + \alpha + (2 - \alpha)(1 - a_2)R] \tag{17.12}$$

This equation is very important for solving games without check-raise, because it is true for not only the two-bet game but games with more bets where folding is allowed but check-raising is not. This is because the addition of more bets changes the value of R , but does not change the fundamental relationships among the lower thresholds. We will see that this relationship greatly simplifies the infinite bets case addressed below.

Once we have this, we can solve for y_1 by substitution and the remaining variables are easily found using Equations 17.2 through 17.10.

Solution:

$$y_2 = (1 - a)^2(1 + 2a) / (7a^2 + 9a + 4)$$

$$y_1 = (1 - a)(8a^2 + 9a + 3) / (1 + a)(7a^2 + 9a + 4)$$

$$y_1^* = 1 - a - a_2(1 - a)^2(1 + 2a) / (7a^2 + 9a + 4)$$

$$y_2^\# = 1 - a$$

$$y_0 = 1 - a(1 - a)(8a^2 + 9a + 3) / (1 + a)(7a^2 + 9a + 4)$$

$$x_2^* = 2(1 - a)^2(1 + 2a) / (7a^2 + 9a + 4)$$

$$x_1 = 2(1 - a)^2(1 + 2a) / (1 + a)(7a^2 + 9a + 4)$$

$$x_1^* = 4(1 - a)(2a^2 + 2a + 1) / (7a^2 + 9a + 4)$$

$$x_0 = 1 - 2a(1 - a)^2(1 + 2a) / (1 + a)(7a^2 + 9a + 4)$$

This is the solution to what is effectively the simplest complete two-bet [0,1] game. Admittedly, this final solution isn't particularly useful - what is useful is the method and the principles that were used in deriving it.

We can look at what the thresholds look like for this game for different pot sizes. The rightmost column in the following table contains a pot size of ϵ , a value frequently used by mathematicians to represent a tiny quantity. Here we use it to show that as the pot converges on zero, the thresholds converge to "no betting" as we expect.

Threshold	$P = \infty$	$P = 1$	$P = 2$	$P = 9$	$P = \epsilon (\sim 0)$
y_2	$1/4$	$2/41$	$2/21$	0.195573441	0
x_2^*	$1/2$	$4/41$	$4/21$	0.391146881	0
x_1	$1/2$	$16/123$	$5/21$	0.426705689	0
y_1	$3/4$	$76/246$	$279/630$	0.655203951	0
x_1^*	1	$20/41$	$68/105$	0.883702213	0

y_1^*	1	$20/41$	$68/105$	0.883702213	0
$y_2^\#$	1	$1/2$	$2/3$	0.9	0
y_0	1	$208/246$	$537/630$	0.934479605	0
x_0	1	$230/246$	$58/63$	0.957329431	0

We continue our discussion with the two-bet game where X can check-raise:

Example 17.3 - [0,1] Game #11

One full street.

Two bets left.

Check-raising is allowed.

The pot size is P and the bet is 1 unit.

As we progress through additional games, we will generally omit listing off the whole solution to the game and calculating each threshold value. As you saw in [0,1] Game #10 (Example 17.2), the solutions to more complex games with finite pots are substantially more complex and not that useful for seeing at a glance. Instead, we will focus on building on the results we have from previous games and showing the areas in which the strategies for more complex games differ from differ less complex counterparts.

In the check-raise game with two bets remaining, we have a number of additional thresholds. However, many of these are easily dismissed because they relate to each other in familiar ways. For example, X will have a check-raise bluff region between $x_2^\#$ and x_1^* . We know from our experience solving indifference equations that the size of this region will be the size of the value check-raising region (x_2) multiplied by the appropriate α value α_2 . Likewise, instead of just a value betting region to the left of y_1 , Y will have a "bet-call" region to the left of y_2^* and a "bet-fold" region between y_2^* and y_1 . These will be sized such that Y folds $1 - \alpha_2$ of his value bets to a raise. This makes X indifferent to raise-bluffing, and so on.

"What's interesting and important here, though, is the change in the value regions for both players. Changes in these value regions will cascade through the other thresholds and impact both players' strategies, but the relationships of the other thresholds to the value regions are fairly simple and straightforward. We could apply the method of writing indifference equations for each threshold, but instead let us simply draw attention to the important thresholds.

In the check-raise game, we know that Y's value-betting strategy after X checks will be impacted by X's ability to check-raise; additionally Y's value-raising strategy will be impacted by the fact that X is checking some of his strong hands and likely betting a little weaker than he does in the game with no check-raise. Recall that these were the major shifts in the game with no folding allowed - Y bet a little tighter after X checked (fearing a check-raise), X bet a little looser (to capture some of the value lost by not inducing weaker hands to bet), and Y raised a little looser (because of X's looser betting standards).

There is an important link between the check-raise game and the no-check raise game. To see this, consider the indifference at x_2 .

For X at x_2 :

Y's hand	X check-raises	X bets-calls
$[0, x_2]$	-2	-2
$[x_2, y_2]$	+2	+2
$[y_2, y_2^*]$	+2	+1
$[y_2^*, y_1]$	+1	+1
$[y_1, y_1^*]$	0	+1
$[y_1^*, y_2^\#]$	0	+2
$[y_0, 1]$	+1	0

$$y_2^* - y_2 + (1 - y_0) = 2(y_2^\# - y_1^*) + y_1^* - y_1$$

To find y_2^* , we consider the indifference at $x_2^\#$:

Y's hand	X check-calls	X check-raise-bluffs
$[0, y_2^*]$	-1	-2
$[y_2^*, y_1]$	-1	-(P + 1)
$[y_0, 1]$	+ 1	+i

$$y_2^* = (P + 2)(y_1 - y_2^*)$$

$$y_2^* = (1 - \alpha_2)y_1$$

We additionally know that $y_2^\# - y_1^*$ is equal to $\alpha_2 y_2$.

So we have:

$$(1 - \alpha_2)y_1 - y_2 + (1 - y_0) = 2\alpha_2 y_2 + y_1^* - y_1$$

In the no-check-raise game, we had the following from the indifference at x_1

$$y_1^* - y_1 = (P + 1) + (y_2^\# - y_1^*) + (1 - y_0)$$

The origin of this equation is that X must be indifferent to bet-folding and check-calling at x_1 . When he bets, he gains a bet against Y when Y has a hand that will call a bet but not bet if checked to. Against this, he loses (P + 1) bets when Y raise-bluffs, and additionally loses the bluffs that he fails to induce by betting. The bluff terms drop off, and we have:

$$(1 - \alpha_2)y_1 - y_2 = (P + 3)\alpha_2 y_2$$

$$(1 - \alpha_2)y_1 = 2y_2$$

$$y_2 = y_1(1 - \alpha_2)/2$$

Recall that in the no-check-raise game, we saw that this ratio $(1 - \alpha_2)/2$ was the ratio between y_2 and x_1 . Here we see that it is the ratio between y_2 and y_1 . Likewise, when we solved the infinite bets game without check-raise (Example 16.3), we saw that r was the ratio between successive thresholds (y_2, x_3, y_4 , etc.), while in the game where check-raise was allowed (Example 16.4), r is the ratio between successive y_n thresholds.

As a general result, we find that when converting [0,1] games from no check-raise to including check-raise, this relationship holds (for simple symmetrical cases). The ratio of Y thresholds to X thresholds in the no-check-raise game is the ratio of successive Y thresholds in the game with check-raise.

We can follow this chain of logic to find that the general equation for y_1 in this game is:

$$y_1 = (1 - \alpha) / [1 + \alpha + (1 - \alpha_2)R] \quad (17.13)$$

This equation is directly analogous to Equation 17.12, but includes the check-raise element of this game.

Calculating the values of the various thresholds is a straightforward process of creating indifference equations and solving them. Bluffing regions have size related by an a value to their corresponding value regions, and so on.

hi the next game, we tackle the more complex problem of allowing infinite raises.

Example 17.4 - [0,1] Game #12

One full street.

Check-raise is not allowed.

An arbitrary number of bets and raises are allowed.

Folding is allowed.

In solving this game, we will relate its components to games we have previously solved. Most of the indifference equations in this game are identical to their counterparts in [0,1] Game #10 (Example 17.2); in fact, the only indifference equation that is invalid is the indifference at y_2 . In that game, the second bet was the last one; hence, Y could not be raised back. We found there that Y raised with half the hands with which X would call. However, here X can additionally raise back, so we must reconsider that equation.

Additionally, we add an infinite number of thresholds to the game, allowing for the arbitrary allowed bets and raises. For each threshold, we have an additional indifference equation. Many of these equations have a clear form, as for example we will obtain for any appropriate n :

$$a_n y_n = (y_n - y_{n-1})^*$$

This is the familiar bluffing-to-betting ratio (based on Equation 11.4), governed by the appropriate α . value and the number of bets to be put in the pot.

When we solved [0,1] Game #10, we identified a key equation that is the general form of x_1 for non-check-raise games.

$$(1 - \alpha)^2 = x_1 [1 + \alpha + R(2 - \alpha)(1 - \alpha_2)]$$

The equation that must be modified to allow for the difference between this game and [0,1] Game #10 is the following:

$$y_2 = R x_1$$

However, this equation need not change except that R has a new value.

Hence, Equation 17.12 still holds. All that's necessary is that we find the appropriate R . We can check this by looking at games we have already solved.

Consider the game where folding was not allowed (Example 16.4). Allowing the pot to go to infinity, we can see that all the α values go to zero. In that game, y_2 was $1/2$ and x_1 was $1/2$. Hence, R was $1/2$. Plugging this into our formula and letting α and α_2 go to zero, we get $x_1 = 1/2$.

In the one-bet game, R was zero because y_2 was at 0. So x_1 equals $(1 - \alpha)^2 / (1 + \alpha)$. Looking back to the solution for [0,1] Game #9, we see that this is the case.

Returning to the game at hand, we find the following indifferences (which should be intuitive to those who have come this far):

$$x_3\# - x_2^* = \alpha_3 x_3$$

The size of the n th bluffing region is α_n times the size of the corresponding value region.

$$x_3\# = (1 - \alpha_2)x_1,$$

When you put in the n^{th} bet for value, you must at least call (or raise) with $1 - \alpha_{n+1}$ of the hands that originally put in that bet (to avoid being exploited by bluff raises).

We also have the following indifference at y_2 . (This is the lowest threshold at which this game differs from [0,1] Game #6.)

$$x_2^* - y_2 = y_2 + (P + 3)\alpha_3 x_3$$

Y picks up a bet at y_2 by value raising when X is between y_2 and x_2^* . He loses one additional bet whenever X is better than y_2 (folding to additional raises) and also loses $(P + 3)$ bets when X is in his reraise-bluff region $x_3\# - x_2^*$.

Solving this equation, we find:

$$(1 - \alpha_2)x_1 - \alpha_3 x_3 = 2y_2 + (P + 3)\alpha_3 x_3$$

Converting α terms to include P :

$$x_1(P + 2)/(P + 3) = 2y_2 + x_3(P + 4)/(P + 5) \tag{17.14}$$

The indifferences here are identical as we increment the thresholds by two bets; hence this equation is a particular case of the general recursion for x_n, y_{n+1}, x_{n+2} as we increment to higher thresholds.

We can take advantage of a principle that once we reach the second raise, the game becomes completely symmetrical in terms of the ratios between thresholds except for the size of the pot. The values of the thresholds are different, of course, because they have been multiplied by the value analogous to r , the same value as in Equation 14.1. We have been calling this value R ; because to this point, we have only considered games where one raise was allowed. However, as the pot grows, the value of this multiplier changes as well.

We will then call this changing multiplier R_p because it changes based on the size of the pot, and we have the following:

$$\begin{aligned}y_2 &= R_p x_1 \\x_3 &= R_p + 2y_2 \\y_4 &= R_p + 4x_3\end{aligned}$$

...and so on.

Substituting into 17.14:

$$x_1(P+2)/(P+3) = 2R_p x_1 + (R_p R_p + 2x_1)(P+4)/(P+5)$$

The x_1 terms drop out, and we have:

$$R_p = \frac{\frac{(P+2)}{(P+3)}}{2 + \frac{(R_{p+2})(P+4)}{(P+5)}}$$

This is the general recursion that describes the relationship between successive values of R_p . To find specific values, we can use the fact that we know that R_p converges on r as P goes to infinity (since this is an approximation of the no-fold game as shown in Example 16.4). Hence, we can choose an arbitrarily large value of n , set it equal to $\sqrt{2} - 1$, and work backwards to the desired P . As it happens, this recursive relation converges quite rapidly. Of course, we can see that for the no-fold case, R_p simply is r for all thresholds, and this game simplifies to its no-fold analogue.

Once we have the value of R_p , we can easily calculate the values of the thresholds by solving for x_1 and then multiplying by the appropriate R_p values. As usual, the bluffing regions and calling regions follow the value regions, and we have a solution.

Key Concepts

- The $[0,1]$ games can be extended to cover more complicated situations, including check-raising, an arbitrary number of bets, and so on.
- In all of these situations, however, the keys to the game are the value thresholds. For each value region, there are corresponding bluffing, calling, folding, and raise-bluffing regions that can be calculated by relating the size of the pot to the size of the value region.
- In games without check-raise, there is a recurrence relation between alternating X and Y thresholds such that $x_1 > y_2 > x_3$ and so on.
- In games with check-raise, there is a similar recurrence relation between consecutive Y thresholds.
- If your opponent could have bet but checks, you cannot be raised, and your opponent may call or fold, the proper strategy is to value bet all hands that your opponent would have value bet and half the hands with which he will call.

Chapter 18

Lessons and Values: The [0,1] Game Redux

At the outset of this book, we described an approach to studying poker that focused on the dimensions of the game; instead of attempting to solve all of poker all at once, an impossible task, we instead focus on the different elements and characteristics of games and attempt to understand how they affect our play, whether optimal or exploitive. The [0,1] game is a clear example of this; over the course of Part III, we have investigated twelve different [0,1] games, each with different betting structures and rules. To this point, we have focused on the solutions to games, and also to the methodologies we could use to solve any game of this type.

However, another important aspect of analyzing games and learning lessons is to consider the value of the game. The *value* of the game is the expectation of a particular player assuming the game is played with *optimal strategies*. By convention, we state the value of the game in terms of Y's expectation; we do this because almost all poker games favor the second player because of his positional advantage. Thus, when we use this convention, values of games are usually positive. We also, as usual, discuss the *ex-showdown* value of the game unless otherwise stated. We often use the letter *G* to denote the game value, whether in a particular region or overall.

We can always calculate the value of a game by brute force; simply calculating all the possibilities for hands that either player would hold, finding the strategies they would employ, the joint probabilities of all possible holdings, and adding up the resultant value distribution. This is often an easy and useful method for solving simple games, such as [0,1] Game #1 (Example 11.2).

X's hand	Y's hand	Probability	Outcome	Weighted Factor
$[0, \frac{1}{2}]$	$[0, \frac{1}{2}]$	$\frac{1}{4}$	0	0
$[0, \frac{1}{2}]$	$[\frac{1}{2}, 1]$	$\frac{1}{4}$	0	0
$[\frac{1}{2}, 1]$	$[0, \frac{1}{2}]$	$\frac{1}{4}$	+1	$+\frac{1}{4}$
$[\frac{1}{2}, 1]$	$[\frac{1}{2}, 1]$	$\frac{1}{4}$	0	0
Total				$+\frac{1}{4}$




Figure 18.1 -Game Value for [0,1] Game #1

So Y's equity in the half-street no-fold game is $\frac{1}{4}$. We can apply this methodology to many of the other games we solved as well. However, this method becomes more and more difficult to implement, as the game gets more and more complex. Fortunately, we can take advantage of some of the properties of optimal strategies to find methods that are more suitable for finding game values for games that are more complex.

Consider this same game. Suppose that we hold Y's strategy constant - that is, Y plays his optimal strategy. We can move X's strategy around, knowing that X can never gain value by doing this; remember, if X can improve his equity, then Y isn't playing optimally.

So we can have X follow a "best response" strategy instead of his optimal strategy. This is convenient for us, because on hands where X is indifferent to different actions, we can choose whichever action is easiest for calculation purposes. *Note that this does not mean that X can play these strategies and achieve these values; Y could exploit him.* We are trying to calculate the value of the game here, not the strategies both players will follow.

Considering full-street games:

[0,1] Game #4 (Example 15.1) was more of an instructive examination of degeneracy than a game in the mold of the other games; its equity is quickly calculated by noting that both players put in one bet with all hands. Hence the ex-showdown value of the game is zero.

[0,1] Game #5 (Example 15.2) was the no-fold game with a single raise allowed, but no check-raising. The solution to this game was:

$$\begin{aligned} x_1 &= \frac{1}{4} \\ y_1 &= \frac{3}{4} \\ y_2 &= \frac{1}{4} \end{aligned}$$

We can easily calculate the equity of this game by holding Y's strategy constant, and finding a best-response strategy for X. The key to the best response strategy is that X can either check or bet between $\frac{1}{4}$ and $\frac{3}{4}$ against the optimal strategy and obtain the same value. If he bets, he gains a bet from Y's hands between $\frac{3}{4}$ and 1, but loses a bet to Y's hands between 0 and $\frac{1}{4}$ (compared to checking). Since the equities from betting or diecking are the same, we can simply use $+\frac{1}{4}$ as Y's equity when X is on $[\frac{1}{4}, \frac{3}{4}]$.

X's hand	Y's hand	Probability	Outcome	Weighted Factor
$[0, \frac{1}{4}]$	$[0, \frac{1}{4}]$	$\frac{1}{16}$	0	0
$[0, \frac{1}{4}]$	$[\frac{1}{4}, 1]$	$\frac{3}{16}$	-1	$-\frac{3}{16}$
$[\frac{1}{4}, \frac{3}{4}]$	$[0, 1]$	$\frac{1}{2}$	$+\frac{1}{4}$	$+\frac{1}{8}$
$[\frac{3}{4}, 1]$	$[0, \frac{3}{4}]$	$\frac{3}{16}$	+1	$+\frac{3}{16}$
$[\frac{3}{4}, 1]$	$[\frac{3}{4}, 1]$	$\frac{1}{16}$	0	0
Total				$+\frac{1}{8}$




Figure 18.2. Game Value for [0,1] Game #5

So the value of the two-bet no-check-raise game to Y is $\frac{1}{8}$ of a unit. This is actually quite a bit of value! To put this into terms that you might be more familiar with, let's consider the total action that Y puts into the pot.

When X bets, Y raises $\frac{1}{2}$ of the time, and calls a bet the rest of the time, for a total of $(\frac{5}{4})(\frac{1}{2})$ action, or $\frac{5}{8}$ of a bet. When X checks, Y puts in one bet $\frac{3}{4}$ of the time, for a total of $(\frac{3}{4})(\frac{1}{2}) = \frac{3}{8}$ of a bet. So the total average action Y puts in from one trial of this game is one bet. For this, he obtains an eighth of a unit, for an edge of 12.5% on his total action.

The reason for this edge is entirely positional; the game is symmetric except that Y acts second. You can see from this simple example that the value of position is inherent in poker. However, this game is not as biased against X as the half-street game was. Recall that in that game, X is reduced to checking and calling. Here, he gains an additional strategic option, that of betting. Y, at the same time, gains the option of raising. It turns out that this game, where X can bet, is worth half as much to Y as the game where X is compelled to check in the dark.

We can further remove the bias against X (in a manner more like regular poker) by looking at the same game, but with check-raise allowed. [0,1] Game #7 (Example 16.4) is the same as [0,1] Game #5 (Example 16.2), except for the addition of check-raise. We should then expect that X will do better because he has gained an additional strategic option.

Let's calculate the equity of this game:

The solution to this game was:

$$\begin{aligned} x_1 &= \frac{5}{9} \\ y_1 &= \frac{2}{3} \\ x_2 &= \frac{1}{9} \\ y_2 &= \frac{1}{3} \end{aligned}$$

Again, we hold Y's thresholds constant, and examine X's equity from his best response strategies.

When X is below $\frac{1}{3}$, we know that no matter what X does, two bets will go in when Y is between 0 and $\frac{1}{3}$. X gains a bet from check-raising (compared to betting) when Y is between $\frac{1}{3}$

and $\frac{2}{3}$, but loses that bet when Y is between $\frac{2}{3}$ and 1. So X is indifferent to betting or check-raising on the entire interval $[0, \frac{1}{3}]$.

When X is between $\frac{1}{3}$ and $\frac{5}{9}$, he no longer can profitably check-raise, because the additional bet he loses to Y's good hands is not made up by the additional bet he gains from Y's mediocre betting hands. Hence, he will simply bet in this interval. Above $\frac{5}{9}$ X checks.

So we can evaluate the table assuming that X plays the following strategy:

Bet hands below $\frac{5}{9}$

Check and call with hands above $\frac{5}{9}$.

This strategy has the same equity as X's optimal strategy against Y's optimal strategy. Y could exploit it, but remember we are only trying to find the game value here, not strategies for the two players to follow.

X's hand	Y's hand	Probability	Outcome	Weighted Factor
$[0, \frac{1}{3}]$	$[0, \frac{1}{3}]$	$\frac{1}{9}$	0	0
$[0, \frac{1}{3}]$	$[\frac{1}{3}, 1]$	$\frac{2}{9}$	-1	$-\frac{2}{9}$
$[\frac{1}{3}, \frac{5}{9}]$	$[0, \frac{1}{3}]$	$\frac{2}{27}$	+2	$+\frac{4}{27}$
$[\frac{1}{3}, \frac{5}{9}]$	$[\frac{1}{3}, \frac{5}{9}]$	$\frac{4}{81}$	0	0
$[\frac{1}{3}, \frac{5}{9}]$	$[\frac{5}{9}, 1]$	$\frac{8}{81}$	-1	$-\frac{8}{81}$
$[\frac{5}{9}, 1]$	$[0, \frac{5}{9}]$	$\frac{20}{81}$	+1	$+\frac{20}{81}$
$[\frac{5}{9}, 1]$	$[\frac{5}{9}, \frac{2}{3}]$	$\frac{4}{81}$	$+\frac{3}{4}$	$+\frac{1}{27}$
$[\frac{5}{9}, 1]$	$[\frac{2}{3}, 1]$	$\frac{4}{27}$	0	0
Total				$+\frac{1}{9}$




Figure 18.3 Game Value for $[0,1]$ Game #7

The equity of the game with check-raise is $\frac{1}{9}$ compared to the equity of the game without check-raise, which is $\frac{1}{8}$. So X gains $\frac{1}{72}$ of a unit in value from the strategic option of check-raising. This isn't a lot. In fact, when we first solved this game, we expected that the value of check-

raising would be significant; we conjectured that it would reduce the game value by half. But looking at this more critically, we find that check-raise doesn't really impact Y's strategy tremendously. He raises a little more loosely (from $\frac{1}{4}$ to $\frac{1}{3}$), and bets a little tighter (from $\frac{3}{4}$ to $\frac{2}{3}$), but neither of these changes create a very large movement in equity from the non-check-raise case.

The final two no-fold games that we considered were the games with infinite raises allowed. [0,1] Game #6 (Example 16.3), where check-raise was not allowed, is the simpler of these games. To find the value of this game, we can take advantage of the game's symmetry once we reach the second bet.

The solution to this game was:

$$x_1 = 1/(1 + 2r)$$

$$x_n = r_{n-1}x_1 \text{ (for odd } n > 1)$$

$$y_1 = (1 + r)x_1$$

$$y_n = r_n - lx_1 \text{ (for even } n > 1)$$

Let's consider first the various regions of interest, working backwards from the weakest hands. We can construct a matrix as follows (X's hands across the top, Y's on the left):

	$[0, x_1]$	$[x_1, y_1]$	$[y_1, 1]$
$[0, x_1]$	* (see below)	+1	+1
$[x_1, y_1]$	-1	0	+1
$[y_1, 1]$	-1	0	0




Figure 18.4. Game Value for [0,1] Game #6

We have marked the area where both players have a hand better than x_1 with a* because the equity in that area is unclear; there are many thresholds where players win different amounts of bets. Specifying the entire range of infinite thresholds would be impossible. However, consider the situation where both players have a hand on $[0, x_1]$. X would bet, and when both players are on this interval, we have a sub-game.

Let's use g to represent the equity for Y in the region marked *

Within the * region, we have:

X's hand	Y's hand	Probability	Outcome	Weighted Factor
$[y_2, x_1]$	$[y_2, x_1]$	$(x_1 - y_2)^2$	0	0
$[y_2, x_1]$	$[0, y_2]$	$(y_2)(x_1 - y_2)$	+2	$2(y_2)(x_1 - y_2)$
$[0, y_2]$	$[y_2, x_1]$	$(y_2)(x_1 - y_2)$	-1	$-(y_2)(x_1 - y_2)$
$[0, y_2]$	$[0, y_2]$	y_2^2	- g	$-gy_2^2$




Figure 18.5. Game Value for [0,1] Game #6 in special region

After one more raise, this game is the same as in [0,1] Game 6, except that the situation has been reversed! Both players have a hand on a particular interval, and the player who made the last raise has equity $-g$.

We can now find g :

$$\begin{aligned}
 g &= y_2(x_1 - y_2) - gy_2^2 \\
 g &= rx_1(x_1 - rx_3) - gr^2x_1^2 \\
 gr^2x_1^2 + g &= rx_1(x_1 - rx_1) \\
 gr^2 + g &= r - r^2 \\
 g(1+r^2) &= r - r^2 \\
 g(2(1 - r)) &= r - r^2 \\
 g &= r/2
 \end{aligned}$$

So the equity in this region where both players have strong hands is $r/2$.

Filling this value into our matrix, we have:

	$[0, x_1]$	$[x_1, y_1]$	$[y_1, 1]$
$[0, x_1]$	$r/2$	+1	+1
$[x_1, y_1]$	-1	0	+1
$[y_1, 1]$	-1	0	0

Taking the weighted average equity of the game, we have:

$$\begin{aligned} \langle Y \rangle &= (1 - y_1)(y_1 - x_1) + rx_1^2/2 \\ \langle Y \rangle &= (1/4)(1/4) + r(1/8) \\ \langle Y \rangle &= 1/16 + r/8 \\ \langle Y \rangle &\approx 0.11428 \end{aligned}$$

We can compare this value to the value $1/8$ obtained from the first game in this chapter, where only two bets were allowed without check-raise. This game is slightly more favorable for X than that one. This is of course due to the fact that X can reraise after Y raises. Even then, Y can put in the fourth bet, X the fifth, and so on, but each succeeding bet has a much smaller impact on the EV. The relatively small change that the addition of infinite bets makes on X's equity can be seen by calculating $x_3 = r^2/(1 + 2r) = 0.09384$. In order for X to put the third bet in, Y must be above y_2 , which occurs only about 22% of the time and X must be above x_3 a parlay of just 2%.

The last of the no-fold games is the infinite bet game with check-raise allowed (see Example 16.4). We can solve this game in a like manner. Recall that in the two-bet check-raise game, X was indifferent between check-raising and betting with all the hands stronger than y_2 in the previous check-raise game. It turns out that this remains the case in the infinite bet game; X is indifferent between check-raising and betting between any two Y thresholds. Hence as our best response strategy for X we will play the simple strategy:

$$\begin{aligned} x_1 &= y_1 \\ x_3 &= y_3 \end{aligned}$$

... and so on.

In the region where both players are better than y_1 , we again find (as in the last game) that the equity of that region is $r/2$. We invite you to verify that this is the case, as well as verifying X's indifference between his various strategy options.

This creates a simple matrix like the previous one:

	$[0, y_1]$	$[y_1, 1]$
$[0, y_1]$	$r/2$	$+1$
$[y_1, 1]$	-1	0

Taking the weighted average equity of this game, we have:

$$\begin{aligned} \langle Y \rangle &= ry_1^2/2 \\ \langle Y \rangle &= r/4 \\ \langle Y \rangle &\approx 0.10355 \end{aligned}$$

Again, we see that the addition of check-raise benefits the first player only a small amount, moving Y's equity from about 0.114 to about 0.104. By contrast, allowing X to bet immediately decreased Y's equity from $1/4$ to $1/8$. We will return to comparing the values of these games shortly, after quickly considering the equity of the finite pot games.

The simplest finite pot game we considered was the half-street game where folding is allowed, [0,1] Game #2 (Example 11.3), whose solution was as follows:

$$y_1 = (1 - \alpha)/(2 - \alpha)(\alpha + 1)$$

$$x_1^* = 2y_1$$

$$1 - y_0 = \alpha y_1$$

To solve this game, we will hold Y's strategy constant and find a best response strategy for X. We will use the notation G_n to represent the value in regions beginning with the weakest ($n = 0$) and proceeding to the strongest.

When X is on $[0, y_1]$, he obviously achieves maximum value by calling, since he even beats some of Y's value bets. So here X calls. He breaks even when Y has a value betting hand, and gains one bet against Y's bluffs.

$$G_2 = \langle X \text{ on } [0, y_1] \rangle$$

$$G_2 = -\alpha y_1$$

When X is on $[y_1, y_0]$, he is actually indifferent with all these hands to calling or folding. He wins against bluffs and loses against value bets. Since he is indifferent, we can arbitrarily assume that he folds. Then Y gains nothing with value bets and wins the pot of P units with bluffs.

$$G_1 = \langle X \text{ on } [y_1, y_0] \rangle$$

$$G_1 = P\alpha y_1$$

$$G_1 = (1 - \alpha)y_1$$

When X is on $[y_0, 1]$, he folds, and Y wins the pot. Y gains half a pot when he is on the same interval as X, so:

$$G_0 = \langle X \text{ on } [y_0, 1] \rangle$$

$$G_0 = P\alpha y_1/2$$

$$G_0 = y_1(1 - \alpha)/2$$

To find the overall equity of the game, we can take the weighted sum of these over the interval lengths:

$$G = y_1 G_2 + (y_0 - y_1) G_1 + (1 - y_0) G_0$$

$$G = y_1(-\alpha y_1) + (1 - y_1 - \alpha y_1)(1 - \alpha)y_1 + \alpha y_1^2(1 - \alpha)/2$$

$$G = (1 - \alpha)y_1 - (2 - \alpha)(1 + \alpha)y_1^2/2 \tag{18.1}$$

We can further simplify this by using the value of y_1 and obtain:

$$G = y_1(1 - \alpha)/2$$

$[0,1]$ Game #9 (Example 17.1) was the finite pot game for a full street. The game value can be found through a similar process. In the region $[0, y_1]$, we can find that the value of X's checks or bets against Y's optimal strategy is the same:

Y's hand	X checks	X bets
$[0, x_1]$	+1	+1
$[x_1, y_1]$	-1	-1
$[y_1, y_1^*]$	0	-1
$[y_1^*, y_0]$	0	0
$[y_0, 1]$	-1	0

The solution to this game (for Y) was:

$$y_1 = (1 - a)/(1 + a)$$

$$y_0 = (1 + a_2)/(1 + a)$$

$$y_1^* = 1 - a$$

We see that Y gains a bet in the region $[y_1, y_1^*]$ when X checks, but loses a bet in the region $[y_0, 1]$, We can show that these regions are equally sized as follows:

$$y_1^* - y_1 = (1 - a) - (1 - a)/(1 + a) = (a - a_2)/(1 + a) = ay_1 = 1 - y_0$$

So X is indifferent to checking or betting against Y's optimal strategy with a hand on $[0, y_1]$. This is convenient as it allows us to reuse all the work from the previous game - the value of this game is equivalent to the value of Game #2 (Example 11.3) where X checks dark except that the value of y_1 is different.

Equation 18.1 holds here:

$$G = (1 - a)y_1 - (2 - a)(1 + a)y_1^2/2 \tag{18.1}$$

Instead of substituting the value of y_1 from the original game, we substitute the new value of y_1 from this game:

$$y_1 = (1 - a)/(1 + a)$$

$$G = (1 - a)^2/(1 + a) - (2 - a)(1 - a)^2/2(1 + a) \tag{18.1}$$

$$G = (1 - a)^2(2 - 2 + a)/2(1 + a)$$

$$G = a(1 - a)^2/2(1 + a)$$

The third finite-pot game we looked at was the game where Y could raise but there was no check-raise (Example 17.2). For this game, we can see that the equations for the following values are the same as in the one-bet case:

$$G_1 = \langle X \text{ on } [y_1, y_0] \rangle$$

$$G_1 = \text{Pay}_1$$

$$G_1 = (1 - a)y_1$$

$$G_0 = \langle X \text{ on } [y_0, 1] \rangle$$

$$G_0 = \text{Pay}_1/2$$

$$G_0 = y_1(1 - a)/2$$

For G_2 , the case where X is on $[y_1, y_2]$ we can examine X's options. We know that if X bets in this region, he is indifferent to his choices when Y raises. Y will always have him beat with value raises, and X will always beat bluffs. So he is indifferent between bet-folding and bet-calling in this region. Comparing the value of X check-calling in this region to his bet-calling, we have:

Y's hand	X check-calls	X bet-calls
$[0, y_2]$	+1	+2
$[y_2, y_1]$	0	0

$[y_1, y_1^*]$	0	-1
$[y_1^*, y_2^\#]$	0	-2
$[y_2^\#, y_0]$	0	0
$[y_0, 1]$	-1	0

Again, we can show that X is indifferent between these choices as follows. One of the indifference equations from this game was:

$$y_1^* - y_1 = (P + 1)(y_2^\# - y_1^*) + (1 - y_0)$$

For bet-calling, X gains:

$$2(y_2^\# - y_1^*) + y_1^* - y_1 = (P + 3)(y_2^\# - y_1^*) + 1 - y_0 = y_2 + ay_1$$

From check-calling, X gains:

$$(y_2 - 0) + (1 - y_0) = y_2 + ay_1$$

Hence on the interval $[y_2, y_1]$ X is indifferent between check-calling, bet-folding, and bet-calling. To those of you who are interested in solving additional $[0,1]$ games, this type of relationship (which we call a **triple indifference**) is common in more complex finite pot games, and is often helpful in solving these games with fewer and simpler indifference equations.

We can most easily find Y's equity when X is on this interval by assuming X checks and calls. Then we have:

$$G_2 = \langle X \text{ on } [y_2, y_1] \rangle = y_2 - ay_1$$

With hands stronger than y_2 , X's best option is to bet-call. By doing so, he breaks even when Y raises for value, gains one bet when Y calls, and gains two bets when Y raise-bluffs. Y's equity, then, is:

$$G_3 = \langle X \text{ on } [0, y_2] \rangle = -(y_1^* - y_2) - 2(y_2^\# - y_1^*)$$

$$G_3 = y_2 - y_2^\# - (y_2^\# - y_1)$$

$$G_3 = y_2 - (1 - a) - a_2y_2$$

$$G_3 = (1 - a_2)y_2 - 1 + a$$

Summing the weighted values of these four regions, we have:

$$\langle Y \rangle = y_2G_3 + (y_1 - y_2)G_2 + (y_0 - y_1)G_1 + (1 - y_0)G_0$$

$$\langle Y \rangle = y_2G_3 + (y_1 - y_2)G_2 + (y_0 - y_1)G_1 + (1 - y_0)G_0$$

$$\langle Y \rangle = y_2(G_3 - G_2) + y_1G_2 + (y_0 - y_1)G_1 + (1 - y_0)G_0$$

$$\langle Y \rangle = -(1 - a)y_2 - a_2y_2^2 + (1 + a)y_1y_2 + (1 - a)y_1 - (1 + a)(2 - a)y_1^2/2$$

For $[0,1]$ Game #11 (Example 17.3), the check-raise game with two bets, we have an analogous situation to the no-fold game. Holding Y's strategy constant, we can show that X is indifferent to check-raising or bet-calling in the region $[0, y_2]$. Hence the value of Game #11 follows the same formula as the value of Game #10; of course the values of y_1 and y_2 are different, and hence Y's equity is decreased.

We're going to stop deriving the game values here; as you can see, the values and the process of deriving them scale in complexity along with the games. However, these methods can be used to

find the value of any game, so we encourage you to attempt to find the values and solutions to other [0,1] games.

Lessons of the [0,1] Game

Now, let's take a deep breath, and review what all these values mean.

Take a look at the following figure, which shows some of the relationships between these games:




Figure 18.6. Game Value for various [0,1] Games

We can see in this figure that some of these games are in fact special cases of each other; for example, Game #1 is actually a special case of the more general Game #2. To obtain Game #1, we simply let the pot become arbitrarily large. This means that the value of α goes to zero. This relationship holds between any finite pot and its no-fold counterpart; we can simply find the value of the game for the finite pot game and take the limit as the pot goes to infinity to find the value of the no-fold game. There are other relationships between the games as well. We can see that some games are obtained by adding strategic alternatives to a different game. For example, Game #7 is simply Game #5, but we give X the opportunity to check-raise. By comparing the equities of the games, we can see the verification of a principle we stated earlier:

Strategic options have non-negative value.

Throughout these games, we find examples of this. Consider the half-street games #1 and #2. In Game #2, the primary change that occurs is that X gains the strategic option of folding to Y's bet. Hence the value of Game #2 (from Y's perspective, remember) is reduced. Likewise for all of the following transitions:

Between Game #1 and Game #4, where X gained the option to bet.

Between Game #4 and Game #5, where Y gained the option of raising.
 Between Game #5 and Game #6, where X gained the option of reraising.
 Between Game #5 and Game #7, where X gained the option of check-raising.

In all of these cases, and in their counterparts in the finite pot games, the player who gained a strategic option gained (or at least retained their previous) equity.

The second principle that we find in the [0,1] game is the value of r . In real poker, there are card removal effects, the effects of multi-street play, and so on that subtly change the value of r . But the idea that you should raise your opponent on the river with about 41% of the hands that he would bet is a practical one.

Consider the following game. Suppose you and a friend decide to play a variant of holdem where you each receive two cards, but then the entire board is dealt, and there is one round of betting (the river). Suppose you structure the betting so that Y can either bet or show down, then X can raise and you allow unlimited raises. The optimal strategy for this game is something roughly like this: the first player bets r of his hands; the second player raises with r^2 of the hands, the first player reraises with r^3 of the hands, and so on. Suppose the board was $J♥ 9♣ 8♠ 3♦ 2♣$. Then the hands could be ranked from top to bottom by value. Each player can have a total of 1,081 hands.

Hand	Number of Ways	Cumulative # of ways
QT	16	16
T7	16	32
JJ, 99, 88, 33, 22	30	62
J9, J8	18	80
J3 or worse two pair	72	152
AJ, KJ, QJ	36	188
JT, J7, J4	60	248
A9-Q9, T9, 97-94	96	344
A9-Q9, T9, 97-94	96	440
Other hands	641	1081

- (1081) $r \approx 448$
- (1081) $r^2 \approx 185$
- (1081) $r^3 \approx 77$
- (1081) $r^4 \approx 32$
- (1081) $r^5 \approx 13$

In doing true optimal strategies for this game, we would be forced to consider card removal effects and account for the fact that the hands are not uniformly distributed, etc. Obviously, you would never stop raising with the nuts in real poker, and so on. But equally, we can see that in the above example, the approximate cutoffs for putting in the n^{th} bet are third pair or better for the first bet, QJ or better for the second bet, first and third pair for the third bet, the second nut straight for the fourth bet, and the nuts for the fifth.

Can we apply this directly to holdem? Not precisely, because the play on previous streets and the resultant distribution of hands changes how the hands are allocated. But using this guideline of 41% (if your opponent cannot, fold - for example the pot is too big and he cannot have a draw) is

a reasonable guess. In games where the opponent can fold, we can adjust this number downward (because the opponent still raises with his best hands but can now fold his worst), to about 35%.

This last example, in addition to showing how we might apply the value of r to real poker, provides an introduction to the last and most important lesson of the $[0,1]$ game. In all the $[0,1]$ games, hands were partitioned into regions, where each player took various actions. We saw that at the threshold points, the players were indifferent to taking the actions on either side. But they are not indifferent to taking all actions most of the time.

When we calculated some of the game values, we held Y's strategies constant and found that in most of his value regions, X was indifferent to taking various value actions. This means that as a practical matter, if Y plays optimally, X has some latitude to make mistakes near his threshold values without it costing much. We call this type of strategy *flexible*. For Y, however, this is not the case. If X plays optimally and Y fails to bet hands a little better than y_1 , Y loses equity immediately. We call strategies like Y's *brittle*. X's strategies are often flexible in $[0,1]$ games because with hands near his betting threshold, if he should check incorrectly, Y will often put in the bet anyway. Such a situation doesn't exist for Y, however.

When you play a side whose strategies are brittle (which happens frequently to both sides of real poker), the selection of hands with which you take each action is very important for optimal play. Organizing hands into regions is both a logical and practical step as well as a high-expectation practice. What this means is that typically in real poker, we will identify cutoffs between various actions and have clear ideas about what our intent in making such a bet or raise is. On the river, these ideas are primarily centered around how many bets we are willing to put in, and whether we will call or fold to a raise. On previous streets, these decisions are primarily oriented toward whether we are betting for value or semi-bluffing, and how we will respond to future action.

Additionally, this means that we typically do not play mixed strategies with very many hands: that is, at the boundaries there are mixed strategies, and there may occasionally be reason to play mixed strategies for the purpose of hiding card removal effects, but normally hands adjacent in value are played in the same way.

Also to be noted is the fact that the $[0,1]$ strategies for finite pot games are filled with regions that exist primarily to make the strategies balanced. For each value bet, there is a corresponding bluff, and a set of hands within the value region that will call a raise and some that will fold a raise. Thinking about real poker play in this way is a very strong tool for enabling you to play in a balanced manner. For every value bet you make, consider the bluffing region that corresponds. How big is it? Are the ratios proper? When raising, do you raise-bluff? With which hands do you do so? Questions like these are the kind that lead to balanced play that is difficult to exploit.

We close this chapter with a look at the last $[0,1]$ game in this book, which provides an interesting result related to high-low games.

Example 18.1 - The $[0,1] \times [0,1]$ Game

In this variant of the $[0,1]$ game, each player receives two real numbers independently and uniformly distributed from 0 to 1. One of these numbers is pre-designated to stand in as his "high" hand, and the other as his "low" hand. Both hands are subject to the usual rules of the $[0,1]$ game; the lower "high" number gets half the pot and the lower "low" number gets half the

pot. As a result, each possible pair (x, y) where x and y are between 0 and 1 represents a possible hand in this game.

For this game we stipulate that neither player may fold, there are two bets left, and check-raise is not allowed. The pot size, as is usual in games where neither side may fold, is irrelevant.

A high percentage of bets in this game are in fact a type of bet that we haven't discussed yet. These bets are called *semi-bluffs*, where a player makes a bet expecting to often win half the pot when called. These bets are extremely common in high-low games, but occur in holdem as well. (As an example consider the play on the river of a two-paired board. Ace-high will often bet expecting to often win half the pot when called by another ace-high.) In high-low games, however, bets of this nature, where a player has a reasonably strong hand one way and a mediocre hand the other make up a substantial part of all bets.

One way of looking at the game is by graphing the game as xy -coordinate pairs in a plane. Instead of a single threshold, however, for each player's strategy, we have a boundary that encloses a region for an action. We initially parameterize the game as having three curves - we call them x_1 , y_1 ; and y_2 by analogy with [0,1] Game #5 (Example 16.2). We also employ the terms "below" the curve as referring to the area closer to $(0,0)$ and "above" the area closer to $(1,1)$. These three curves, then, are related as follows:

$$(0, 0) < y_1 < x_1 < y_2 < (1, 1)$$

We have indifference equations at each boundary as well. However, generating these indifference equations is a little more complex than previously. In this case, we will begin by assuming X's strategy first, deriving Y's strategy from that, and then retaining and trying to improve X's equity by varying his strategy to prove that the strategy we assumed for X was in fact optimal.

Assumption: x_1 lies on the line $y = 1 - x$.

This assumption may perhaps make intuitive sense to the reader, by analogy with x_1 in the one-dimensional [0,1] games that was $1/2$ in the no-fold-allowed case.

Next consider the boundary y_2 . Y must be indifferent between calling (above y_2) and raising (below y_2). The next step is to identify the regions that generate expectation for either player. Note that in many regions the pot is split and neither player gains. For any point (x, y) on y_2 :

When X's hand in is a rectangle with vertices $\{(0, 0), (0, y), (x, 0), (x, y)\}$, Y loses one bet by calling, but two by raising. This rectangle is the only place where X scoops Y. Y in turn scoops X in a triangular region bounded by the vertices $\{(x, 1 - x), (1 - y, y), (x, y)\}$. See Figure 18.7.




Figure 18.7. Region comparisons for hands in the $[0,1] \times [0,1]$

Each of these regions is worth one bet, so setting their areas equal:

$$\begin{aligned} xy &= \frac{1}{2} (1 - y - x) (1 - x - y) \\ 2xy &= 1 - y - x - y + xy + y^2 - x - x^2 + xy \\ 1 &= x^2 + y^2 - 2x - 2y + 2 \\ 1 &= (x - 1)^2 + (y - 1)^2 \end{aligned}$$

Students of algebra will recognize this as the equation of a circle of radius 1, centered at (1, 1). So we find that y_2 is a quarter-circle of radius 1 centered at (1, 1).

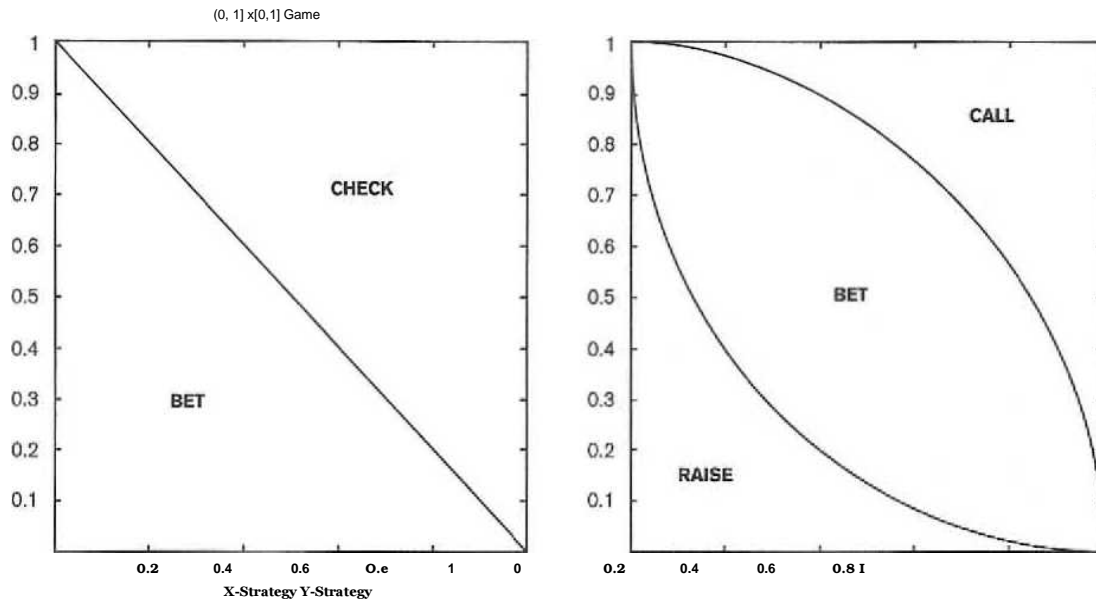
Turning to y_1 , we look at Y's indifference between checking and betting; Y gains a bet by betting when X is in a rectangle with vertices $\{(1, 1), (1, y), (x, 1), (x, y)\}$ and loses a bet when betting in a triangle with vertices $\{(x, y), (x, 1 - x), (y, 1 - y)\}$.

Setting the areas of these regions equal:

$$\begin{aligned} (1 - x)(1 - y) &= \frac{1}{2} (y - (1 - x))(x - (1 - y)) \\ 1 - x - y + xy &= \frac{1}{2} (y - 1 + x)^2 \\ 1 - x - y + xy &= \frac{1}{2} (y^2 - y + xy - y + 1 - x + xy - x + x^2) \\ 2 - 2x - 2y + 2xy &= y^2 - 2y + 2xy - 2x + x^2 + 1 \end{aligned}$$

$$1 = y^2 + x^2$$

This is the equation of a circle of radius 1 centered at (0, 0). Plotting these three graphs, we see that the strategies form a "football" about the centerline $y = 1 - x$.



We now return to our assumption and show that x_1 does in fact lie on the line $y = 1 - x$. To do this, we can use the fact that we know that along x_1 , X is indifferent to betting or checking. Along this boundary, X gains a bet by checking (over betting) when Y is below y_2 , because he does not get raised. At the same time, he loses a bet when Y lies above y_1 because he fails to gain a bet from Y. Of course, for some hands above y_1 and below y_2 , there is no difference because the pot is split.

We know that the size of the two regions A and B must be equal. Intuitively we can see that along the diagonal from (0, 1) to (1, 0) there will be symmetry between the two regions. Since the diagonal is actually $y = 1 - x$, we have our optimal strategies.

This gives rise to at least one interesting answer: if it's a two-bet high-low game with no qualifier, no card removal effects, no check-raise, and neither player may fold, what percentage of hands should the second player bet? The answer is $\pi/4$!

These restrictions are quite onerous and so this game is not a true analogue to real-life high-low games. However, it is at least sometimes the case that when playing high-low games (for example stud), we find ourselves in a position very much like the "no fold" restriction. The distribution of hands with which we have arrived at seventh street is such that neither player can fold, due to the high probability that the opponent has a one-way hand. Even weakened in this way, this result indicates a significantly higher betting frequency than is commonly thought, particularly by the player who is acting in last position.

Key Concepts

- Comparing the values of various games gives us insight into the importance of various concepts.
- Adding strategic options to a one player increases the value of the game for that player.
- When considering the value of a game, we can often hold one player's strategy constant and allow the other player's strategy to maximally exploit it in order to simplify the game value calculation.
- Some strategic options are brittle - that is, small deviations from optimal cost equity immediately. Others are flexible - that is, against the optimal strategy, they can move around in a range.
- Strategic options have non-negative value.

Chapter 19

The Road to Poker: Static Multi-Street Games

In the preceding chapters, we have considered a wide variety of games of the half-street and full-street varieties. One of the primary virtues that these games have is that they are relatively simple to solve; and games that we can solve can give us insight into the impact of changes to a particular dimension of the game. However, real poker is seldom played in half-street or full-street ways, and as we shall see in this chapter and the next, adding additional streets changes strategy in fundamental ways.

We begin by considering what we call *static games*. These are games where there are multiple streets of betting, but neither player's hand changes value. These games give us insight into how bluffing works across multiple streets without considering the complexity of draws. After covering a few games of this type, we will move on to the more difficult and realistic games where draws are present.

The most basic static game of two streets that corresponds with our treatment of half street and full street games is the clairvoyant game.

Example 19.1 - The Two-Street Static Clairvoyant Game

Two full streets.

Static hand values.

Pot size of P bets.

Limit betting.

Y is clairvoyant.

Recall that in the one-street clairvoyant game, X never bet because such bets would never have positive value. Since Y was clairvoyant, he could exploit any such bets that X made with perfect information. The same situation occurs here. On either street, X will only check and call because Y can respond to any bet with perfect accuracy. This is the case here only because the game is static; that is, Y knows his ending hand value now. If there were draws where Y did not know the outcome, there might be situations where X would bet.

In the one street game, we found that the solution to the clairvoyance game was that Y would bet all his winning hands, and bluff with a fraction a of his losing hands. X would respond by folding a of his hands, and calling with the rest. This was true except for the cases where Y had such a strong distribution of hands that he could bet all hands and X would be compelled to fold.

In the two-street game, we have an additional street of betting. This creates some additional strategic options for the two players. First consider Y's options.

Check both streets (strategy Y_0)

Check the first street and bet the second (strategy Y_1)

Bet the first street and check the second (equivalent to strategy Y_1)

Bet both streets (strategy Y_2)

Clearly Y will prefer to bet both streets with his nut hands; this maximizes value. Hands that employ strategies Y_1 and Y_0 will be bluffs, as well as some hands that employ strategy Y_2 . Y will also prefer to play the second version of Y_1 where he bets the first street and checks the second. The check-first version is of course exploitable because X will know that Y does not have a nut hand because of his initial check.

X's strategy must make Y indifferent between these three options when Y has a dead hand. X has three options in response to Y's bets:

Fold immediately (strategy X_0)

Call the first street and fold the second (strategy X_1)

Call both streets (strategy X_2).

We have the following matrix (values from Y's perspective):

	X_0	X_1	X_2
Y_0	0	0	0
Y_1	+P	-1	-1
Y_2 (bluff)	+P	+(P+1)	-2
Y_2 (value)	0	+1	+2

When Y uses strategy Y_2 , we will call his fraction of value betting hands y_v and his fraction of bluffs y_b . We will call X's frequency of using his three strategies x_0 , x_1 , and x_2 , and Y's frequency of using his three strategies y_0 , y_1 , and y_2 .

Then we have the following indifference between Y_1 and Y_0 :

$$0 = Px_0 - x_1 - x_2$$

We know that $x_1 + x_2 = 1 - x_0$, because the total frequencies for both players must sum to one. Hence:

$$0 = 1 - x_0 - Px_0$$

$$x_0 = 1/(1+P) = \alpha$$

So on the first street, X plays a strategy that is the same as his strategy in the one-street game. He folds α of his hands, and calls with the remainder.

Next, between Y_1 and Y_2 (bluff), we have:

$$aP - x_1 - x_2 = aP + (P+1)x_1 - 2x_2$$

$$(P+2)x_1 = x_2$$

We have $x_1 + x_2 = 1 - a$, so:

$$1 - a - x_1 = (P+2)x_1$$

$$x_1 = a_2(1 - a)$$

On the second street, as well, X plays a strategy that is very much like his strategy in the single street game. After folding α of his hands on the first street, he has $1 - \alpha$ hands left. The pot has grown by two bets because of the betting on the first street, so on the second street he now folds a_2 of his hands, and calls with the remainder. This is a very important idea:

The non-clairvoyant player plays each street independently as if it were a single street game.

We now turn our attention to Y's strategy of betting and bluffing. Y must make X indifferent between his three strategies x_0 , x_1 , and x_2 . We have the following indifference between x_0 and x_1 .

$$\begin{aligned} P y_1 + P y_b y_2 &= -y_1 + (P + 1) y_b y_2 + y_v y_2 \\ (P + 1) y_1 &= y_v y_2 + y_b y_2 \\ y_1 &= a y_2 \end{aligned}$$

This indifference shows the second important result of this game.

The hands that Y bluffs once and then gives up on have the a bluffing ratio relationship to the total hands that Y will bet for value and bluff on the second street.

Next, the indifference between X_1 and X_2 :

$$\begin{aligned} -y_1 + (P + 1) y_b y_2 + y_v y_2 &= -y_1 - 2y_b y_2 + 2y_v y_2 \\ y_v y_2 &= (P + 3) y_b y_2 \\ y_b &= a_2 y_v \end{aligned}$$

So on the second street, Y value bets and bluffs in the typical bluffing ratio out of epe hands he has left. Here the ratio is a_2 because the pot is two bets larger.

Let's recap this idea, because it is an important one. On the final street, Y will have a set of betting and bluffing hands appropriate to the pot size at that point. On the preceding street, he "pretends" that all those hands are value hands and bluffs appropriately to the current pot size. Then on the final street, he gives up on his "bluffs" from the first street, and only continues to bluff out of the "value hands" he carried forward from the last street.

This often leads to Y bluffing a larger number of hands than we saw in the one-street games. For example, suppose the pot is 2 bets before the first street and Y has just a small number of winning hands. Then we have the following:

On the river, the pot will be 4 bets. Y will want to value bet all his winning hands and $1/5$ as many bluffs on the river. So he "pretends" that his value bets plus another $1/5$ as many bluffs are value hands. On the first street, the pot is 2 bets, so he wants to bluff with his "value hands" and $1/3$ as many dead hands. Since he has $6/5$ as many "value hands," he bluffs with the intention of giving up with $2/5$ as many hands as he has nut hands on the first street. The total number of dead hands he bets on the first street is $3/5$ as many as he has value hands.

He then bets all these hands, gives up on $2/3$ of his dead hands after the first street, and continues with $1/3$ of the dead hands and all the value bets on the second street. The second player cannot exploit this strategy, as we showed that he is indifferent between calling once, twice, or not at all.

This result generalizes to any number of streets and to any bet size, subject to the usual restrictions on strength of distribution. Suppose that we play the following game:

X and Y ante 2 units.

Both players receive hands, but both hands are revealed only to Y.

A complete board is dealt but revealed only to Y.

There are three streets - a flop where the bets are 1 unit, and a turn and river where they are 2 units.

If there's a showdown, the board and both hands are exposed and the showdown is as usual.

Again in this game, X cannot profitably bet on any street; the solution is just as it was in the two-street game, except with an additional street and different bet sizes.

With winning hands, Y will bet on each street, and the pot size going into the river will be $4 + 2 + 4 = 10$ units, with a bet size of 2. On the river, Y will bet with the $\frac{1}{2}$ of hands with which he will win, plus an additional $\frac{2}{12}$ as many, for a total of $\frac{7}{12}$. This makes a total of $\frac{7}{12}$ of his hands.

On the turn, the pot will be 6 bets, with a bet size of 2. So Y will have $\frac{7}{12}$ of his hands designated "Tor value," and will additionally bluff $\frac{1}{4}$ as many hands as that, or an additional $\frac{7}{48}$. So his total betting hands on the turn will be $\frac{35}{48}$.

On the flop, the pot will be 4 bets, with a bet size of 1. Y will take his $\frac{35}{48}$ "value bets" and bluff an additional $\frac{1}{5}$ of that, or $\frac{7}{48}$. So his total betting hands on the flop will be $\frac{42}{48} = \frac{7}{8}$.

The solution is that Y bets $\frac{7}{8}$ of his hands total and gives up on $\frac{1}{12}$ of his dead hands after the flop and $\frac{7}{48}$ of his dead hands after the turn. X simply folds α of his hands on each street. The value of this game is easily found by assuming X folds immediately (since he is indifferent). $\frac{1}{8}$ of the time, X wins the pot for a net to Y of 1.5 units. It's amazing that X can retain any equity at all in such a biased game.

We said that this extended to all bet sizes. In games where the clairvoyant player controls the bet size (such as no-limit), we can use the maximization techniques from previous chapters to find optimal bet sizes.

Example 19.2 - The Two-Street Static Clairvoyant No-Limit Game

Two full streets.

Static hand values.

Pot size of 1 unit.

No-limit betting (stacks of N units).

Y is clairvoyant.

We'll first consider the two-street case; the idea here generalizes to any number of streets. We saw in the previous game that the non-clairvoyant player plays each street as if it were a single-street game, folding α of his hands and calling with the remainder. Likewise, the clairvoyant player bets and bluffs using α -ratios to the hands he carries forward from street to street.

These relationships hold even when the bet sizes are not constant and can be determined by the clairvoyant player. As a result, we can find the clairvoyant player's equity in terms of bet sizes s_1 and s_2 and attempt to maximize this value. We have a matrix similar to the previous game, supposing the pot is 1 unit.

	X_0	X_1	X_2
Y_0	0	0	0
Y_1	+1	$-s_1$	$-s_1$
Y_2 (bluff)	+1	$+(1 + s_1)$	$-(s_1 + s_2)$
Y_2 (value)	0	$+s_1$	$+(s_1 + s_2)$

Writing the same indifference equations as in the previous game, we have:

$$X_0 - s_1 X_1 - s_1 X_2 = 0$$

$$x_0 = s_1(1 - x_0)$$

$$X_0 = s_1 / (1 + s_1)$$

Again, we see that X folds α of his hands (recall that $\alpha = s/(1 + s)$ as well as $1/(P + 1)$).

$$\begin{aligned} x_0 + (1 + s_1)x_1 - (s_1 + s_2)x_2 &= x_0 - s_1x_1 - s_1x_2 \\ (1 + 2s_1)x_1 &= s_2x_2 \\ x_2 &= x_1(1 + 2s_1)/s_2 \\ x_2 + x_1 &= x_1(1 + 2s_1 + s_2)/s_2 \end{aligned}$$

Now let's consider what α_2 is for this game.

After the first street, the pot size will be $1 + 2s_1$ (the original pot, plus a bet and a call). So α_2 would be $s_2/(1 + 2s_1 + s_2)$. Substituting this into the above equation, we have:

$$\begin{aligned} x_2 + x_1 &= x_1 / \alpha_2 \\ x_1 &= \alpha_2(x_2 + x_1) \end{aligned}$$

We saw this result previously. X plays each street as though it were a single street game, folding an appropriate α of his hands on each successive street. Of all the hands that call one bet ($x_2 + x_1$), X folds α_2 for the second bet.

The indifferences for X are likewise the same:

	X_0	X_1	X_2
Y_0	0	0	0
Y_1	+1	- s_1	- s_1
Y_2 (bluff)	+1	+ (1 + s_1)	- ($s_1 + s_2$)
Y_2 (value)	0	+ s_1	+ ($s_1 + s_2$)

$$\begin{aligned} y_1 + y_2y^b &= -s_1y_1 + (1 + s_1)y_2y^b + s_1y_2y^v \\ (1 + s_1)y_1 &= s_1y_2y^b + s_1y_2y^v \\ y_1 &= y_2s_1 / (1 + s_1) \\ y_1 &= \alpha y_2 \end{aligned}$$

Here we see, as before, that the number of hands that Y bluffs is equal to α times the number of hands he will both bluff and value bet on the river.

$$\begin{aligned} -s_1y_1 + (1 + s_1)y_2y^b + s_1y_2y^v &= -s_1y_1 - (s_1 + s_2)y_2y^b + (s_1 + s_2)y_2y^v \\ (1 + s_1)y_2y^b + s_1y_2y^v &= - (s_1 + s_2)y_2y^b + (s_1 + s_2)y_2y^v \\ (1 + 2s_1 + s_2)y_2y^b &= s_2y_2y^v \\ (1 + 2s_1 + s_2)y^b &= s_2y^v \\ y^b &= \alpha_2y^v \end{aligned}$$

And on the river, the bluffing ratio is simply α_2 . So this solution holds for any bet sizes (provided that the distributions are not so strong). To find the proper bet sizes s_1 and s_2 , we maximize Y's equity for the two bet sizes.

We can make a nice simplification that additionally helps us understand the principle of optimal bet sizing. Suppose that Y bets s_1 into a pot of 1. If X were to call, then the pot would grow on this street from 1 to $(1 + 2s_1)$. We define r_1 as this ratio; the ratio that the pot will grow if the bet s_1 is called, in a similar fashion, we can define r_2 as the ratio by which the pot will grow if the bet s_2 is called, from $(1 + 2s_1)$ to $(1 + 2s_1 + 2s_2)$, or

$(1 + 2s_1)/(1 + 2s_1 + 2s_2)$ and so on:

$$r_n = (p_n + 2s_n)/p_n$$

Note that this r isn't necessarily the same as the previously discussed r (Equation 14.1).

We know from the above definition of r that for any street, the following is true:

$$r = 1 + 2s$$

(where 1 is the pot size coming in to this street).

From this, we can derive the following useful transformation:

$$r - s = 1 + 2s - s$$

$$r - s = 1 + s$$

$$r - s = (2 + 2s)/2$$

$$r - s = (1 + r)/2$$

We can also find the following:

$$\alpha = s/(1 + s)$$

$$\alpha = s/(r - s)$$

Let M_n represent the amount of hands which are "for value," whether this means value bets on the next street or if this is the final street, hands which are the nuts. On any street, the amount of hands bet (bluffs and value bets) B_n for that street is:

$$B_n = s_n M_n / (r_n - s_n) + M_n$$

$$B_n = M_n (r_n - s_n + s_n) / (r_n - s_n)$$

$$B_n = M_n r_n / (r_n - s_n)$$

$$B_n = 2M_n r_n / (1 + r_n)$$

This is an important result; it provides a "factor" that we can apply to the streets in order to find the number of bluffs on each street.

On the second street of our game, then, the total number of hands bet (if Y has y winning hands overall) will be:

$$B_2 = 2y r_2 / (1 + r_2)$$

while the hands bet on the first street will be:

$$B_1 = 2B_2 r_1 / (1 + r_1)$$

$$B_1 = 4y r_1 r_2 / (1 + r_1)(1 + r_2)$$

Let's review all that math to distill it to a couple of results.

We defined values r_n that are the ratio of the pot after a bet and a call on that street to the pot before that street.

We found that for each street, we can express the total number of hands that are bet (bluffs or value bets) with the following formula:

$$B_n = 2M_n r_n / (1 + r_n)$$

where M is the number of hands that will be carried forward and bet on the next street (if there are streets remaining) or that are nut hands (if this is the final street).

If Y bets s_1 and s_2 at optimal frequencies, then X will be indifferent to calling or folding on any street. We can assume, therefore, that he simply folds immediately, which means that Y wins the pot with all his bluffs. We also know that X won't actually fold all his hands - in fact, he will call optimally. Since this is the case, Y will want to maximize the total amount he bets overall when he has a nut hand (and therefore also when he bluffs). As a result of this, we know that $s_1 + s_2$ will equal a constant of N; that is, Y will bet his entire stack on the two streets when he bets twice. Given that, we can calculate Y's equity by assuming that Y does not fold. Since the pot is one unit, Y maximizes his value by maximizing the fraction of hands that he bets.

The total number of hands Y bets is:

$$B_1 = \frac{4y r_1 r_2}{(1 + r_1)(1 + r_2)}$$

The $4y$ term is of course a constant, so we need to maximize the following expression:

$$r_1 r_2 / (1 + r_1)(1 + r_2)$$

$r_1 r_2$ is a constant; we can see this as follows:

$$\begin{aligned} r_1 &= (1 + 2s_1) \\ r_2 &= (1 + 2s_1 + 2s_2) / (1 + 2s_1) \\ r_2 &= (1 + 2N) / (1 + 2s_1) \\ r_1 r_2 &= 1 + 2N \end{aligned}$$

So we are left to maximize

$$1 / (1 + r_1)(1 + r_2)$$

We know that both r_1 and r_2 are positive; hence, we can simply seek to minimize the denominator.

$$\begin{aligned} (1+r_1)(1+r_2) &= 1 + r_1 + r_2 + r_1 r_2 \\ (1+r_1)(1+r_2) &= 2 + 2N + r_1 + r_2 \end{aligned}$$

The constant terms fall out, and we are left to minimize $r_1 + r_2$ subject to the constraint that $r_1 r_2$ is constant. The arithmetic-geometric mean inequality shows us then that this expression is minimized when $r_1 = r_2$.

In fact this result generalizes to more than two streets. Using the same logic as we used above, we can find that the equity for Y for a three street game is:

$$B_1 = \frac{8y r_1 r_2 r_3}{(1 + r_1)(1 + r_2)(1 + r_3)}$$

However if $B_1 \geq 1$, then what is happening is that Y does not have enough bluffs. Y should simply bet the growth of pot. And X should fold *all* hands.

By reducing out terms as before, we are left to minimize the expression $r_1 + r_2 + r_3$ subject to the constraint that $r_1 r_2 r_3$ is a constant. This occurs when $r_1=r_2=r_3$. This result additionally applies to any number of streets.

We call this betting progression (growing the pot by the same amount on each street) *geometric growth of pot*. This result is very significant to our understanding of no-limit bet sizing; in fact, as we will see in the next chapter, we use this as one of the bases of our entire bet-sizing scheme in no-limit holdem. When the game is static and one player is clairvoyant, the optimal bet size is to make the pot grow by the same amount on each street, such that on the last street the entire stack has been bet.

We will provide a short example game to contrast this strategy (against an optimal opponent) with another strategy that might seem reasonable but performs dramatically more poorly.

Example 19.3

Suppose we play the following game:

X and Y start with stacks of \$185 and each antes \$5.

Y draws a single card from a full deck.

There are three streets of no-limit betting.

If there is a showdown, then Y wins if his card is an ace or a king, and loses otherwise.

This game is equivalent to the clairvoyant game we discussed above, but with three streets. First, we can calculate the showdown value of this game; that is, suppose there were no betting. Then Y would win $\frac{2}{13}$ of the time for a total gross equity of \$1.54 per play and a net loss of \$3.46.

In order to make this game profitable, Y must make up \$3.46 in equity from the betting. We will solve this game using the solutions we found earlier for two bet sizing strategies:

Y_1 : Y bets one-third of his stack on each street.

Y_2 : Y bets the geometric growth of pot on each street.

For Y_1 , of course, each bet will be \$60. For Y_2 we can take advantage of the fact that $10r_1r_2r_3 = 10r_3 = \$370$. So r is the cube root of 37, or about 3.3322.

On each street, $p_n + 2s_n = rp_n$

Street 1:

$$\$10 + 2s_1 = \$10r$$

$$s_1 = \$11.66$$

Street 2:

$$\$33.32 + 2s_2 = \$33.32r$$

$$s_2 = \$38.86$$

Street 3:

$$\$111.04 + 2s_3 = \$111.04r$$

$$s_3 = \$129.48$$

X will play his optimal strategy of calling $p_n/(p_n + s_n)$ of the time on each street. Doing this results in his strategy values as follows:

Street	Y's bet (Y_1)	X's call freq.	Y's bet (Y_2)	X's call freq.
1	\$60	14.29%	\$11.66	46.2%
2	\$60	68.42% -> 9.77%	\$38.86	46.2% -> 21.3%
3	\$60	80.6% -> 7.88%	\$129.48	46.2% -> 9.84%
Total	\$60	7.88%	\$90	4.06%

We know that Y breaks even on his bluffs - on each street he bluffs appropriately to the number of hands he will carry over to the next street.

On his value hands, then, Y obtains the following value:

$$\langle Y_1 \rangle = (14.29\%)(\$60) + (9.77\%)(\$60) + (7.88\%)(\$60) = \$19.16$$

$$\langle Y_2 \rangle = (46.17\%)(\$8.34) + (21.31\%)(\$22.26) + (9.84\%)(\$59.40) = \$26.41$$

So Y does better by more than \$7 per winning hand by betting the geometric growth. On a per-hand basis, this is enough to swing the game from negative (using Y_1) to positive (using Y_2). It is worth noting that the limit of ex-showdown value (no matter how much betting is left) must be equal to X's showdown value; if Y could extract more than that value from him, X would simply fold and surrender his \$5 ante.

We also briefly look at a graph of the r constant against the stack sizes for two and three streets:

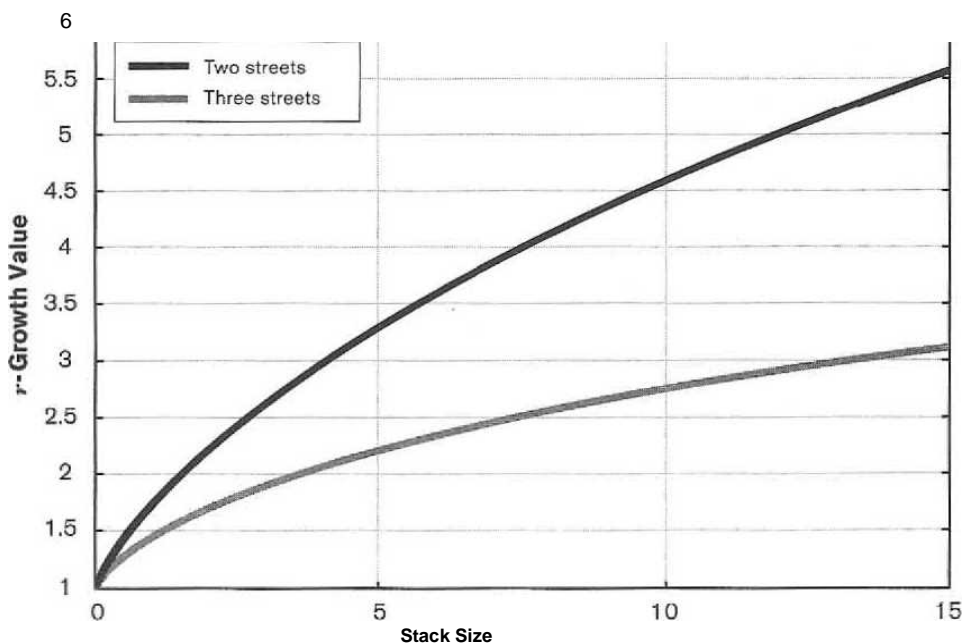


Figure 19.1. Geometric Growth of pot for two and three streets

When r is equal to 3, then s is equal to 1 and the bets are pot-sized. You can see that in order for pot-sized bets to be optimal in the three-street case, the stacks must bet 13 times the initial pot. At stacks of 3.5 times the pot, half-pot bets are sufficient to get all the money in geometrically.

Auction Games (Games with Infinite Streets)

The last static games we will consider are not precisely like poker games; in fact, the structure of their solutions is rather different from the structure of some of the preceding games. However, they do provide us with at least one important principle and have some interesting properties of their own. These are games with infinite streets (or arbitrarily large numbers of streets). We call

these *auction games* as well because in the limit as the number of streets goes to infinity, they converge onto a different type of game altogether.

Example 19.4

We begin by considering a clairvoyance game of an arbitrary number of streets.

M full streets.

Static hand values.

Initial pot size of 1 unit.

No-limit betting (stacks of S units).

Y is clairvoyant.

Extending the geometric growth of pot results to M streets, we know the following will be true:

The final pot size (if all bets are called) will be $1 + 2S$ (that is, the current pot size plus both stacks).

$$P = 1 + 2S$$

The growth on each street will be the Mth root of P (this is the geometric growth of pot for M streets), so

$$G = P^{(1/M)}$$

The bet size s (in current pots) on each street will be:

$$s = (G - 1)/2$$

Suppose that we call the pot on the current street 1. Y will bet the correct number of hands for value yv and bluff with $s/(1 + s)$ times as many hands. His total betting hands, then, will be $y_v(I + 2s)/(I + s)$.

We can state this as follows. The ratio of the hands bet on street $m+1$ to the hands bet on street m will be:

$$r_b = (1 + 2s)/(1 + s)$$

You can think of this value as the "bluff multiplier." Suppose we have a two-street game with 1 unit stacks and a pot of 1. Growing the pot geometrically would mean growing the pot from 1 to 3 in two streets. So r for this game is $\sqrt{3}$ and the bet size for each street is $(\sqrt{3} - 1)/2$, or about 0.366. To find the bluffs here, we simply figure that Y has y value bets. So on the river, he will bluff $s/(1 + s)$ times y . The expression $s/(1 + s)$ for this bet size is about 0.268. Hence, Y will bet $1.268y$ hands on the river, and $(1.268)^2 y$ hands on the turn. Substituting 0.366 for s in the above expression for r_b , we find that r_b is 1.268.

Suppose Y has a total number of nut hands for the entire game y . Let y_m be the number of hands he bets on street $m + 1$, while y_M is the total number of value hands, or $y_M = y$. Also let y_0 be the number of hands he bets on the first street.

Then we have:

$$y_m = y_0 r_b^m$$

$$y_M = y_0 r b^M = y$$

$$y_m = y r b^{m-M}$$

$$y_0 = y r b^{-M}$$

We may, however, run into a problem here. Suppose that $y_0 > 1$. Then we have a situation we've discussed earlier where Y doesn't have enough bluffs, since adding enough bluffs would make him bet more hands than he has. Hence, he can simply bet and claim the pot and X should optimally fold.

Otherwise, X calls with his typical frequency of $1/(1 + s)$. Let x_m be the number of hands he calls with on the $m + 1^{\text{th}}$ street. Then:

$$r_c = 1/(1 + s)$$

$$x_m = r_c^m$$

In either of these cases, assuming that Y plays optimally, X is indifferent to calling or folding. So the value of the game is equal to $y_0 - y$, which is the pot for every hand Y bluffs initially.

Next we consider what seems to be an entirely different type of game:

Example 19.5 - The Auction Game

Static hand values.

Initial pot size of 1 unit.

Stacks of S units.

Each player makes a bid secretly, and the bids are simultaneously revealed.

The bids are s_x and s_y and must be on the interval $[0, S]$. Both bids are put into the pot.

If $s_x > s_y$, then X wins, if $s_y > s_x$, then Y wins.

If $s_y = s_x$, then there is a showdown. (This will usually occur when the bids are both S).

This game is not quite so poker-like as some other games we have examined. However, let's be a little abstract. Suppose we have a poker game with static hand values and M streets. Let's suppose we are going to play a strategy that is inspired by the auction game.

For each hand we can hold, we are going to pick a number s_x that is some fraction of our total stack S.

On each street, we are going to bet an equal piece of our remaining stack $\epsilon = S/M$ until we reach our threshold value s_x , after which point we will just check and give up. ϵ is the Greek letter epsilon, which we introduced earlier and represents a very small amount. As we shall see, in this game, epsilon will go to zero as the number of streets grows. If our opponent puts in a raise of any amount, we will call if the raise is to an amount smaller than our threshold value and begin the next street by betting ϵ for our remaining stack and the remaining streets. If the raise is to an amount higher than our threshold, we will fold.

Now suppose our opponent is the nemesis. How much can he gain from exploiting our strategy? In order to find out where our threshold value is, he must call the initial bets that led up to our check. Thus, he cannot effectively exploit us by calling along and bluffing us out of the pot. In fact, the only gain he can realize is when our strategy forces us to bet an amount that is slightly larger than s_x . Suppose that we have 100 streets, and s_x is 0.601. For the first sixty streets, we bet 0.01 of the initial pot. On the sixty-first street, however, we still have not reached our threshold.

Now we cannot bet a smaller amount without exposing ourselves to exploitation. So we again bet 0.01, making our total bet 0.61. This is 0.009 more than we were actually willing to risk, and against almost all of the hands that make it to this point, we lose this additional 0.009 of value.

We can assess the value of a raise to the nemesis as well. Suppose that we bet on a particular street, making the total amount we have bet equal to x . Since we bet this street, the nemesis knows that our s_x value must be less than $x - \epsilon$. Suppose he then raises to y units. There are three cases:

$s_x < y$. In this case, we fold. However, the nemesis could have made more money ($s_x - x$ at least) by simply calling along until we gave up.

$s_x > y$. In this case, we call. After the raise, however, our bet size for each street has been reduced - instead of betting $(S - x)/M_{\text{cur}}$ per street the rest of the way, we will bet

$(S - y)/M_{\text{cur}}$, a smaller amount. Since we already showed how the nemesis can gain up to ϵ against our strategy, he certainly would not want to make it smaller.

$s_x = y$. In this case, we simply call and then check and give up on later streets. The nemesis gains nothing.

In no case can the nemesis gain anything by raising - in fact, he ought to simply call down as well.

Another way of seeing this is to consider two players who play this strategy against each other. The following is a result from game theory: if we have a strategy pair (X, Y) , and we know the value of the nemesis of X and the nemesis of Y , the game value must be between those two values. This is because the optimal strategy for X must do worse than the nemesis against Y , and vice versa. Hence, the value to the nemesis of our employing this betting structure instead of the optimal structure is at most the value of the game played optimally $G - \epsilon$.

We know that as the number of streets M approaches infinity, ϵ will go to zero, because M is the denominator of the expression for $\epsilon = S/M$. So we know that the value of any game converges on the value of the auction game that corresponds to the parameters of that game. Likewise, we know that for any number of streets, playing an auction strategy gives up at most ϵ , or S/M .

Example 19.6

Next, we consider the clairvoyant auction game.

Y is clairvoyant.

Auction game rules with stacks of S and initial pot 1.

Let Y have y winning hands. Then there are three types of hands total: X 's hands, which are all alike and Y 's hands, which are winning hands and losing hands. Y will evidently always choose $s_y = S$ when he has a winning hand; that is, Y will always be willing to bet his entire stack when he is going to win the pot. Let y be the amount of winning hands that Y has. Let $x(s)$ be the fraction of hands with which X bids at least s , and $y(s)$ be the fraction of hands with which Y bids at least s .

Y 's strategy will be to bid S with nut hands, and to choose s_y for bluffing hands such that he gradually gives up on bluffing hands until he is left with $y(S) = y$. To make Y indifferent between bidding two amounts s and $(s + \Delta s)$ with a dead hand, we have the following. Δs , or

"delta s" is generally used to denote changes (small or large) in s. So here Y is indifferent between bidding s and bidding an amount greater by Δs.

$$-\Delta x = x(s + \Delta s) - x(s)$$

(We use this notation because as s increases, x will decrease).

$$-x\Delta s + \Delta x\Delta s + (1 + 2s)(-\Delta x) = 0$$

$$(1 + 2s + \Delta s) - \Delta x = x\Delta s$$

$$-\Delta x / x = \Delta s / (1 + 2s - \Delta s)$$

as $\Delta s \rightarrow 0$, we have:

$$-\Delta x / x = \Delta s / (1 + 2s - \Delta s)$$

$$-\Delta x / x \approx \Delta s / (1 + 2s)$$

$$-dx/x = ds/(1 + 2s)$$

Integrating both sides, we have:

$$-\ln(x) + C = \ln(1 + 2s)/2$$

$$\ln(Cx^2) = \ln(1 + 2s)$$

$$x = k / \sqrt{(1 + 2s)}$$

Of course, $x(0) = 1$ as all hands must bid at least zero. Hence $k = 1$, and we have:

$$x(s) = 1 / \sqrt{(1 + 2s)}$$

In the same way, Y must make X indifferent between s and s + Δs:

$$(-\Delta y) (1 + 2s + 2\Delta s) + y(\Delta s) = 0$$

$$-\Delta y/y = \Delta s / (1 + 2s + 2\Delta s)$$

As $\Delta s \rightarrow 0$:

$$dy/y = ds/(1 + 2s)$$

$$y(s) = k / \sqrt{(1 + 2s)}$$

We know that $y(0) = y_0$, so:

$$y = k / \sqrt{(1 + 2s)}$$

$$\text{So } y(s) = y_0 \sqrt{\frac{1+2s}{1+2s_0}} \text{ and } x(s) = 1/\sqrt{(1 + 2s)}$$

However in this case, too, y_0 may be greater than 1, as we saw when looking at the $x(s)$ values

So the solution to the game is:

Solution:

$$y(s) = y_0 \sqrt{\frac{1+2s}{1+2s_0}} \quad (\text{or } 1, \text{ whichever is smaller})$$

$$x(s) = 1/\sqrt{(1 + 2s)} \quad (\text{or } 0 \text{ if } y_0(\sqrt{(1 + 2s)}) > 1)$$

To recap what these solutions mean, these are the total fraction of hands with which X and Y will bid at least s in the auction (if their entire stack is S). Suppose that the stacks are 10 total and Y has a winner 50% of the time. Then Y will bid at least 1 unit with:

$$y(1) = (0.5)(\sqrt{(1 + 2(10))/(1 + 2(1))})$$

$$y(1) = (0.5)(\sqrt{7})$$

$$y(1) \approx 1.323$$

Since this value is greater than one, Y will bid at least 1 with all hands. He will bid at least 5 with:

$$y(5) = (0.5)(\sqrt{(1 + 2(10))/(1 + 2(5))})$$

$$y(5) = (0.5)(\sqrt{21/11})$$

$$y(5) \approx 0.691$$

Here Y will bid at least 5 with 0.5 winning hands and 0.191 of his hands as "bluffs," and so on.

We can also show that these solutions are the limit as $M \rightarrow \infty$ of the clairvoyant poker game with M streets. We will ignore here the cases where Y lacks enough bluffs, and concentrate on the more usual case where both sides have mixed strategies on the first street.

To avoid confusion with regard to the values of s , let

$$s_n = (G - 1)/2$$

which is the one-street betting ratio. (See Example 19.4)

We can use the following approximations:

$$s_n = (G - 1)/2 \approx \ln(G)/2 = \ln P / 2M$$

We know from the solution to the clairvoyant poker game (Example 19.1) that $x_m = r_c^m$.

$$x_m = r_c^m$$

$$x_m = (1 + s_1)^{-m}$$

$$x_m \approx e^{-ms_1}$$

$$x_m \approx \exp\left(-\frac{m}{2M} \ln P\right)$$

$$x_m \approx P^{-\frac{m}{2M}}$$

We also know that the pot size after this street will be:

$$1 + 2s_1 = G_m$$

$$1 + 2s_1 = P^{\frac{m}{2M}}$$

Substituting, we have:

$$x_m = (1 + 2s_1)^{-1/2}$$

$$x_m = x(s)$$

We can show that the same holds true for $y(s)$ - that the solution for the auction game in the limit as M goes to infinity is the same as the clairvoyant no-limit game with M streets.

Key Concepts

- The clairvoyant game can be extended to more than one street. The key feature of the multi-street game is that the clairvoyant player carries forward a set of hands from street to street, bluffing many more hands on earlier streets and giving up on a fraction of those hands on each street.
- In the no-limit multi-street clairvoyant game, the optimal bet size is the geometric growth of pot - growing the pot by the same fraction on each street.
- As the number of streets increases, positional value goes to zero - the auction game is symmetric.
- As the auction game (and geometric growth of pot) show, it is optimal to bet small fractions of the pot in static games. Non-static games do not necessarily follow this principle.
- Auction strategies assign a value s_x , to each hand. This is a reasonable way to think about no-limit games in general, as having a distribution of hands that each has a value. Then the goal is to correlate the pot size with our hand strength.
- The non-clairvoyant player plays each street independently as if it were a single street game,
- The hands that Y bluffs once and then gives up on have the a bluffing ratio relationship to the total hands that Y will bet for value and bluff on the second street.

Chapter 20

Drawing Out; Non-Static Multi-Street Games

In the previous chapter, we discussed some clairvoyant multi-street games where the hand values do not change from street to street. We saw that in that case the non-clairvoyant player checked and called on each street as though it were a single street game. The clairvoyant player employed a strategy of "stacking" bluffs, and giving up on some of them on successive streets, such that on the river he had an appropriate mixture of hands for the single-street river strategy.

In this chapter, we continue to add more complexity by adding draws to the games. What this really means is that one or the other player may have a hand which has some probability of being best after the next street. In some cases both players will know what cards have come (we call this an *open draw*). This is more like a situation in holdem, where the river card is faceup and shared. We also have *closed draws*, which are similar to seven-card stud, where the river cards are revealed to only one player.

We begin by considering once again a clairvoyant game. This game actually contains no draws; instead, it might be the end position of a hand where there were draws,

Example 20.1 - The Clairvoyance Game (Open Draws)

Consider the following game:

It's limit holdem and the pot contains 6 big bets.

X has an exposed pair of aces.

Y has a hand from a distribution that contains $\frac{7}{8}$ hands with two hearts and $\frac{1}{8}$ dead hands.

The board has come:

Q♥ 7♦ K♥ 4c 8♥

Using the solutions from the original clairvoyance game (Example 11.1), we can see that Y will value bet all of his flushes, and then bluff enough dead hands so that X is indifferent to calling. Using the familiar ratios we calculated when studying that game, we have the following strategies:

X checks and Y bets all of his flushes ($\frac{7}{8}$) and a times as many dead hands, or $\frac{1}{8}$ of his hands total. X folds α of his hands ($\frac{1}{7}$) and calls with the remaining $\frac{6}{7}$.

Example 20.2

But now consider this related game:

It's the turn in limit holdem and the pot contains 4 big bets.

X has an exposed pair of aces.

Y has a hand from a distribution that contains $\frac{1}{10}$ hands with two hearts and $\frac{9}{10}$ dead hands.

For simplicity's sake, suppose the flush draw comes in exactly 20% of the time. We'll also ignore the fact that the flush comes in less frequently when Y actually has the Bush draw because there are fewer hearts left in the deck.

The board has come:

Q♥ 7♦ K♥ 4♣

How will the action on the turn go? In Part II. we looked at games like this. So it should be clear that the aces will bet. Since Y has no hands that are ahead of the aces, he will simply call when he has pot odds (with his flush draws). He will also call with some dead hands, hoping to bluff if the flush comes in. The total number of dead hands he calls with is equal to the amount of hands he will bluff on the river if the flush comes in. We can find the number of bluffs by figuring out what the pot size will be. After a bet and a call on the turn, the pot contains six big bets. Y will want to bluff on the river (if the flush card comes) with $\alpha = 1 / (P + 1) = 1/7$ as many hands as he will value bet, or $1/70$ of his hands total.

Now the river comes the 8♥. We have now arrived at the situation in Example 20.1 above: the pot contains six bets, Y's distribution contains a mixture of flushes and dead hands, and the board is identical. Y now cheerfully bets all his flushes and bluffs with $1/7$ as many dead hands - which is all the dead hands he brought to the river. Before going on, you should try to figure out how often X has to call here. The single-street solution is of course that X should call $6/7$ of the time. However, the fact that we paused and turned it into a quiz should alert you that this is probably not the right answer. Suppose X decided not to call at all. Y did, after all, bet an amount on the river that would make X indifferent to calling or folding.

We know, however, that if X did decide to simply fold all the time, Y could likely exploit that. Since Y is left with only the set of hands that he wants to have to play optimally on the river, he would have to change his play on a previous street to carry forward more bluffs. But now he has a problem. The problem is that in order to carry forward bluffs to the river, Y has to call a bet on the turn with a dead hand in order to bluff on the end. However, $4/5$ of the time, he will be unable to bluff because the flush doesn't actually make it and he has no credible threat to have a strong hand. On those hands, he just loses one bet. It turns out that he can in fact do this if X never calls on the river when the flush comes in. He loses 1 bet $4/5$ of the time, but gains 5 bets when the flush comes in, for a total win of $1/5$ of a bet.

So X must find a different strategy. What he is trying to do here is not to make the flush draws indifferent on the river to bluffing with a dead hand, but to make the dead hands indifferent to the entire play of calling the turn (with nothing), hitting the flush, and bluffing the river. We can find this value algebraically. Let x be X's calling percentage. If the dead hands fold immediately on the turn, their expectation is 0. To make them indifferent to the entire play, the expectation of that play must also be zero. The dead hands lose 2 bets when the flush hits and they are called, win 5 bets (4 in the pot, plus the bet from the opponent on the turn) when they are not called, and lose 1 bet when the flush misses.

$$\langle \text{dead hand, fold} \rangle = 0$$

$$\langle \text{dead hand, play} \rangle = p(\text{call})p(\text{flush})(-2) + p(\text{flush})p(\text{fold})(5) + p(\text{no flush})(-1)$$

$$\langle \text{dead hand, play} \rangle = x(1/5)(-2) + (1/5)(1 - x)(5) + 4/5(-1)$$

$$\langle \text{dead hand, play} \rangle = x(-7/5) + 1/5$$

$$x(-7/5) + 1/5 = 0$$

$$x = 1/7$$

This leads to two principles that are at the core of analyzing multi-street play.

The first is:

In order for a bluff to be credible, it must come from a distribution of hands that contains strong hands as well.

When considering calling ratios, the cost of playing the distribution to that point must be considered. Marked made hands can call less often on the river against marked draws when the draw comes in if the marked made hands bet and were called on previous streets. The made hands effectively extracted their value on an earlier street, and the opponent cannot exploit the made hand effectively because the draw doesn't come in often enough.

This principle has many implications for real poker play; the first is that it's frankly terrible to find oneself in a situation with a marked open draw. You can see from this previous example - X only called for value with $\frac{1}{7}$ of his hands on the river, instead of $\frac{6}{7}$, but Y gained nothing from compensating bluffs. X's strong position on the turn enables him to give up when the flush comes in because Y paid so much for the opportunity to bluff on the river.

In fact, you might notice that Y's partial clairvoyance doesn't help him very much at all here; this occurs primarily because his distribution of hands is so weak. Optimal strategies would normally never result in a situation where a hand was marked as a draw; in fact, our fellow poker theorist Paul R. Pudaite has proposed a "Fundamental Theorem of Chasing" that conjectures that given sufficient potential future action, every action sequence must have some positive chance of containing hands that are the nuts for all cards to come.

The second principle is also of prime importance in understanding how to apply game theory to poker.

Multi-street games are not single-street games chained together; the solution to the full game is often quite different than the solutions to individual streets.

In the first game, we considered the play on the river in isolation as a problem unto itself. However the solution to the two-street game is different, even though after the play on the turn, the second game reduced to an identical situation to the first. In fact, this is almost always the case. This is one of the most important differences in the analysis of games such as chess or backgammon and the analysis of poker. In those games, if we arrive at a particular position, the move order or the history of the game is irrelevant to the future play from this point. This is not the case in poker.

This critically important point also leads us to one of the great difficulties of doing truly optimal poker analysis - no situation exists in a vacuum. There was always action preceding the play on a particular street that should be taken into account. There are always more complex distributions of hands that must be taken into account. In fact, we really cannot label a play or solution "optimal" unless it is taken from the full solution to the game. This difficulty is in some ways insurmountable due to lack of computing power and also because of the multiplayer nature of the game. However, we still look for insight through the solving of toy games and then try to adjust for these factors in real play.

Example 20.3 - The Clairvoyance Game (Mixture of Draws)

Suppose that instead of a distribution of hands for Y that contains only draws and dead hands, he has a distribution which contains x% flush draws (the main draw), y% weak closed draws - for example underpairs that don't have sets yet, and the remainder z% are dead hands. For the sake of simplicity, presume that the flush hits 20% of the time and the weak draws hit 4% of the time (these events are independent - both could happen, or neither).

The weak draws don't have pot odds (even considering the extra bet they might win on the river) to call a bet on the flop. However, since Y is calling with additional hands in order to balance his bluffs on the river, he might as well (in fact, he must to be optimal) call with all of his weak closed draws first, and finally with some dead hands.

So X bets, and Y calls with a mixture of the three types of hands. However, Y still has some equity on the river when the flush doesn't come in, because 4% of the time he has hit a set with a non-flush card. He will bet his sets for value and bluff with some dead hands. (Busted flush draws, weak draws that didn't come in, and dead hands from the start are all dead hands on the river).

In the previous game (Example 20.2), it was clear which hands X had to make indifferent by calling on the river. X could not prevent Y from profitably calling and betting the river for value when he made his flush. Hence, he made Y indifferent to calling with a pure bluff. Here, however, he must decide between trying to make Y indifferent with his weak closed draws or with his completely dead hands. This choice depends on the total fraction of each of these hands that are present.

If $x = 10\%$ and $z = 90\%$, we have the game we previously solved.

Suppose that $x = 10\%$ and $y = 90\%$ (there are no dead hands). Then clearly X must make Y indifferent with his weak closed draws. He does this by calling enough in both cases that Y is indifferent to calling and bluffing the river. Say that x_f is X's calling frequency when the flush comes in, and x_n is X's calling frequency when another card hits.

Now let's consider a single of Y's closed draws. Y has a number of strategies he can employ. We'll transform multiple decisions into a matrix because it makes solving the game easier. We'll denote Y's strategies as the following:

[<turn action>/<river action if the flush comes in>/< river action if the flush misses >]

So for example [F//] is a turn fold, while [C/B/K] means to call on the turn, bluff the river if the flush comes in and check the river if the flush misses. In all cases, if the hand hits a set, it will value bet on the river. If this occurs, then us equity is 5 bets (the bets in the pot and from the turn) plus the appropriate implied odds (x_f calls if the flush came in, x_n calls if it missed). We'll assume that the set is made with a flush card $1/4$ of the time.

$$V_{\text{set}} = (3/4)(5+x_n) + (1/4)(5+x_f)$$

Y's action sequence	Y's equity	Y's equity assuming $x_n = 1$ and $x_f = 1$
[F//]	$-(1/25)(4) = -0.16$	
[C/B/B]	$(1/25)V_{\text{set}} + (19/100)(-2x_f+5(1-x_f)) + (77/100)(-2x_n+5(1-x_n))$	
[C/B/K]	$(1/25)V_{\text{set}} + (19/100)(-2x_f+5(1-x_f)) - 77/100$	
[C/K/B]	$(1/25)V_{\text{set}} - 19/100 + (77/100)(-2x_n+5(1-x_n))$	
[C/K/K]	$(1/25)V_{\text{set}} - 19/100 - 77/100$	$72/100$

The values $1/25$, $19/100$, and $77/100$ are the respective probabilities of hitting a set (whether a flush comes or not), hitting a flush card but not a set, and missing everything.

We can immediately see that calling and checking down whatever comes must be dominated. Suppose that x_n and x_f were both 1, the most favorable situation for a set. Then $\langle [C/K/K] \rangle = \frac{1}{25}(6) - \frac{96}{100} = -\frac{72}{100}$, which is substantially worse than simply folding immediately.

We also know that the total number of bluffing hands when the flush comes will be greater than the total number of bluffing hands when the flush does not come (because we will have more value hands). Therefore, there must be some hands that are played in the manner $[C/B/K]$ in addition to those which are played in the manner $[C/B/B]$, (If any hands are played $[G/K/B]$, we can simply remove these hands and convert $[C/B/K]$ to $[C/B/B]$).

Since we have both types of hands, bluffing on the river when the flush does not come in must be a mixed strategy.

We therefore have:

$$\begin{aligned} \langle [C/B/B] \rangle &= \langle [C/B/K] \rangle \\ \frac{77}{100}(-2x_n + 5(1 - x_n)) &= -\frac{77}{100} \\ 5 - 7x_n &= -1 \end{aligned}$$

We additionally know that Y must bluff at least sometimes on the river since he does not always have the nuts. Hence we have the following:

$$\begin{aligned} \frac{4}{100}V_{\text{set}} + \frac{19}{100}(-2x_f + 5(1 - x_f)) - \frac{77}{100} &= \frac{16}{100} \\ \frac{1}{25}((\frac{3}{4})(5 + \frac{6}{7}) + (\frac{1}{4})(5 + x_f)) + \frac{19}{100}(5 - 7x_f) &= \frac{61}{100} \\ x_f &= \frac{3}{7} \end{aligned}$$

This is a pretty interesting result as well: by adding draws that are only $\frac{1}{4}$ as strong as the flush draws, Y is able to force X to call on the river three times as often. Here we see the power of closed draws. When only closed draws are possible (the flush didn't come in), X must call $\frac{6}{7}$ of the time to keep Y indifferent. When the flush does come in however (an open draw), X needn't call so much, only $\frac{3}{7}$ of the time.

For Y's play, he need only play enough weak draws so that he has the appropriate amount of bluffing hands when the flush comes in. If he plays y weak closed draws, he will have a made hand $\frac{1}{9}y + \frac{1}{10}$ of the time on the river, and his a should be $\frac{1}{7}$. Hence,

$$\begin{aligned} (\frac{1}{7})(\frac{1}{9}y + \frac{1}{10}) &= \frac{8}{9}y \\ \frac{8}{9}y &= \frac{1}{63} + \frac{1}{70} \\ y &= \frac{9}{550} \end{aligned}$$

(This value is a little more than $\frac{1}{70}$, the total number of dead hands Y played in the previous game).

Now let's consider a more general case, where Y's distribution contains some of each type of hand. The shape of his distribution determines the strategy of the game - X must make a decision regarding which type of hand to make indifferent. Y needs to add "bluffing" hands to support his flush draw calls. Y will of course prefer to add hands with equity compared to hands with none. So he will add the weak closed draws to his calling distribution until they are exhausted. Only then will he add dead hands and attempt to bluff when the open draw comes in.

In a way, this is similar to the [0,1] game. Y's distribution has regions; a region of flush draws, a region of weak closed draws, and a region of dead hands. He has to take actions in certain frequencies for purposes of balance, to make X unable to exploit him. Likewise, however, he will choose hands that have positive value with which to do this. At some point, there will be a threshold hand that will have a mixed strategy - at that hand he will be indifferent.

Sometimes, that threshold hand will be a weak closed draw (as we saw when there were no dead hands). Sometimes, that threshold hand will be a dead hand (as we saw when there were no weak closed draws). The key is: if there are not enough weak closed draws in Y's distribution, then X cannot make him indifferent to calling and bluffing with those closed draws. Y will always have greater equity from doing that than doing something else. This occurs in just the same way that X cannot make him indifferent to calling with flush draws; X cannot bet enough to prevent him from making money. The presence of the flush draws allows the weak closed draws (up to a certain number) to profitably call on the turn because they obtain positive value (compared to checking) by bluffing when the flush comes in. This positive value enables them to overcome the fact that they are calling without sufficient pot odds to do so.

So suppose that $x = 30%$, $y = 1%$, and $z = 69%$. Then X will make Y indifferent to calling and bluffing with a dead hand. Y will simply call with all his weak closed draws. But if $x = 30%$, $y = 50%$, and $z = 20%$, then X makes Y indifferent to calling and bluffing with a weak closed draw.

This is a question that people who are familiar with our work on optimal play often ask us: "Which draws should we make indifferent to calling and bluffing?" The answer is, in a sense, the best one you can! It is possible to extend the above methodology- to include any combination of draws that do not contain the nuts. We can take this concept further by combining the clairvoyance game with the two-street game and looking at a situation where a known, exposed hand plays against a known distribution with cards to come.

Example 20.4

Consider the following situation:

X and Y play holdem on the turn.

X's hand is exposed.

Y holds some hand y chosen randomly from a distribution of hands Y .

There are P units in the pot.

Now we can directly map Y and y to probabilities; Y knows X's hand, and he knows all the possible cards that can come on the river, so his chance of being ahead in a showdown (with no other betting) is simply the probability that he will have the best hand after the river card (for the sake of this discussion, we'll ignore ties).

For example, suppose we have the following:

A's hand: $A\spadesuit 4\clubsuit$

The board: $A\heartsuit J\spadesuit 8\diamondsuit 7\clubsuit$

Y's distribution of hands: $\{99+, AQ+\}$

Y has the following types of hands:

Hand Type	Frequency	Probability of winning
AA	$\frac{1}{39}$	1
KK	$\frac{6}{39}$	$\frac{2}{44}$
QQ	$\frac{3}{39}$	$\frac{2}{44}$
JJ	$\frac{3}{39}$	1
TT	$\frac{6}{39}$	$\frac{6}{44}$
99	$\frac{6}{39}$	$\frac{6}{44}$
AK	$\frac{8}{39}$	$\frac{41}{44}$
AQ	$\frac{6}{39}$	$\frac{41}{44}$
Total	1	

We can sum the frequencies of each of the unique probability outcomes (without regard to the actual hand type) and create an alternate distribution \bar{Y} , which for the purposes of this game is equivalent to Y :

$$\{(1, \frac{4}{39}), (\frac{41}{44}, \frac{14}{39}), (\frac{6}{44}, \frac{12}{39}), (\frac{2}{44}, \frac{9}{39})\}$$

The reason that this is equivalent to the previous distribution is that all the hands with one probability of winning are interchangeable, even though they have different cards that "make" them, and even though both players will see the river card. This is because on the river, the second player has perfect information about the hand. Therefore, he will never pay off a bet incorrectly, so the first player cannot gain from inferences about the river card by betting or raising.

There are, of course, two streets remaining, but play on the river is fairly simple. If the hand gets to the river, the player with the made hand checks (if he is first). The player with the distribution of hands then value bets all his winning hands, plus a fraction of losing hands equal to his winning hands times $1/(P + 1)$. This is the familiar betting-to-bluffing ratio that we saw first in the one-street clairvoyance game (Example 11.1). Then the opponent calls $P/(P+ 1)$ of the time with his exposed hand, and there is a showdown.

Y therefore breaks even on his bluffs, and wins one bet when he is called with a winner. The bets he wins on the river are "implied odds" from the play on the turn, but because the play on the river is essentially automatic if both players play optimally, we can actually incorporate that directly into calculations about play on the turn. We define the value of a winning concealed hand to be the pot size on the river plus $P/(P + 1)$ bets, the amount that the concealed hand wins by value betting and bluffing appropriately on the river. We use the notation P' to represent this value.

$$P' = P + P/(P+1)$$

X has to pay off here with $P/(P+ 1)$ of his hands because we constrained the game to contain only a half street on the turn.

Let's assume that in the above example after the turn betting there are five bets in the pot. Now X will bet his winning hands, and be called $\frac{5}{6}$ of the time. So the value of the hand is the pot (5 bets) plus $\frac{5}{6}$ of a bet = $\frac{35}{6}$.

Now using our mapped distribution, we can define some more things. Suppose we order the probabilities from best to worst. So in the example distribution, the probabilities are:

$4/39$ of the hands have equity 1.
 $14/39$ of the hands have equity $41/44$
 $12/39$ of the hands have equity $6/44$
 $9/39$ of the hands are have equity $2/44$

We now create a function that identifies the probability of a particular hand based on the rank of that hand:

$y(t)$ = the winning probability of the best hand in \bar{Y} , after removing the t best hands from \bar{Y} .
 $(0 \leq t \leq 1)$

$y(0)$ = the winning probability of the best hand in \bar{Y} .
 $y(1)$ = the winning probability of the worst hand in \bar{Y} .

In our example distribution, we have:

$y(0) = 1$
 $y(1/4) = 41/44$
 $y(1/2) = 6/44$ (the median hand)
 $y(1) = 2/44$

Another function that is quite useful for analyzing this game is the average equity of the best t hands, which we denote by $y^\theta(t)$. This is a measure of the probability of a range of hands; it is the expected value of the designated subset of the distribution. So for example, $y^\theta(1/2)$ is the average of $y(t)$ for the top 50% of hands.

$y^\theta(t)$ (the value of the t^{th} hand) is often discontinuous and ill-behaved.

We can see some values of the $y^\theta(t)$ function for our example distribution:

$y^\theta(0) = 1$
 $y^\theta(1/4) = (16/156)(1) + (23/156)(41/44) \approx 0.23995 / 0.25 \approx 0.95980$
 $y^\theta(1/2) = (8/78)(1) + (28/78)(41/44) + (3/78)(6/44) \approx 0.44231 / .5 \approx 0.88462$
 $y^\theta(1) \approx 0.48951$

$y^\theta(t)$ is a much different and much more well-behaved function; it turns out that this function has some nice properties.

- Continuous.
- Decreasing for all t , and strictly decreasing as long as $y(t) > y(0)$.

Another way of expressing $y^\theta(t)$ (which shows its continuity) is as follows:

$$y^\theta(t) = \left(\int_0^t y(t) dt \right) / t$$

$y^\theta(1)$ is an important value - the average equity for all the hands is \bar{Y} . This value will be quite significant in determining the play. For example, suppose that $y^\theta(1) = 1$. Then Y would bet the turn and X would fold (because X has no equity in the pot!). Similarly, if $y^\theta(1) = 0$, then Y could not win and X would simply claim the pot. However, between these two extremes, there are cutoffs such that the players employ mixed strategies designed to prevent exploitation.

Example 20.5 - The One-and-a-Half-Street Clairvoyance Game

One half-street on the turn (X checks dark), followed by a full street on the river.

Y's hand is exposed.

Y's hand is randomly chosen from a distribution Y

There are P bets in the pot.

No raising is allowed*

* There is one extremely limited case where Y might bet and X check-raise if it were allowed, when the relationship between the distribution of Y, the adjusted pot on the turn, and the adjusted pot on the river meets a quite narrow set of criteria. We ignore this case for simplicity's sake here by simply removing the raise option.

To solve this game, we must identify the optimal strategies. Consider intuitively how the play should go. If Y's distribution is very strong (such that Y doesn't have "enough" equity), then will just bet and take the pot, as X lacks pot odds to call. On the other hand, if Y's distribution contains a mixture of strong and weak hands, Y will bet some set of his better hands and check the rest. And if Y's distribution contains only weak hands, then he will just check and play the river.

First, let's find some of the threshold values. We'll calculate all equities from Y's point of view as usual. Assume Y bets all hands. If he is called, the pot on the river will be

$P+2$ bets and his hand will be worth $(P+2)^t$ on the river if he hits a whining card. It costs him one bet to make this play. The betting on the river is already included in the t . If Y bets and X folds, then Y wins P units.

X should fold immediately if the value of folding is less than the value of calling.

$$y^0(1)(P+2)^t - 1 \geq P$$

$$y^0(1) \geq (P+1)/(P+2)^t$$

or

$$y^0(1) \geq (P+1)/(P+3)/(P+2)(P+4)$$

So if Y's overall hand strength (chance of winning) from all the hands in his distribution is greater than the value $(P+1)/(P+2)^t$, then Y can bet and claim the pot. X can't profitably call on the turn at all.

We'll call that Case 1.

Case 1.

Condition: $y^0(1) \geq (P+1)/(P+2)^t$

Action: Y bets and X folds.

Something else comes up here. Suppose Y's distribution is such that this case is satisfied, but checks and then bets the river? Can X sometimes call then?

The answer is no. The reason is that by doing so, X can be exploited. Suppose that Y has $1/5$ hands that always win, and the remaining $2/5$ of hands are close to always winning, such that Case 1's condition is true. If Y can occasionally check the nuts on the turn, bet the river and occasionally get called, then he can employ the following exploitive strategy:

Bet all hands that don't have 100% equity (X must fold) and win P units.

Check the nuts and bet the river.
 Hope to get called and win more than P units.

Since X's optimal strategy cannot allow this exploitation. X's optimal strategy is to fold at any future point during the hand. This is actually a general rule:

If your opponent's distribution and yours are such that your opponent can take some action and you must fold and surrender the entire pot, then at any later point in the hand, it is at least co-optimal to fold to any future action.

Avoiding this situation is another reason to play properly balance strategies throughout the hand.

In any event, at the opposite end of the spectrum, let us consider the situation where X should always call a bet. If X will always call, then Y's equity from checking a hand is simply $y(P^t)$. That is, Y checks the turn and if he wins on the river, the value of the hand is P^t . Y's equity from betting a hand y is similarly $y(P+2)^t - 1$. Then Y should bet when:

$$y(P+2)^t - 1 \geq y(P^t)$$

$$y \geq 1/((P+2)^t - P^t)$$

What this means is that if X will call all the time, then Y can profitably bet all hands such that y is greater than this value. For most P values, this expression is a little less than $1/2$. At $P=3$, for example, it evaluates to around 0.48. It should be clear that Y profits by betting when he has more than 50% equity in the pot. However, from this expression, we can see that Y can profitably bet with a little less than 50% equity in the pot, because of the implied odds his clairvoyance affords him on the next street.

Now we can make use of one more piece of notation. We will often want to designate "the set of hands where y is greater than some value y_0 ," in terms of how many hands there are or what fraction of the hands satisfy this condition. So the set of hands $\{y \mid y \geq y_0\}$ corresponds to a different set $\{y(t) \mid t \leq \mathcal{T}\}$ for some value \mathcal{T} . We'll call this value $\mathcal{T}[y_0]$. Let's go back to our example distribution \mathbf{Y}' :

$$\{(1, 4/39), (41/44, 14/39), (6/44, 12/39), (2/44, 9/39)\}$$

Suppose we want all the hands where x is greater than $y_0 = 1/2$. This set of hands includes all the 100% winners and all the $41/44$ winners $\{AA, JJ, AK, AQ\}$. We can also define this set of hands with a t value that corresponds to the cutoff value, t values below this threshold will have equity greater than or equal to y_0 , and t values above this threshold will have less. Since these hands make up $18/39$ of the total, $\mathcal{T}[1/2] = 18/39$. This is the t value such that $y(t) \geq 1/2$ for all t values less than $\mathcal{T}[1/2]$, or more straightforwardly, it's the fraction of hands with equity greater than or equal to $1/2$.

$y(1/2)$ is the equity of the best hand left after removing the best $1/2$ of hands. So there are 39 hands total. Removing the best 19 of them, the best hand left is $6/44$. So that's what $y(t)$ is, the best hand after removing the best t hands. $\mathcal{T}[x]$ is like the inverse of this. So suppose we have an *equity* $1/2$, and we want to know how many hands have at least that much equity. In this case it's $18/39$. $\mathcal{T}[x]$ is the fraction of hands with at least x equity.

So in this case, we'll define:

$$t_0 = \mathcal{T}[1/((P+2)^t - P^t)]$$

t_0 is the threshold where if X calls all the time. Y should bet all hands that have greater equity. If Y does bet all these hands, then X can call all the time as long as Y's equity with the hands he bets is less than $(P+1)/(P+2)^t$.

To summarize this case, if X calls all the time, then Y can come up with a set of hands to bet with a threshold of t_0 . His equity with that set of hands is $y^\theta(t_0)$. As long as this value is low enough that X always has odds to call on the turn, X just calls all the time.

Case 2.

Condition: $y^\theta(t_0) \leq (P+1)/(P+2)^t$

$t_0 = \mathcal{T}[(1/(P+2)^t - P^t)]$

Action: Y bets and B calls.

So those are the two extreme cases - where Y has almost all of the equity, and where Y doesn't have enough strong hands to prevent X from calling no matter what strategy he employs.

The third case, however, is the most interesting one, because X utilizes a mixed calling strategy. That is, the set of hands that Y bets is too strong for X to call all the time, but not strong enough for X to simply fold. Instead, X calls often enough that Y can't just bluff his bad hands with impunity, but folds often enough that Y can't exploit him by betting strong hands exclusively.

So X is going to have some "folding frequency," which is the % of the time he will fold on the turn. We will call this value α . Here, however, since it is the turn, applying the typical $\alpha = P/(P+1)$ is incorrect. Remember, X's hands are all alike, so his play can be expressed as a single number (call/fold frequency on the turn). On the river, of course, he simply calls $P/(P+1)$ of the time. Y's strategy, on the other hand, includes betting different types of hands with different values.

If X will fold α of the time and call $1 - \alpha$ of the time, we can find Y's strategy in terms of α . Y will bet when the equity of betting is higher than the equity of checking.

$$\langle Y, \text{bet hand } x \rangle = \alpha P + (1 - \alpha) (y (P + 2)^t - 1)$$

$$\langle Y, \text{check hand } x \rangle = y P^t$$

$$\alpha P + (1 - \alpha) (y (P + 2)^t - 1) \geq y P^t$$

$$y [(1 - \alpha)(P + 2)^t - P^t] \geq 1 - \alpha(P + 1)$$

The left side of this equation is positive as long as $\alpha P < 2$, or $\alpha < 2/P$.

The right side of this equation is positive as long as $\alpha > 1/(P+1)$. So if X chooses his folding frequency a greater than $1/(P+1)$, then Y just bets all the time. This could be optimal (in Case 1 from above), but if that case doesn't apply, then Y is exploiting X. We can also see that if X chooses a frequency α that's high enough, then Y bets whenever his equity from betting is higher than checking:

$$y [(1 - \alpha)(P + 2)^t - P^t] \geq 1 - \alpha(P + 1)$$

$$y \geq (1 - \alpha(P + 1))/[(1 - \alpha)(P + 2)^t - P^t] \tag{20.1}$$

We can use another fact about optimal strategies. If X is playing a mixed strategy of calling and folding with his hand, then he must be indifferent to doing those two things. Otherwise, he could simply improve by moving his calling frequency. So suppose that Y is betting his optimal set of hands \mathcal{T} . Then we have!

$$y^0(\mathcal{T})(P+2)^t - 1 = P$$

$$y^0(\mathcal{T}) = (P+1)/(P+2)^t$$

Now we return to the properties of $y^0(t)$. This function is decreasing, so we know that $y^0(1) \leq y^0(t) \leq y^0(0)$. We also have an important value $(P+1)/(P+2)^t$. This value is the value of $y^0(t)$ such that if Y bets all hands better than T, X is indifferent to calling or folding. So one of the following must be true.

Possibility 1: $y^0(1) \geq (P+1)/(P+2)^t$

Possibility 2: $y^0(0) \leq (P+1)/(P+2)^t$

Possibility 3: $y^0(\mathcal{T}) = (P+1)/(P+2)^t$

Possibility 1 is within Case 1, where Y bets all the time and X just folds.

Possibility 2 is within Case 2, where Y only bets his hands that have positive equity, and X calls all the time.

Possibility 3 is a new case, Case 3. We know that $y(t) \leq y^0(t)$, because $y(t)$ is just the worst hand that makes up the equities in $y^0(t)$. If $y(\mathcal{T}) > 1/((P+2)^t - P^t)$: then the threshold hand (the hand that makes X indifferent to calling) is still a "good hand." So Y bets and X still calls because he still has odds, and this is actually Case 2.

If $y(\mathcal{T}) < 1/((P+2)^t - P^t)$, however, then we have Equation 20.1, which we found earlier, and shows that Y will only bet hands that have positive equity from betting:

$$y \geq (1 - \alpha(P+1)) / [(1 - \alpha)(P+2)^t - P^t]$$

Substituting $y(\mathcal{T})$ for y and solving for α , we get:

$$\alpha = (1 - y(\mathcal{T})[(P+2)^t - P^t]) / (P+1 - y(\mathcal{T})(P+2)^t)$$

In this case, Y chooses a \mathcal{T} such that X is indifferent to calling or folding against the whole range of hands better than T. And X calls and folds with a frequency that makes Y indifferent to betting with exactly $y(\mathcal{T})$.

Case 3.

Condition: Neither Case 1 nor Case 2 hold.

Y bets hands such that $y^0(\mathcal{T}) = (P+1)/(P+2)^t$

X folds $\alpha = (1 - y(\mathcal{T})[(P+2)^t - P^t]) / (P+1 - y(\mathcal{T})(P+2)^t)$ of his hands and calls with the rest.

Now let's look at some examples. First, we return to our example hand. Recall that the situation was:

X's hand: A♠ 4♣

The board: A♥J♠ 8♦ 7♣

Y's distribution of hands: {99+, AQ+}

Y has the following types of hands:

Hand Type	Frequency	Probability of winning
AA	$\frac{1}{39}$	1
KK	$\frac{6}{39}$	$\frac{2}{44}$
QQ	$\frac{3}{39}$	$\frac{2}{44}$
JJ	$\frac{3}{39}$	1
TT	$\frac{6}{39}$	$\frac{6}{44}$
99	$\frac{6}{39}$	$\frac{6}{44}$
AK	$\frac{8}{39}$	$\frac{41}{44}$
AQ	$\frac{6}{39}$	$\frac{41}{44}$
Total	1	

This converted into a secondary Y distribution of:

$$\{(1, \frac{4}{39}), (\frac{41}{44}, \frac{14}{39}), (\frac{6}{44}, \frac{12}{39}), (\frac{2}{44}, \frac{9}{39})\}$$

Say the pot is 5 bets.

First of all, we can look for which case we are in.

Case 1:

$y^0(1)$ is about 0.48 (as we found earlier)

$$(P+1)/(P+2)^t = (6)/(63/8) = 48/63, \text{ or about } 3/4$$

So Case 1 doesn't hold because $y^0(1) \geq (P+1)/(P+2)^t$.

This should make sense, because Y's distribution here isn't strong enough that he can just blast away with all his hands and X has to fold. X has a big share of the pot if Y bets all hands, so he won't fold outright.

Case 2:

First, we can find t_0 .

$$t_0 = \mathcal{T} [(1/(P+2)^t) - P^t] = \mathcal{T} [8/15]$$

$\mathcal{T} [8/15]$ is the t value for which $y(t)$ is greater than $8/15$. That includes all the hands with 100% equity and all the hands with $41/44$ equity, or $18/39$ of the hands. If the following is true, we have Case 2:

$$y^0(t_0) \leq (P+1)/(P+2)^t$$

$$y^0(18/39) \leq 48/63$$

The average equity of the top $18/39$ of Y's hands is just:

$$y^0(18/39) = (\frac{4}{18})(1) + (\frac{14}{18})(\frac{41}{44}) \approx 0.94697$$

Case 2 isn't true either, because the inequality above isn't satisfied. This also should make sense. Y has some pretty strong hands in his distribution in addition to a lot of weaker ones. If X were to call all the time, Y could selectively bet his strong hands to exploit X.

Since neither Case 1 nor Case 2 hold, we must be in Case 3.

Now what we must do is find a \mathcal{T} value for Y, such that if Y bets, X is indifferent to calling or folding.

$$y^0(\mathcal{T}) = (P+1)/(P+2)^t$$

$$y^0(\mathcal{T}) = 48/63$$

We can work backwards into this:

$$\{(1, 4/39), (41/44, 14/39), (6/44, 12/39), (2/44, 9/39)\}$$

$$48/63 = [(4/39)(1) + (14/39)(41/44) + (y)(6/44)] / (18/39 + y)$$

(where y is the number of $6/44$ equity hands that Y will bet)

$$48/63(18/39 + y) = [(4/39)(1) + (14/39)(41/44) + (y)(6/44)]$$

$$y \approx 0.13655$$

$$T = y + 18/39 \approx 0.59808$$

(You may wish to independently verify that $y^0(\mathcal{T}) = 48/63$)

So Y bets his best 59.8% of hands (approximately). We must also find X's calling frequency. (He seeks to make Y indifferent between bluffing and checking at T).

$$\alpha = (1 - y(\mathcal{T})[(P+2)^t - P^t]) / (P + 1 - y(t)(P+2)^t)$$

$$\alpha = (1 - (6/44)(15/64)) / (6 - (6/44)(63/8))$$

$$\alpha \approx 0.19651$$

So X folds about 19.6% of his hands.

This game has three important features that help to inform our decisions in real poker.

- Implied odds are an important part of optimal strategy in games with more than one street. These games are often more difficult to solve because of their increased complexity. Hence, game theory examples often consider the play on a single street and show how to make the opponent indifferent on the one street. In this game, the P^t idea introduces a quantitative measure of the implied odds. We can see from this that the strategy for each individual street is not the same as the strategy for that street in isolation. For example, on the turn, the pot contains 5 bets. Usually, if a player folds more often than $1/(P+1)$ of the time, he is exploitable by bluffing. However, in this case, with a pot size of 5, it is correct for X to fold almost $1/5$ of his hands (instead of $1/6$). This is because X is at a disadvantage not only on this, but on future streets. When he folds, he doesn't have to pay off on the river. Hence, he can fold more often optimally. We cannot emphasize this concept enough - multi-street games are not single-street games chained together; the solution to the full game is often quite different than the solutions to individual streets.
- The distributions that the players hold cause the play to be significantly different. When Y is quite strong, X can just fold; when Y is weaker, X can just call all the time. X can sometimes play accurately against Y's distribution.
- The players cannot make each other indifferent to taking certain actions with certain hands, in fact, there is usually just one threshold where they are indifferent, but that

threshold forms the linchpin of the entire strategy. In this case, the values for T and α represent those thresholds. X can vary the value of α fairly widely without giving up value as long as Y continues to play optimally. However, if X strays from the optimal α value, Y can exploit him by bluffing or not as appropriate.

More Advanced Topics In Games With Draws

The previous game was a fairly general case of a game where X could not bet on the turn. We have done some work on games with even more complexity; we omit describing some of these games because they are substantially more difficult, and often involve many cases related to the shapes of the distributions. However, we will briefly touch upon some of the principles we see from studying more complex games.

One important principle we find throughout looking at games with different types of draws is that concentrating equity in a few hands is stronger than having a distribution that contains many draws of moderate value. Contrast the following two distributions - one that contains 25% the nuts and 75% dead hands and one that contains 100% hands that have 25% equity. In terms of showdown equity, these distributions are identical; $y(I)$ for both of them is 0.25. However, the distribution that contains nut hands can play much more effectively in a two-street game, extracting additional value from his winning hands. When he holds a homogeneous set of draws, he basically has to wait until the draw comes in or doesn't in order to effectively bet and bluff. Distributions that contain more hands with high showdown value are in general more like clairvoyance than distributions that contain hands with medium-strength. And clairvoyance is effectively the limit of how well you can do playing a game.

As a result, it is important, as we said earlier, to keep one's distributions balanced with different types of hands, and to allow nut hands to balance draws whenever possible when deciding to carry hands from street to street.

In full two-street games we also see semi-bluff raising. When Y 's distribution is sufficiently weak or draw-heavy, X will bet out on the turn. Y should then raise with strong hands and balance those raises with semi-bluffs, while continuing to call with weaker draws and folding the weakest draws of all (beyond the bluffing hands he needs to carry forward, as in the first games we studied in this chapter),

A third principle we see is that in clairvoyant situations, open draws are often much weaker than closed draws. When a draw is closed, the non-clairvoyant player must fold just α of his hands. Contrast this to the results we saw in open draw games, where the non-clairvoyant player was able to fold more than half his hands, even when there were some closed draws present.

The complexity of multi-street games provides a tremendous challenge for us as poker analysts; there are many factors to be considered about the shape of the distributions, the size of the pot, and the structure of the game. And these are only two-street games where the bet size is the same on each street. If we had a full understanding of all the ways in which distributions made up of hands of varying strengths and varying draw potential interacted across multiple streets in optimal play, we would have a solution to any poker game. As it is, we simply continue to try to take out tiny chips from that large tree, by analyzing one situation at a time.

Key Concepts

- There are two principles at the core of analyzing multi-street play. The first is that in order for a bluff to be credible, it must come from a distribution of hands that contains strong hands as well. The second principle is also of prime importance in understanding how to apply game theory to poker. It is that multi-street games are not single-street

games chained together; the solution to the full game is often quite different than the solutions to individual streets.

- If a player has a distribution that contains only draws and does not threaten to have a strong hand in a particular situation, it is typically very bad for his expectation.
- Closed draws are much more powerful than open draws because the opponent is forced to pay off to value bets more often.
- Some draws (or made hands) are so strong that it is impossible to make them indifferent to calling a value bet. When choosing which draw to make indifferent, the answer is often the strongest draw that you can.
- Implied odds have a very important impact on action frequencies when draws are possible. In the clairvoyant game with arbitrary draws, the entire game revolves around equity in the implied pot - which contains additional value once the river is completed.
- If your opponent's distribution and yours are such that your opponent can take some action and you must fold and surrender the entire pot, then at any later point in the hand, it is at least co-optimal to fold to any future action.

Chapter 21

A Case Study: Using Game Theory

Over the course of Part III, we have spent a lot of time looking at various toy and real games, attempting to obtain insight into how the various dimensions of the game will affect optimal strategy. One of the hardest things to do, however, is to put these largely theoretical concepts into practice at the table. We stated at the outset of this book that we would avoid writing a manual on actual play, and even in this chapter, we will speak largely in generalities (particularly because of the complexity of factors that affect our decisions at each point). However, it seems that the best way to illustrate the process of applying theory in practice is to provide an overview and some examples of the way that *we* play particular situations in a particular game, demonstrating ways in which some of the various results we have obtained through the course of this book can be applied in practice.

Of course, we are not claiming that our strategy is optimal, unexploitable, or perfectly balanced. In fact, we know that this is not the case, because we know how difficult it is to play in an entirely balanced manner even in the simplest games. But hopefully you will see how our decisions are primarily influenced by two concepts: playing with balance and having an appropriate level of aggression for our value actions. The game we have chosen is the poker game that is most popular at this writing: no-limit holdem.

In discussing any no-limit game, the stack sizes are of paramount importance. In certain parts of this discussion, we will discuss situations and ideas that cover different stack sizes, while in others, we will limit the discussion to a specific stack size. In our discussion here, we will ignore our opponents' tendencies or tells as a source of information. Because poker is a multiplayer game, there are no optimal strategies (more on this in Part V). But we are searching for strategies that are near-optimal, or at least balanced, which can be profitably played against a wide range of opposition with little or no information.

A thorough discussion of this nature of a game as complex as no-limit holdem could likely fill more than one volume the size of this one; so we will try to keep our comments brief. You will also notice that we do not discuss "hands" here; a particular hand is of very little significance to us. Playing a particular betting sequence (including the hand distribution you carry to that sequence) is what near-optimal play is all about, and so we describe situations without reference to the cards we hold.

We begin by discussing preflop play (when we are the first player in).

Preflop Play (When First In)

The problem of selecting a strategy for play preflop is not an easy one. There are a number of factors that we must consider when choosing the hands with which we will enter the pot, and the strategy we will employ when doing this. Consider the case where we are first in; that is, all the players in front of us have folded

A preflop strategy is essentially a probability distribution for each hand, including how often we will raise (and to how much) with a hand, how often we will call with a hand, and how often we will fold the hand. Our experience examining games such as the [0,1] game indicates that many of these decisions will be pure; that is, most hands will not utilize mixed strategies. In general, hands that have mixed strategies will be threshold hands. In holdem, however, card removal effects are of some significant concern, and so we expect to find that more hands have mixed strategies than only those hands at the thresholds.

Identifying this strategy is no simple task, especially because even though we can break the strategy down into different elements, these elements are strongly tied together. For example, we can subdivide our preflop strategy into three major areas.

- Which hands to play
- Whether to raise or call with those hands
- What amount to raise with hands that will raise

Each of these items is somewhat complex in its own right: but they are also directly tied together. For example, suppose we are considering whether we should play 88 from six off the button with 20 big blinds in our stack. Our overall strategy for playing hands from this position will have a strong influence on our decision to play or not play this hand. Suppose we make a strategic decision to limp with all hands with which we will continue. Then 88 would seem to have positive equity when played with such a strategy. However, if our strategy is to raise to four times the blind from this position with all hands we will play, 88 is likely a money loser and should be folded.

And yet, in addressing the balance of the strategic alternatives we choose to employ across our entire strategy, we need information about each particular hand. Suppose for a moment that we decide to raise first in with any hand with which we will enter the pot, and to fold other hands. Then the only task that remains is to select the hands which merit a raise based on their expectation as part of a distribution of hands that plays in this manner. We might decide to limp in with any hand we will play preflop, folding the rest. Then all we need decide is whether to play a hand by calling or not, a decision that rests on the expectation of each of our hands when played as part of a distribution of hands that limps in this situation all the time. A third possibility is that we choose to employ a strategy of limping sometimes and raising sometimes. If we choose this alternative, then we are forced to consider the balance of both our raising and calling hands to ensure that our strategy does not give away enough information that it can be exploited. The point here is to consider the value of our entire distribution and find the distribution that maximizes our total expected value per hand.

This circularity makes defining a near-optimal preflop strategy extremely difficult. What we do in practice, however, is to choose to hold one of these variables constant, and from that make inroads to the others. It is our preference to play a strategy that never calls first in outside the blinds. There are a number of reasons for this. First, from each position, we will be playing a distribution of hands that is strong relative to the random hands left behind us. As a result, we expect to have the best hand, and benefit from more money going into the pot preflop. Second, we believe that the value of putting pressure on the big blind to call a raise in order to see the flop is higher than the potential lost value from raising instead of calling when our opponents hold strong hands. As we saw throughout Part III, making the opponent indifferent is an important principle of optimal play; in a sense we are attempting to create some land of indifference for the blind in cases where it is folded around to him. Third, we generally feel that the potential loss from the fact that we can play fewer hands in a raise-only strategy than we could in a strategy that was mixed between calling and raising is more than offset by the information hiding that playing a uniform strategy affords us.

The next major concern is that of position. Play from the front is perhaps the most poignant example of the power of position. Six other (presumed near-optimal) players will act after us, and each has the possibility of picking up a strong hand. Entering the pot with weak or marginal hands will likely result in being raised or reraised when one of those players has a strong hand. Therefore, we must take care in the hands that we select to play.

One of the principles of play in the front is to be selective, because of the threat of all the hands behind us.

As we move around toward the button, there are fewer and fewer random hands behind us, and so we can raise with weaker and weaker hands. On the button, then, we can play some wide variety of hands that rate to have positive expectation against the blinds, including the value of position on later streets. Still, though, we must take care to play only loosely enough to add hands that enhance the value of our distribution.

To find these hands, we consider our distribution as containing zero hands. Then we can add hands that have positive expectation one by one, until we can no longer add hands with positive expectation to our distribution. For each hand, however, we consider the value of the hand in the context of the new distribution. To understand what we mean by this, consider a simple example. Suppose that the first hand we add is AA, the strongest hand. Just assume temporarily that we are choosing to raise all hands first in, and that we raise to double the big blind in first position with AA and no other hands.

In our quest to be unexploitable, we can even take the position that we are going to announce to the table our strategy beforehand. Suppose we did this and our strategy was to raise AA alone. Then clearly our opponents would never reraise our raises, and would likely fold all their hands, provided the stacks were not terribly deep. So our equity here with AA would be something approximating the blinds. Obviously we can do better; suppose we add KK to the mixture. Our distribution is still quite strong, but the blind can probably call with hands containing a single ace. Still, the equity of each raise we make with a distribution of {KK+} is going to be worth approximately the blinds.

Suppose we go on adding hands to our distribution: QQ, AKs, Aco, and so on. At some point, we will have a weak enough distribution of hands such that our opponents will want to play with us especially out of the blind, but even in the field. When this occurs, it should be evident that though the value of the hands on the borderline is relatively small, the value of AA has increased. So when we raise with a distribution like {99+, AQ+}, AA has much higher value than when we raised it out of a distribution of {KK+}. In poker, it is generally very difficult to reduce the value of the pre-existing hands in a distribution by adding new, weaker hands. In fact, the pre-existing hands normally gain value by doing this. Constructing a distribution that maximizes the total value of our hands should be our goal.

In doing this, we must also take into account card removal effects and the effects of the community cards on the future play. In most of the games we have studied so far (such as the [0,1] game), hand values have been monotonic; that is, a particular hand is either stronger or weaker than another. However, the nature of the community cards in holdem weakens this assumption. Take, for example, the two hands AJs and 98s. Against a random hand (and against most normal distributions of hands), AJs is stronger. However, a difficulty arises because the flop has such a significant impact on hand values. There are many flops where the threat of holding 98s will add value to the play of the entire distribution of hands. This holds for many different types of hands. We saw in the chapters on multi-street play (Chapters 19 and 20) how important it is to threaten to have strong hands. Particularly when the stacks are deep, this effect becomes more and more pronounced. As the stacks deepen, we add more and more hands (particularly in fractional increments, using mixed strategies) to our raising distributions. The types of hands we add are hands which threaten to have or make the nuts on typical types of flops, including weaker pairs, AXs, suited connectors, and so on.

But this generally does not extend to playing weak hands for "deception" purposes. Optimal or near-optimal play does not contain such loss-leader plays. The ability to threaten to hold the nuts in a particular situation is worth some equity to all the hands in our distribution - but playing a hand with inferior showdown potential costs equity as well. With gigantic stacks (say, one million times the blind), it would indeed be indicated to play 52s under the gun with some positive probability, in order to threaten to have the nuts on boards with A34 of that suit. However, it would not be indicated to ever play 52o from under the gun.

The final piece of our first-in preflop discussion is to consider raise amounts. To understand the logic that we use to define the amounts we will raise, we begin by considering the conflict between the raiser and the blind. The blind is a key player in every hand. This is because the blind is the player most likely to call a raise from the field, due to the favorable odds afforded him by the dead money in the pot. There are two primary situations that face the raiser. When our opponents have a strong hand in the field, they will likely reraise or at least call. When no one behind us has a strong hand, it will be typical that the pot will be contested heads-up between the raiser and the blind. In choosing a raise amount, we will want to strike a balance between these two scenarios. In the first situation (where an opponent has a strong hand), we generally want to keep our investment in the pot small. By doing this, we lose less on the hands we have to fold to his reraise, while still retaining the option to commit more chips when we have stronger hands. Conversely, in the second scenario, we want to make our raise amount appropriate to put some pressure on the big blind while extracting the appropriate value. Raising to (for example) ten times the big blind with a strong distribution would force him to fold many hands where he would have enough equity to call a smaller raise, while if we were permitted by the rules to raise to 1,05 units, he would have an easy call with any two cards. We choose raise amounts that are intended to give the big blind a difficult decision when the field folds to our raise when he holds a hand of only moderate strength. We cannot make him indifferent to calling with his better hands - rather, we could do this, but it would sacrifice too much equity overall.

This means that our raise sizes should be varied according to the strength of our distribution, and as our position becomes later and later, we open with weaker and weaker hands. This is because the threat of players behind us is diminished as the number of potential hands behind us lessens. As a result, we can open with a weaker distribution of hands. To achieve the effect we just described when opening on the button with our weaker distribution, we have to raise a larger amount. So in our raise-only strategy, we choose a set of raise amounts that varies on position. We choose not to vary our raise amounts based on the strength of our hand - hopefully you realize that the information hiding implications of raising different amounts with different hands are indeed troublesome. The approximate amounts that we raise into a pot that contains only the blinds and no antes steadily increase from the minimum raise of doubling the big blind from positions five off the button or more to tripling the blind on the button. Antes generally cause us to raise a little more in order to attempt to preserve the same ratios for the blind. For example, suppose we would raise to 1200 to win 600 (200-400 blinds) from the button. Then the blind would be getting 1800 to 800 to call. Adding antes of 25 (eight players in), so that the pot is now 800, would make a similar ratio (if we raised to x) of $(x+800)$ to $(x-400)$. This occurs where $x = 1350$. So the antes of 200 result in a larger raise amount of about 150.

To recap, then, we employ a preflop strategy that has the following elements.

- When first in, we raise with any hand with which we intend to enter the pot.
- Our raise amounts vary from the minimum of doubling the blind (when five or six off the button) to approximately three times the blind on the button. When there are antes present, we modify our raise amounts to keep the same ratios of call to pot size.

- We play distributions of hands such that all the hands we play have positive expectation (as far as we can estimate) against the remaining field of near-optimal non-colluding players.

Jamming First In (Preflop Play with Short Stacks)

When we discussed raise amounts in the previous section, we assumed that the stacks were of sufficient size that we would always want to raise to some fraction of our stack and anticipate later play. However, as the size of our stack relative to the blind shrinks, we find ourselves in a situation where we may not want to raise to the smaller amounts described previously. Suppose that we have a stack size of just three big blinds and we are five off the button. We decide to raise with some hand. According to the guidelines of the previous section, we would raise to just double the big blind.

However, this would leave us in a very silly-looking position were the blind to call our raise. The pot would be $4\frac{1}{2}$ big blinds, and we would have remaining a stack of just one big blind. We would be hard-pressed to fold post-flop in such a situation. As a result, our opponent can effectively force us to commit at his option, either preflop or on the flop. Contrast this with the situation that would occur if we raised all-in preflop. Then our opponent lacks options: he can play or fold. Recall from Chapter 12 that strategic options have non-negative value. So, we should seek to prevent our opponent from having a strategic option at no cost. In this case, raising all-in is indicated in lieu of raising to two-thirds of our stack.

But suppose our stack is four, or six, or fifteen big blinds? At some point, our situation changes, and raising to the smaller amount is again indicated. This threshold is scarcely static - it clearly changes with position. But we can attempt to estimate where this threshold is in order to play more accurately preflop. We will take as our example a raise from early middle position (four off the button) and work backward from there. Suppose we are playing a distribution of hands (near the cutoff between jamming and raising a small amount) that is as follows:

{77+, AJs+, AQ+, KQs}

We can calculate the equity of this distribution against a typical reraising distribution, such as {TT+, AQs+, AK+}. Our raising distribution has showdown equity against this distribution of 41.9%. Using this number, it would seem that we would be committed to call a reraise as long as we had pot odds to call with this equity.

But this is a mistake - we are not interested here in how our entire distribution does against the reraiser. We know we are going to call an all-in reraise with aces! We are interested in the equity of our worst hands. In this case, our worst hands of each type are 77, which has 33.0% equity, AJs, which has 32.0% equity against the reraiser, and KQs, which has 31.88%.

Suppose that we identified a threshold x , which is the stack size such that if we raised the big blind of one to two units and were reraised all-in, we would be compelled to call. If we have a stack lower than x , then we remove a strategic option (flat calling the raise) from our opponent by jamming instead of raising to a small amount and therefore pick up whatever value that strategic option would have afforded him.

If we raise to 2, and someone behind us in the field raises to x , then we are calling $x-2$ units to win $x + 3.5$ units. In order for it to be correct to call this bet with all hands, our equity in the resultant pot with our weakest hand must be greater than the cost of calling.

$$(0.3188)(2x + 1.5) > x - 2$$

$$x > 6.83$$

So for this example, we find that a direct lower bound for x is a little less than seven big blinds. Of course, this is only applicable to this particular distribution of raising hands and this particular reraising distribution. We can treat this problem a little more generally. Consider the following game theory ideas:

- We should be indifferent to jamming with the weakest hand in our distribution. This is because our choice is between folding and jamming. Since this hand is the threshold between folding and jamming, the EV of jamming should be approximately zero.
- At the stack size that is the threshold between jamming and raising to a smaller amount, we will be indifferent between these two options (although perhaps we will play different hands).

Suppose we have a near-optimal opponent behind us. (Ignore for a moment hands left to act after him.) If he knows that this is the stack size at which we are indifferent between jamming and raising a smaller amount, he should be willing to commit with the same distribution of hands when we raise a smaller amount as when we jam. If he is not, we may exploit him by jamming with particular hands in our distribution and raising to a smaller amount with others.

Use N to designate the distribution of hands with which he will commit all-in against our raise to x units. Then N contains all hands that have positive equity confronting our jam. Suppose we jam a stack of 6 units with the aforementioned range: (77+, AJs+, KQs). Then N consists of all hands that have at least $\frac{12}{27}$ of the pot, or 44.44%. This distribution is: (99+, AKs, AKo). Against this range, our worst hand is 77, which has just 29.91% of the equity in the pot. However, had we raised to a smaller amount here, we would still be committed to calling an all-in reraise, as we have odds:

$$0.2991(13.5) - 4 > 0$$

Suppose we increased the stack size now to 8 units. Then the equity threshold for reraising rises to $\frac{16}{35}$ of the pot, or 45.7%. This excludes 99 from the reraising range, leaving us with (TT+, AKs, AKo). Against this distribution, 77 is still the worst, with 31.41% of the pot. If we raised to two units, a call is no longer indicated:

$$(0.3141)(17.5) - 6 < 0$$

Again, we see that somewhere between six and eight units is an appropriate cutoff if our only concern is whether we will be forced to call a reraise when we are in early position. We can perform the same calculation on a much wider raising range, such as when we are on the button:

{22+, A2+, K2s+, K8o+, Q8s+, Q9+, J8s+, JT, T8s+, 97s+, 87s}

Suppose here that we raise to 3 units out of our stack of 10 units, with the intention of calling any raises. Then N contains all hands that have positive equity against our distribution for 10 units. Hands in N have at least $\frac{20}{43}$, or 46.5% of the pot; this distribution is:

{33+, A2s+, A5o+, KTs+, KTo+, QJs}

This is a much wider distribution! After a reraise from this distribution, we would be calling 7 units to win a pot of 14.5. Our worst hand equity wise here is K8o, which has 33.4% equity against the reraising distribution. We can see that

$$(0.334)(21.5) - 7 > 0$$

So at this stack size, it is indicated to call with all hands after a reraise at a stack size of 10. In addition to this, there are other potential benefits to pushing all-in; we must compare the equity of pushing all-in as a strategy to the equity of raising to a smaller amount (assuming we do not want to break information bidding by raising two separate amounts).

There are a number of difficulties associated with this process. One is that it is fairly difficult to accurately estimate equity against a distribution when there is postflop play. A second difficulty is that many of the hands that have sufficient showdown equity to make them playable are very difficult to play after the flop. Consider the most basic example of this, AX offsuit. Suppose that we raise to a smaller amount, as above, such as 3 units from our stack of 12 units. On the flop, the pot will be 7.5 units and we will have 9 units left in our stack, a little more than a single pot-sized bet.

This stack size leaves little room to maneuver; after the opponent checks we have a difficulty. Obviously if an ace has hit the flop we are willing to commit all our chips; the same is true if we paired our kicker. But what about the large number of flops where we have flopped ace-high only? On these flops we have a substantial amount of equity - far too much to simply give up - but we cannot play in a manner that is terribly favorable to our implied odds. Our opponent will have flopped a pair $\frac{1}{3}$ of the time and some type of flush worth committing with an additional $\frac{1}{6}$ th of the time. Hence, our opponent will be able to commit accurately with his best 50% of hands. In this sense, he has better clairvoyance than we do; and he has effective position. Since he can automatically check, he effectively forces us to act first on the flop. Any strategy we employ in this situation has its drawbacks.

Checking behind risks giving a free card that beats us when we have the best hand, and exposes us to being bluffed out on the turn. To balance that, we get another card to try to make a pair as well, while taking back effective position.

Betting a small amount and folding to a check-raise all-in allows our opponent to steal the pot very frequently with hands that we beat; he can effectively turn his 40 or 50% equity in the pot into the whole pot with draws.

Betting a small amount and calling a check-raise does better against draws, but loses equity against made hands. By playing this strategy, we effectively allow our opponent to be selective about the hands with which he will commit. The same situation occurs if we simply jam the flop behind; we allow our opponent to be selective in calling our jam and allow him to identify the strength of his hand before doing this.

All these difficulties, however, are obviated by jamming preflop. At some stack size, however, it's clear that we expose too much (or alternatively, jam with too few hands) by jamming preflop. Hence, we must strike a balance between jamming outright and raising to a lesser amount. We can iterate the above analysis for a long time, looking for breakeven points. In fact, if we could estimate postflop equity very- accurately, we could likely arrive at a stack size for every position.

But as a practical matter, we advocate jamming with all hands (in lieu of raising to a lesser amount) for stack sizes of approximately 6 times the pre-raise pot from the front and 7 times the pot from near the button.

Play on the Flop

After preflop play has concluded and the flop has been dealt, a great deal of information about the eventual values of the hands has been revealed. Using this information to play accurately on the flop by assessing the shape of the distributions accurately is a critical skill in playing no-limit holdem well.

There are two major elements we focus on in examining the effect of the flop on future play. The first is who *the previous aggressor* was; that is, who was the player who last bet or raised preflop? This player will generally (in an environment of near-optimal play) possess a stronger distribution of hands than his opponent. Also included in his distribution will be weaker hands that balance his stronger hands. The second is the *texture* of the flop, a term used loosely to describe the approximate likelihood that either player's hand contains draws, and the strength of those draws.

For example, aboard like 9♦ 8♦ 7♣ contains a large number of potential draws. The shape of distributions (no matter how carefully balanced preflop) is altered by the nature of the board. We call boards like this one either *drawish* or *dynamic*. By contrast, a board like K♠ 8♣ 2♦ contains very few draws; in fact, most of the draws on this board will be five or six-outers to pairs or two pair. This type of board is more *static*. The extreme example of a static board is a board such as A♥ A♣ A♦, where there are just a few pair draws and few classes of hands overall. The ranks and texture of the flop affect the play of the distributions in ways that are extremely significant.

One aspect that is highlighted by the idea of the previous aggressor is that the distributions in real poker on the flop (or after some action has taken place) are nearly always asymmetric -that is, in contrast to many of our toy games, the players hold different sets of hands. Consider the situation where an early position raiser confronts a caller in the big blind. The blind will hold a significantly weaker distribution on the flop than the raiser, because he will have called with a much looser set of hands because of his reduced price for entering the pot. Even in situations where the stacks are deep and the raiser has a less tight set of hands, the blind's distribution will still be much weaker.

In fact, the raiser's range in this position is so strong that he likely has an autobet on most flops: that is, he should bet with all his hands if checked to. In this case, the blind should check all his hands (we can show, although we omit the proof here, that if the raiser has an autobet, then it must be at least co-optimal to check as the first player). We can see that this leads to a mode of play that gives deference to the previous aggressor (the raiser). It may be protested that the raiser can exploit this strategy of autochecking by checking behind, but we believe that the raiser will have higher equity utilizing an autobetting strategy than he would attempting to exploit by checking behind. This is because his distribution is so strong that he wants to either extract value or claim the pot; this situation is loosely analogous to the situation where a player has few bluffs, and so his opponent can reasonably fold a large portion of his hands (more than a) because he need not fear being exploited.

As the blind here, in addition to our weak hands, we will additionally have hands with value that we want to continue playing, such as pairs, two pair; sets, flush draws, straight draws, and so on. After the early raiser autobets, we must construct a strategy that maximizes value as well as being balanced. To do this, we must identify our strategic options and the hands that best fit those strategic options. After checking and having our opponent bet we can raise, call, or fold. It's clear from preceding material (and intuitively as well, we hope) that we will fold our worst hands. Additionally, we have the options of calling or raising. Assume the stacks here are quite deep.

A thorough, treatment of these options for all different stack sizes, raiser positions, and so on is beyond the scope of this chapter. However, we can identify some principles that may help to guide us in identifying a balanced strategy. We can generally say the following:

Both we and our opponent will take actions in "action regions" that will have threshold hands. (These thresholds may more formally be sets of hands due to card removal effects, but we can still visualize in this way.) In constructing a balanced strategy, our parameterization of regions must be such that we can create the proper indifferences between actions at the threshold points so that he cannot exploit us by moving his thresholds.

This might not be that easy to see, but suppose that we consider a check-raise. Our opponent has (for most stack sizes) three reasonable options. The first is to fold to our check-raise. The second is to call, and the third is to jam (or reraise). For each of these regions, there will be thresholds - in fact, because his non-folding actions will contain semi-bluffs to balance them, there will be multiple thresholds. Additionally, because this is not a static game, hands have different values - a draw value (which is the distribution of the hand's strength on the river) and a current value (which is the distribution of the hand's strength right now).

When we raise, then, we want to raise a mixture of hands such that our opponent cannot exploit us by moving his thresholds. Suppose that player Y has a threshold between y_1 and y_2 . Then his opponent must have hands in different regions such that the threshold hand is indifferent to playing y_2 or y_1 . This is a little difficult to ensure in practice, but we can effectively substitute the idea that if Y attempts to exploit him by expanding any one of the regions and playing that option more often, the opponent must have hands such that one or more regions of his hands benefit enough from that movement to offset or more than offset the loss to the other regions.

Suppose we divide all the blind's hands into the following categories:

- 1) Strong hands (i.e., sets, two pair, straights, and the like).
- 2) Medium strength made hands with weak draws (one pair hands).
- 3) Weak made hands with very strong draws (flush draws with additional value, such as a gutshot or pair).
- 4) Weak made hands with medium draws (flush draws or straight draws).
- 5) Weak hands without at least medium draws.

In truth, the only way to assess the relative weights of all these possibilities would be to calculate the optimal strategy' for a game with preflop distributions specified - even this is beyond current computing power. However, we can use the above to understand what a reasonable and balanced mixture of check-raising hands might be.

Hands that benefit the most from the opponent folding are draws, particularly weak ones. In fact, the weaker the draw, the more value is gained when the opponent folds. This of course must be balanced against the additional lost value against the opponent's stronger hands and even semibluffs.

If the opponent calls with more strong hands (that is, reraises less and calls more), we lose value with strong hands, as well as losing value with medium strength hands. Draws, however, benefit because they are not forced to put in additional money or fold to a reraise.

If the opponent calls with more weak hands (that is, folds less and calls more), we benefit with strong hands and medium strength hands, while losing value with weak draws.

If the opponent reraises more often, strong hands benefit, while draws are punished - except for the special case of very strong draws. These hands can react to a reraise by pushing all-in, as we saw in Part H.

What we seek to do in constructing a balanced strategy here is to identify the hands that have high value in one action, balance those hands, and also ensure that the remaining hands (which will play the other branch of check-calling) are also balanced.

Very weak draws have high value in check-raising, because they gain so much from the opponent folding. In fact, it is not unusual for many very weak draws to be unplayable if there was no raise option, but to have positive expectation from raising. Likewise, check-raising with a sizable portion of strong hands is also indicated. This is in keeping with the idea of correlating the size of the pot to the strength of your hand. The third type of hand, which (along with strong hands) helps to defend against looser reraising, is very strong draws, such as open-ended straight flush draws and gutshot straight flush draws. These hands have a very effective response to a reraise - to shove all-in. Recall that we discussed these types of hands in Part II. They benefit from getting all the money in, even against strong opposing hands. The worst outcome for them is to get some of the money in and then have to play more streets with a short stack relative to the pot size. This is because their value degrades on each street that they do not hit. In addition, when they do hit, they are less likely to be paid off. Likewise, since we can play these high-value draws in this way, we can play some very strong hands in this manner - this picks up value if our opponent calls looser (such as with one-pair hands) expecting us to be unbalanced toward playing strong draws.

If we construct a strategy of check-raising with some very weak draws (with the intention of giving up to a raise or a call if we miss on the turn), with some very strong hands (we will need some strong hands to balance our calling distribution), and very strong draws, we will include the hands that do the best by check-raising and will have achieved balance. This leaves hands of medium strength, some strong hands, and weaker draws (such as ace-high flush draws) with which we can call. This second set of hands is also balanced because it creates a situation where after the turn card, the distribution of hands will be such that the opponent cannot exploit by betting more or less often. This is because the combination of strong hands from the original distribution and hands that have become strong from the draws will be sufficient to prevent this type of exploitation.

That's a lot of ideas and things to think about, and all we've considered here is one decision point, in a fairly simple situation where one player is autobetting. Additionally, we neglected to specify bet/raise amounts, nor have we considered the texture of the flop, and so on. Playing strong exploitive poker is hard; it requires that you gather information in situations where that information is highly ambiguous and often misleading, calculate the inference from that information, and not expose yourself to counter-exploitation too much. But playing near-optimal poker is hard as well. Considering situations rather than hands and ensuring that one's play is balanced among the various alternatives is very difficult.

We next consider some ideas related to bet sizing. The texture of the flop is extremely important in the approach we take to bet sizing because of how it impacts the shape of the distributions. The presence of two suited cards on the board adds a prominent set of hands (flush draws) into the possible hands for both players. Since the draw is open, this will affect the play significantly. Consider first, however, a fairly static board, such as KK4, where we have been the preflop aggressor with a moderately strong distribution. On this type of flop, our tendency is to bet approximately the three-street geometric growth of pot, assuming that the stacks are moderately sized (such that geometric growth is between $\frac{1}{3}$ of the pot and a pot-sized bet). This betting

scheme gives us a fairly good balance between putting chips at risk when the opponent is strong and extracting value from his weaker draws. Since the quality of the draws on this board is very poor, we need not bet a lot in order to make the borderline draws indifferent.

When the stacks are larger, we often restrict ourselves to betting less than the pot; instead, we might use a four-street geometric growth scheme, allowing for a potential raise from our opponent to take up one of the streets. On static boards, we never overbet the pot, preferring to take advantage of the static nature of the board to provide value to our distribution.

When the board contains two flush cards or connected cards, however, the situation is different. First of all, a substantial fraction of the hands our opponent will continue with will have no value on the river. In order to make us indifferent to "bluffing" the opponent need only consider his hands that beat a bluff. What happens here is that on the river, the opponent with the weaker distribution, because of his knowledge of whether he held a weaker draw or a stronger one, has an advantage in the betting on the river. He is effectively playing with more clairvoyance than we are on the river - and when the flush misses, he need not pay us off with those draws at all. We will refer in general to "the obvious draw", which is usually the flush draw on a board where such is possible, or a straight draw on a board like $J\clubsuit T\heartsuit 4\clubsuit$.

Because of this enhanced clairvoyance effect, we are more inclined to bet larger amounts on earlier streets, offsetting the ex-showdown disadvantage we will have on later streets. On boards that have only a moderate amount of draws, such as $K\spadesuit 7\heartsuit 2\heartsuit$, we might simply increase our bet amount by perhaps half of the three-street geometric growth of pot. Hence, if we would bet $\frac{1}{2}$ of the pot geometrically, we might bet $\frac{3}{4}$ of the pot on a board like this one. When the board is extremely draw-oriented, such as the board discussed briefly earlier ($9\spadesuit 8\spadesuit 7\diamondsuit$), we will often make our bet amount the two-street geometric growth of pot, attempting to maximize against our opponent's draw-heavy distribution. And when the board has three flush cards, the distributions are even more impacted by the flop; here we often use the two-street geometric growth of pot as well.

Each of these bet amounts requires adjustments in the set of hands we play. If we are betting a larger amount, we should do so with fewer hands in order to prevent exploitation. This requires that we effectively balance our play by checking both hands with little value and hands with a lot of value. Take, for example, a three-flush flop such as $K\clubsuit 7\clubsuit 4\clubsuit$. On this type of flop, it is likely advisable (assuming that our bet amount on this flop would be approximately the pot) to check behind hands such as AQ with no club and the like. To balance this, then, we will need to check behind some hands with strong showdown value.

The danger of checking behind is that we give weaker hands a free card to beat us. Here the free card danger is exaggerated by the nature of the flop - giving a free card to a lone club can be disastrous. In trying to balance our play here, we look for hands that are not terribly hurt by giving a free card to a lone club, but have strong showdown value if blanks hit. Hands such as $A\clubsuit Ax$ or $A\clubsuit Kx$ fit this description perfectly; in a sense they are likely the best made hand and the best draw. By checking behind on the flop, we in a sense induce a bluffer semi-bluff on the turn. It is extremely important that we gain value with some of our hands to offset this.

Play on the Turn

We next consider play on the turn. This section will be substantially shorter than the section on flop play, mostly because many of the concepts that apply on the turn are the same as those that were discussed under flop play. The importance of balancing our play between the various strategic options, distributing our strong hands across our strategic options so that we threaten to have a strong hand from any action sequence, and so on are all just as important in turn play as

on the flop. The two areas in which turn play differs substantially from flop play are in the realization of draw equity and the complement of that situation, when a blank has come off on the turn.

Consider first the situation where no obvious draw has come in. The hand values of "made hands" remain more or less the same, with the draws that missed being downgraded substantially because they generally have just around 20% chance of coming in. Obviously, the rank of the card that hit on the turn can change the value of some medium- strength hands, either by making two-pair or by making a higher one-pair. Ordinarily, the stronger distribution (from the previous street) has the advantage in this situation, as it is often less likely to have the draw. Typically, our continuation betting strategy is structured as follows:

Suppose that on the flop, we had bet a set of hands. Of these hands, we will now check with following classes of hands:

- A few very strong hands.
- Hands that were of moderate strength on the flop, but unproved on the turn (for example, hands which hit two-pair).
- Some hands of moderate strength that will now check-call. (This group often includes pairs that are not top pair or overpairs or the like.)
- Some weaker hands that will now check-fold. (Often weak draws are in this category, or hands with just one overcard or the like.)

The hands that we continue to bet are:

- The other very strong hands.
- Hands of decent value, such as top pair or two-pair (with the intention of calling a raise).
- Stronger semi-bluffs (with the intention of calling a raise to balance the above).
- Some weaker semibluffs (that will fold to a raise).

The selection of hands here is not set in absolutes - obviously the nature of the board, the ranks of cards on it, and so on; each have an important effect on the selection of particular hands. What's important is that for each action we might take, we must select a distribution of hands such that our play cannot be exploited by an opponent. We know of many players who are extremely vulnerable to exploitation in a sequence where they raise preflop and get called behind. The flop comes down and they bet and get called, and they then check and fold to a bet from behind far too often relative to their checking frequency. They could make themselves less exploitable by checking more strong hands and hands that are able to call a bet on the turn.

When the draw does come in, many of the same principles apply. However, because many of the semi-bluffs have now been promoted to strong hands, many hands which would have been check-calls are demoted to weaker bets, while hands that would have been weaker bets might be demoted to check-folds. Likewise, when the draw comes in, often an additional draw is created. Suppose that the card that made the draw on the turn was a flush card. Then the bare ace of the flush draw is promoted to the nut draw against made flushes or other made hands.

In optimal strategy, having the backdoor nut flush draw on the flop (the bare ace of the two-flush on the board) is of real importance. Holding that card has two important effects. First of all, it creates the potential for a runner-runner nut draw, which is of course important for equity. Additionally it decreases the chance that the opponent is holding a strong flush draw when he pushes all-in on the flop. Suppose that you hold $A\clubsuit K\spadesuit$ on a flop of $K\clubsuit 7\clubsuit 2\heartsuit$, and the sequence evolves such that your opponent shoves all-in for a large raise on the flop. If you held the $A\heartsuit$,

your opponent's distribution might contain some $Ac Xc$ flush draws, expecting to often have 12 outs when called (9 flush cards and 3 aces). However, since you hold the $A\clubsuit$, you might give a little more weight to folding because the opponent's distribution contains none of these hands although he does not know this.

Card removal effects like this additionally have profound implications for river play, the topic we will consider in the next section.

Play on the River

The river is the simplest street to play on; it is also the street where the pot is largest and making mistakes has the highest cost. The river is the primary street where deficiencies of previous hand distribution selections turns into actual lost profit, though in the way that most players typically think of poker, this is an invisible effect. It is not too much of an exaggeration to say that the goal of a poker hand is to arrive at the river with a distribution that contains a mixture of strong, medium strength hands that did not reach their full potential due to the cards that fell, and weak hands and a pot size correlated to the strength of that distribution.

Play on the river, then, is quite comparable to the $[0,1]$ game with asymmetric distributions, except for the fact that the hand ranks are discrete instead of continuous, and there are card removal effects. In this light, we find ourselves often following through on value betting the geometric growth of pot on the river with hands that are appropriately strong.

However, taking a lesson from the no-limit AKQ game, we additionally employ "preemptive bets" on the river. These are small bets that we make with a mixed distribution of hands that have some showdown value, but don't want to call a large bluff on the river, and hands that are very strong. By doing this, we force the opponent to either raise us with their bluffs, in which case we win more money with our strong hands, or give up on those hands, in which case we gain from not losing the pot when they would have bluffed. Against that we lose additional money when the opponent raises, as we typically have to fold against a raise on the river. These preemptive bets are often as little as $1/10$ to $1/6$ of the pot.

The most common scenario in which we will make these types of bets is when the obvious draw has come in but our distribution contains a significant number of hands that were made by the flush card as well. In these cases, our distribution still contains a number of hands that are of mediocre strength and don't want to call a big bluff on the end, but still have value. Thus we make preemptive bets with a mixture of flushes and pairs (usually), along with standard value bets with some flushes and other hands.

Card removal also plays a significant role in deciding which hands to bluff. In general, when selecting from among the hands that have little enough showdown value to bluff, we should select hands which make it more likely that our opponent does not have a calling hand. In the same way, we should pay attention to our cards when deciding whether to call down an opponent who might be bluffing.

Consider a seven-card stud hand where the opponent has a marked flush draw and we have aces up. On the river, the opponent is effectively clairvoyant (assuming we do not have a full house). He will bluff to make us indifferent to calling if we do not catch one of his flush cards, because that is the most likely scenario. He cannot make us indifferent to calling if we catch one of his flush cards; hence we should simply call all the time when we do catch one, and an appropriate additional fraction of the otherwise.

Key Concepts

- Sometimes it may seem that you are being exploited on a particular street while playing near optimally. However, often this occurs because the opponent has sacrificed equity on a previous street in order to obtain the position where he can exploit on this street. For example, consider the game we studied in the multi-street section where the exposed hand needs only call $\frac{1}{6}$ of the time on the river because his opponent was forced to call a turn bet in order to see the river.
- For every action sequence, there should be corresponding balancing regions. Value, semi-bluffing (when there are cards to come) and bluffing regions should generally be seen together in a strategy.
- In every action sequence, provided there is a significant amount of action remaining, it is imperative to threaten to have a strong hand so that you cannot be exploited by the opponent wantonly overbetting the pot.
- The texture of the board affects the hand distributions profoundly, in constructing distributions for play on early streets, it's necessary to understand the expectation across all possible boards. This, of course, is impossible in practice due to computational infeasibility.

Part IV: Risk

*“My ventures are not in one bottom
trusted, Nor to one place; nor is my
whole estate Upon the fortune of this
present year: Therefore my merchandise
makes me not sad.”*

William Shakespeare, *"The Merchant of Venice"*

Chapter 22

Staying in Action: Risk of Ruin

Throughout this book, we have focused almost exclusively on expectation as a guide to proper play. Our assumption, stated at the outset, was that all players were adequately bankrolled for the games and that their only concern was to maximize EV. For this chapter, however, we suspend this assumption and consider instead the important question, "What does it mean to be adequately bankrolled for a game?"

Of course, there is no one single answer to this question; all players have different tolerances for risk. For some people, the risk of losing their entire poker bankroll is unacceptable: for others, playing poker may be just like a startup business, where a relatively high chance of failing is acceptable in pursuit of profit maximization. We can, however, apply quantitative methods to the problem by making assumptions, figuring out the implications of those assumptions, and building a model. One model that we use frequently in assessing bankroll requirements is called the *Risk of Ruin* (RoR) model.

A substantial part of the following was originally worked out by Bill Chen and Tom Weideman on rec.gambling.poker.

This model assumes the following:

- We repeatedly play a fixed game with a fixed distribution of outcomes \mathbf{X} ,
- Each trial is independent and randomly selected from \mathbf{X} .
- We have a starting bankroll b .
- The bankroll b is modified by the outcome of each trial.

We play the game indefinitely, or until the bankroll is reduced to zero.

For the assumptions outlined above, there are always only two possible outcomes for an infinite number of trials. The first possibility is that the bankroll grows without bound; that is, we win more and more and our long-term win rate approaches our win rate in the game. The second is that we lose the entire bankroll at some point. The value we are interested in obtaining is the probability that the starting bankroll will be lost under these conditions. We call this function the *risk of ruin function*, $R_X(b)$. We often shorten this to $R(b)$ when dealing with a single game and its corresponding distribution.

This function has certain properties that will allow us to describe its behavior and eventually find its value for a particular game.

Property 1.

If there is no probability of a negative outcome, then $R(b)=0$ for all b .

If we never have a negative outcome as a result of playing the game, then there is no chance of losing the entire bankroll.

Property 2.

If the game has negative expectation, then $R(b) = 1$ for all b .

This property is a result of the fact that the game will be played indefinitely. If the game has negative expectation, then the expected amount lost will grow without bound as the number of

trials goes to infinity. Contrast this with the situation where expectation is positive. Here the amount won will grow without bound except in the case where we lose the entire bankroll first.

Property 3.

If a game has any possible negative outcomes in \mathbf{X} , then $R(b) > 0$ for all b .

This property is easily justified as follows. Let's say that the probability of a certain negative outcome X_L is p , and the value of that negative outcome is v . If we have a bankroll B , then we can lose the entire bankroll immediately by having the outcome X_L occur B/v times in a row. This probability is simply $p^{B/v}$. This probability might be incredibly small, but as long as there exist any possible negative outcomes, $R(b) > 0$.

Property 4.

If a game has positive expectation and bounded negative outcomes, $R(b) < 1$ for all b .

Suppose X has positive expectation. Then let X_n be a sequence of independent trials. Define the cumulative result:

$$c_n = \sum_{j=1}^n x_j \text{ then we know that with probability 1: } \langle X \rangle = \lim_{n \rightarrow \infty} \frac{c_n}{n} > 0$$

Suppose that $R(b)=1$. Then also with (ruin) probability 1, we know that there will be a cumulative result that is negative for some C_n . Furthermore for each C_n we know there will be an $m > n$ such that $C_m < C_n$. This means we have a subsequence:

$$0 > C_{n1} > C_{n2} > C_{n3} > \dots$$

This contradicts the original assumption that X has positive expectation.

Property 5.

The risk of ruin function for the sum of two bankrolls is equal to the product of the individual risk of ruin values for each bankroll.

$$R(a + b) = R(a)R(b) \tag{22.1}$$

This property forms the basis for the derivation of the RoR function. If we have a bankroll of \$500, and some risk of ruin given that size bankroll, what is the chance that we go broke with \$1000? After we have lost the first \$500, which occurs with some specified probability, we now have another \$500 bankroll with which to play. Since losing the bankrolls are independent events, the joint probability of losing two \$500 bankrolls (equivalent to losing one \$1000 bankroll) is simply the probability of losing a \$500 bankroll squared.

There are some definitional issues with this statement; for example, say we have a game that consists of a \$100 wager, but we only have \$50. In this sense, we have been "ruined" already. However, we take the position that from an RoR standpoint, since we can put two \$50 bankrolls together to form a \$100 bankroll to play the game, \$50 has a lower risk of ruin than \$0 for this game. When we consider poker games, this distinction effectively never comes up, so we simply neglect it.

Example 22.1

Let us look at a simple example. Suppose we play the following game. We roll a die, and if the the comes up 1 or 2 we lose \$100. If the die comes up 3-6, we win \$100. This is a great game! We win more than \$33 per toss. But suppose we only have \$100. What is our chance of

eventually losing this \$100 if we play the game indefinitely? Using the properties above, we can actually solve this problem.

On the first roll, we have a $\frac{1}{3}$ chance of going broke, and a $\frac{2}{3}$ chance of growing our bankroll to \$200. Suppose we call \$100 "one unit." Then the risk of ruin of a bankroll of \$100 is $R(1)$, and the risk of ruin of a bankroll of \$200 is $R(2)$.

$$R(1) = \frac{1}{3} + \frac{2}{3} R(2)$$

We know From Property 5 above that $R(a + b) = R(a) + R(b)$, so:

$$R(2) = R(1)R(1)$$

Substituting and solving:

$$\begin{aligned} R(1) &= \frac{1}{3} + \frac{2}{3} R(1)^2 \\ 2R^2 - 3R + 1 &= 0 \\ (2R - 1)(R - 1) &= 0 \end{aligned}$$

This gives us two possible risk of ruin values, $\frac{1}{2}$ and 1. We know that this game has positive expectation and therefore by Property 4 above, the risk of ruin must be less than one. Hence, $R(1)$ for this game is $\frac{1}{2}$.

The reader can actually verify this by trials. Roll dice and play the game until either losing the initial \$100 stake or winning \$1000. If you repeat this process enough times, the results will converge on $\frac{1}{2}$,

Reminding to Equation 22.1:

$$R(a + b) = R(a)R(b)$$

We can do some algebraic manipulation which will help us to expose an important property of $R(x)$.

First, we take the natural logarithm of each side:

$$\ln R(a + b) = \ln R(a) + \ln R(b)$$

If we make up a function $f(x) = \ln R(x)$, then we have:

$$f(a + b) = f(a) + f(b)$$

This relationship shows that f is linear. We can see, for example:

$$\begin{aligned} f(1+1) &= f(1) + f(1) \\ f(2) &= 2f(1) \\ f(2+1) &= f(2) + f(1) \\ f(3) &= 3f(1) \end{aligned}$$

$$f(n) = nf(1)$$

...and so on.

So $f(x)$ has the form $f(x) = -\alpha x$, where α is a constant. (This is not to be confused with the α of Equation 11.1, the ratio of bluffs to value bets.) The risk of ruin for every game has this form, and we call α the **risk of ruin constant**. This causes α to be positive, α is actually equal to the natural logarithm of the risk of ruin of a bankroll of 1 unit. Because we know that $R(x)$ is between 0 and 1 for all x , we add the negative sign by convention, as $\ln x$ is negative for x between 0 and 1.

$$a = \ln (R(1))$$

α is a constant that tells us how much one unit of bankroll is worth in terms of risk of ruin. Suppose that we have one unit of bankroll. Then our risk of ruin is $e^{-\alpha}$. (e is a common mathematical constant - approximately 2.71 - that appears frequently throughout Part IV). We use the notation $\exp(x)$ and e^x interchangeably for ease of reading. Suppose that we have two units of bankroll. Then our risk of ruin is $e^{-\alpha}$ times $e^{-\alpha}$, or $e^{-2\alpha}$. And our risk of ruin for N units of bankroll is just $e^{-N\alpha}$.

So we have:

$$\ln R(x) = -\alpha x$$

$$R(x) = e^{-\alpha x}$$

So the form of the risk of ruin function is an exponential depending on the bankroll and some constant that is a function of the structure of the distribution of outcomes of the game.

Let's reconsider the example we had. Based on what we showed previously, we know that the risk of ruin for the 1 unit bankroll is $1/2$. The game where the bankroll is \$100 and the outcomes are +\$100 and -\$100 is equivalent to the game where the bankroll is 1 unit and the outcomes are +1 units and -1 units.

$$R(b) = e^{-\alpha b}$$

$$R(1) = 1/2$$

$$1/2 = e^{-\alpha}$$

$$-\alpha = \ln 1/2$$

$$\alpha = \ln 2$$

There is one more important property of the risk of ruin function.

Property 6.

For one play, $R(B) = \langle R(B + x) \rangle$, where x is a random variable selected from the outcome distribution \mathbf{X} .

This property states that the risk of ruin for a bankroll B is the expected value of the risk of ruin of the bankroll after one additional trial. This is because the risk of ruin model assumes that you will play the game indefinitely, one dial at a time. So if you have a bankroll B with a risk of ruin $R(B)$ right now, and you play the game once, $R(B)$ is the weighted average of your possible bankrolls after this trial. This may be a little confusing, but an example should help to show its simplicity.

Example 22.2

In our example game, we have the following outcomes:

+1, probability $2/3$

-1, probability $1/3$.

What this property says is that for any bankroll,

$$R(B) = \langle R(B + x) \rangle$$

$$R(B) = \frac{1}{3}R(B - l) + \frac{2}{3}R(B + l)$$

This relationship describes the relationship between risk of ruin and the expected outcomes of the game. After one trial, there is some number of discrete possible bankrolls, each corresponding to a particular outcome of the game. The expected value of the risk of ruin of each of those bankrolls must equal the current risk of ruin. This game is repeated an infinite number of times.

Using this, we have:

$$R(B) = \langle R(B + x) \rangle$$

$$e^{-\alpha B} = \langle e^{-\alpha B} e^{-\alpha x} \rangle$$

On the right side, $e^{-\alpha B}$ is a constant, so we can pull that outside of the expectation brackets:

$$e^{-\alpha B} = e^{-\alpha B} \langle e^{-\alpha x} \rangle$$

$$1 = \langle e^{-\alpha x} \rangle \tag{22.2}$$

We can use this important relationship to find the risk of ruin constant for simple risk of ruin problems directly, such as the example we have been working on:

$$\frac{1}{3}e^{\alpha} + \frac{2}{3}e^{-\alpha} = 1$$

$$e^{\alpha} + 2e^{-\alpha} = 3$$

$$1 + 2e^{-2\alpha} = 3e^{-\alpha}$$

Letting $x = e^{-\alpha}$

$$1 + 2x^2 = 3x$$

$$2x^2 - 3x + 1 = 0$$

$$(2x - 1)(x - 1) = 0$$

$$x = \frac{1}{2} \text{ or } x = 1$$

$$e^{-\alpha} = \frac{1}{2} \text{ or } e^{-\alpha} = 1$$

When we solve equation 22.2 for any game, one of the solutions is always $\alpha = 0$. For games with $EV < 0$, this is the correct solution to use, as $R(b) = 1$ for these games. However, for games with $EV > 0$, the answer is the second root, in this case $\alpha = \ln 2$. Again, this matches our previous solutions.

We should take some time here to review what we've done so far. First, we created a set of assumptions about a function called the risk of ruin function, which we called $R(b)$. Then we enumerated some properties of this function, including the extremely important property $R(a + b) = R(a)R(b)$. From this and the other properties of the function, we were able to derive a more direct formula for the risk of ruin of any game, $R(b) = e^{-\alpha b}$, where α is a constant that is a property of the game being studied. However, at this point, we had no reasonable way to find α other than solving for $R(l)$ directly.

Next, we stated an additional property of the function $R(b)$, which explained the relationship of the $R(b)$ to the expected value of risk of ruin after a single trial. Using this relationship, we were able to cancel out the bankroll-dependent portion of the risk of ruin function and show that *we can find the value of a directly from the parameters of the game, without solving any risk of ruin for any particular case.* This makes an entire class of problems tractable that would have been difficult under other circumstances. We continue by observing some things about the value of a .

We can make some interesting observations about the behavior of the function $\langle e^{-ax} \rangle$, which we will call $g(a)$ for now.

First, we know that $g(0) = 1$.

Second, we can take the partial derivative of g with respect to a to obtain:

$g'(a) = \langle -x e^{-ax} \rangle$. $g(0)$ then $= \langle -x \rangle$, Since we know that the game has positive expectation, $g'(0)$ is always negative. So at 0, g is 1 with a downward slope.

Third, we can find $g''(a) = \langle x^2 e^{-ax} \rangle$. This value is always positive.

Fourth, we can see that if there are any negative outcomes in \mathbf{X} , then $g(x)$ goes to infinity as a goes to infinity, because all terms with positive outcomes go to zero while those with negative terms increase without bound.

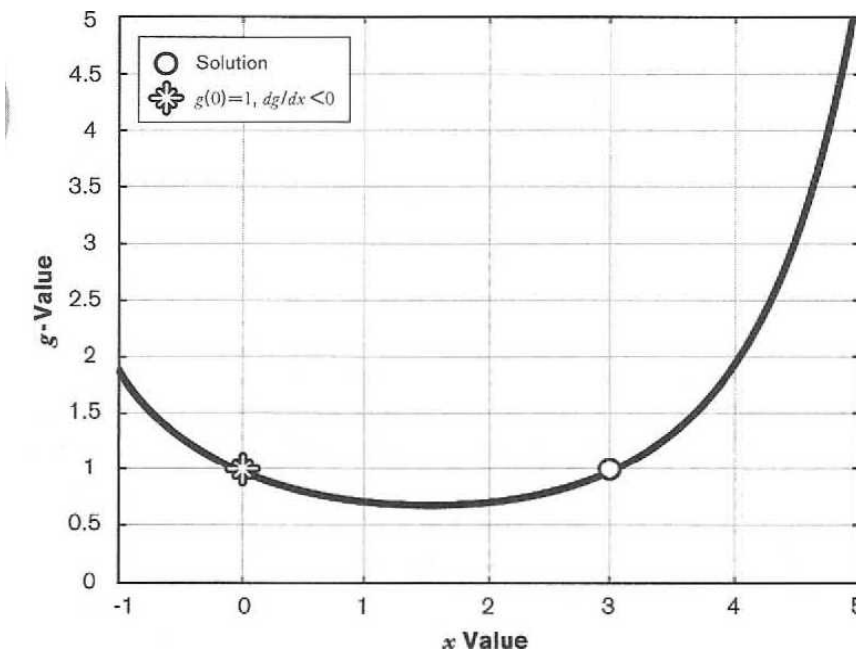


Figure 22.1. A sample of $g(x)$ for a game with positive EV.

Taking all these findings together, the graph of g must begin at 1, slope downward but with upward concavity, and eventually increase without bound as a goes to infinity. Therefore a solution to Equation 22.2 exists where $a > 0$, and that solution is unique. This obviates a potential difficulty that could arrive if we had several different a values that all satisfied Equation 22.2. Once we find a non-zero value of a , we know that it is the correct value to use.

This suggests, then, a method for solving any risk of ruin problem where the distribution of outcomes is known, which is to solve Equation 22.2 for the expected distribution of outcomes. We consider a more complex and real-world example next.

Example 22.3

Consider the case of a player on a limited bankroll. He is fairly risk-averse, as he has neither another means of support nor any way to raise an additional bankroll. He has decided that \$200+15 sit-and-go tournaments are the best mixture of return and variance reduction for him. These tournaments pay out the following:

- 1st place: \$1000
- 2nd place: \$600
- 3rd place: \$400

His expected finishes in the tournaments are as follows:

Finish	Probability	Net
1 st	0.12	\$785
2nd	0.13	\$385
3rd	0.13	\$185
4th-10th	0.62	-\$215
Total	1	EV = \$35/tournament

What is his risk of ruin for various bankroll sizes?

The method to use here is to set up Equation 22.1 for this player's outcome distribution.

$$0.12e^{-a(785)} + 0.13e^{-a(385)} + 0.13e^{-a(185)} + 0.62e^{-a(215)} = 1$$

Now this equation is not easily solvable by hand; however, there are many tools, such as Microsoft Excel, that include "solver" utilities that can solve such problems by converging on a solution. By putting the above equation into one of these, we have:

$$a \approx 0.000632$$

$$R(b) = \exp(-0.000632b).$$

Looking at the results of this for some different possible bankroll sizes:

Bankroll Size	Risk of Ruin
500	72.9%
1000	53.2%
2000	28.3%
3000	15.0%
4000	8.0%
5000	4.3%
7500	0.9%
10000	0.2%

A Word of Caution

Before we proceed any further, let us make particular note of an important assumption that underlies risk of ruin. This assumption is that the player will put all the money won in the game back into the bankroll and let it grow without bound. When we look at risk of ruin values as percentages, the temptation is to view these values as an ongoing and permanent chance that we will go broke given a certain bankroll that generates motley that we skim off the top. We cannot be too clear about this misconception.

The risk of ruin of a bankroll that is maintained at a particular level and from which all the profit above that level is taken out of the bankroll is 1, or a 100% chance.

If any money is taken out of the bankroll even if it is from profit, then the risk of ruin of that bankroll will naturally increase. Risk of ruin takes into account that some of the paths include making enough money initially to sustain a long losing streak in the future. There are many reasonable strategies for extracting money from a bankroll in a manner that allows us to rationally consider risk, but all of them increase the risk of ruin.

Half-Bankrolls

Those readers familiar with the sciences will recognize the structure of the risk of ruin function as being similar to that of other functions, such as radioactive decay. A radioactive substance contains (usually) certain isotopes of an element that are unstable. The unstable isotopes decay in a random way over a period of time, called the half-life. During this amount of time, the random decay process causes half of the radioactive matter inside a sample to decay. However, after two half-lives, the sample is not completely decayed. Instead, during the second period, half of the *remaining* sample decays, leaving $1/4$ of the original, and so on. The sample never truly decays completely, but simply approaches zero over a large number of half-lives.

We have analogized the process of risk of ruin to this; both are exponential functions and behave in similar ways. A *half-bankroll* is the amount of bankroll required such that the risk of ruin of the bankroll is exactly one-half. Then in the same way, the RoR of two half-bankrolls is $1/4$ of three half-bankrolls is $1/8$, and so on.

To find a formula for generating half-bankrolls, we simply solve the RoR equation such that the RoR is $1/2$:

$$\begin{aligned} e^{-ab} &= 1/2 \\ -ab &= -\ln 2 \\ b &= (\ln 2)/a \end{aligned}$$

So the size of a half-bankroll is $(\ln 2)/a$. We can approximate the half-bankroll for the sit-and-go tournament situation described earlier; recall that $\ln 2$ is about 0.693 and a was about 0.00632. So the size of a half-bankroll is approximately \$1100. Looking at the table of risk of ruin values above, we see that if we interpolate the risk of ruin for \$1100, \$2200, \$3300, etc, we see that they follow the geometric pattern of 0.5^b , which we expect.

Half-bankrolls are simply a special case of the more general risk of ruin formula, and can be of some use for approximating bankrolls required for various levels of risk of ruin. We often use these as a shorthand for the more complex risk of ruin formula as they can easily be compared across different games. Once we have solved for any particular risk of ruin, we can easily find the half-bankroll and then put together multiples of the half-bankroll for any desired risk of ruin.

Risk of Ruin for Normally Distributed Outcomes

Next we consider the risk of ruin formula for a game in which the distribution \mathbf{X} is normal, with a mean μ and a standard deviation σ .

Calling $v(x)$ the value function and $p(x)$ the probability density function for that value function, we have the following:

$$\langle v(x) \rangle = \int_{-\infty}^{\infty} p(x)v(x)dx \quad (22.3)$$

This identity provides the basis for extending our basic risk of ruin formula to this case where the distribution is normal.

The normal distribution is defined by its characteristic density function:

$$N(x) = \left(\frac{1}{\sigma\sqrt{2\pi}}\right) \exp\left(\frac{-(x-\mu)^2}{2\sigma^2}\right)$$

where μ is the mean of the distribution and σ its standard deviation.

We begin with equation 22.2:

$$\langle e^{-ax} \rangle = 1$$

Combining this with equation 22.3, we have:

$$\int_{-\infty}^{\infty} e^{-ax} N(x) dx = 1$$

$$\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp(-ax) \exp\left(\frac{-(x-\mu)^2}{2\sigma^2}\right) dx = 1$$

$$\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(\frac{-[(x-\mu)^2 + 2\alpha\sigma^2x]}{2\sigma^2}\right) dx = 1$$

Considering briefly only the exponent of e , we have:

$$\frac{-(x-\mu)^2 + 2\alpha\sigma^2x}{2\sigma^2}$$

$$\frac{-(x^2 - 2x\mu + \mu^2) + 2\alpha\sigma^2x}{2\sigma^2}$$

$$\frac{-(x^2 - 2x\mu + 2\alpha\sigma^2 + \mu^2)}{2\sigma^2}$$

$$\frac{-(x^2 - 2x(\mu - \alpha\sigma^2) + \mu^2)}{2\sigma^2}$$

We can then complete the square of the numerator of this function to obtain:

$$\frac{-(x^2 - 2x(\mu - \alpha\sigma^2) + (\mu - \alpha\sigma^2)^2 + \mu^2 - (\mu - \alpha\sigma^2)^2)}{2\sigma^2}$$

$$\frac{-[(x - (\mu - \alpha\sigma^2))^2 + \mu^2 - (\mu - \alpha\sigma^2)^2]}{2\sigma^2}$$

Substituting this back into our original equation:

$$\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(\frac{-[x - (\mu - \alpha\sigma^2) + \mu^2 - (\mu - \alpha\sigma^2)^2]}{2\sigma^2}\right) dx = 1$$

$$\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(\frac{-[x - (\mu - \alpha\sigma^2)]^2}{2\sigma^2}\right) \exp\left(\frac{-[\mu^2 - (\mu - \alpha\sigma^2)^2]}{2\sigma^2}\right) dx = 1$$

As the second exponential term of this equation contains no x variables, we can move it outside the integral and we are left with:

$$\exp\left(\frac{-[\mu^2 - (\mu - \alpha\sigma^2)^2]}{2\sigma^2}\right) \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(\frac{-[x - (\mu - \alpha\sigma^2)]^2}{2\sigma^2}\right) dx = 1$$

We also have the following:

$$\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(\frac{-[x - (\mu - \alpha\sigma^2)]^2}{2\sigma^2}\right) dx = 1$$

This is true because this equation is the integral of the density function of a different normal distribution, where the mean is $\mu - \alpha\sigma^2$ instead of simply μ .

So we are left with:

$$\exp\left(\frac{-[\mu^2 - (\mu - \alpha\sigma^2)^2]}{2\sigma^2}\right) = 1$$

Taking the natural logarithm of both sides:

$$(-\mu^2 - (\mu - \alpha\sigma^2)^2)/2\sigma^2 = 0$$

$$(-\mu^2 - (\mu - \alpha\sigma^2)^2) = 0$$

$$(\mu - \alpha\sigma^2)^2 = \mu^2$$

$$\pm\mu = \mu - \alpha\sigma^2$$

$$0 = a \text{ OR } 2\mu/\sigma^2 = a$$

So we again have our familiar two solutions for a . One is where $a = 0$, and the other (the critical value for this case), is where $a = 2\mu/\sigma^2$.

Substituting this back into our generic risk of ruin formula, we have:

$$R(b) = \exp\left(\frac{-2\mu b}{\sigma^2}\right) \quad (22.4)$$

This important result allows us to solve risk of ruin problems for games with normal distributions. However, poker results are not in and of themselves normally distributed; in order to attempt to apply this risk of ruin formula to poker, we must recognize the limitations of this approach. Poker games are not precisely normal. In general, poker hand results are clustered around particular outcomes, such as losing 1 small bet, losing 2 small bets, and so on. However, we know that the Central Limit Theorem allows us to conclude that the distribution of samples of sufficient size will converge on normal. The term "sufficient size" begs for clarification. Our ability to use the risk of ruin formula in the previous section is dependent upon the size of sample that allows us to assume normality.

If we have a distribution (of any shape) \mathbf{X} for a game and we have samples of a certain size such that the samples are roughly normal, we can calculate the risk of ruin of a game using the sampling distribution \mathbf{X}' . This will provide an approximation of the risk of ruin of the game: however, this risk of ruin is often biased because of the difference between the sampling distribution and the underlying distribution.

One of the primary sources of bias here is *skew*. Skew is a cousin of variance. When we looked at the variance of a distribution, we considered the distance of each outcome from the mean and squared it. This resulted in positive numbers for each possible outcome, which were summed to find the variance of the distribution. Skew instead sums the cubes of the distances, which causes the values to have sign. Distributions where the weighted sum of the cubes of the distances is positive are called skew positive and those where the sum is negative are called skew negative. When calculating skew, we also divide the sum of the cubed distances by the standard deviation of the distribution cubed.

The normal distribution has zero skew; it is symmetrical around the mean. By contrast, tournament distributions are highly skew positive. It is easy to see this by considering a 100-player, winner-take-all tournament. A player of average ability in this tournament would win 99 buyins 1% of the time and lose 1 buyin the other 99%. The *expectation* of this tournament would be zero, the *variance* would be 99 buyins²/tourney, and the *skew* of the tourney would be about +9.84. Skew effectively measures the tendency of the distribution to have outliers or larger values to one side of the mean or another.

Distributions that have a lot of skew cause problems for the normal approximation method of calculating risk of ruin. We can attempt to calculate the risk of ruin for this tournament in two ways: first by using the normal approximation, and second by calculating a directly. Since risk of ruin is only directly applicable to games with positive expectation, we will consider a player with a full buyin overlay; that is, a player who wins 2% of the tournaments (as described above) that he enters.

First, the normal approximation:

$$w = 1 \text{ buyin/tourney}$$

$$\sigma = 14.036 \text{ buyins/tourney}$$

$$R(b) = e^{-2b/197}$$

For a convenient value of 100 buyins, we have a risk of ruin of about 36.2%. Contrast this with the results we get from solving for a directly:

$$\langle e^{-ax} \rangle = 1$$

$$\frac{1}{50} (e^{-99a}) + \frac{49}{50} (e^a) = 1$$

Solving this equation by numerical methods, we get $a \approx 0.01683$.

Using this value in our RoR formula:

$$R(100) = e^{-1.683} = 18.6\%$$

So by solving directly for the risk of ruin of the game, we obtain a much lower risk of ruin. This is because the distribution of outcomes for a tournament has so much positive skew. It is possible to win 99 buyins in a single tournament; however at most we can lose 1 buyin. The normal approximation assumes that there will be a substantial number of larger negative outcomes, such as -5 or -10 buyins. However, the fact that these outcomes do not exist causes the normally distributed risk of ruin formula to overestimate the chance of initial ruin.

For this reason, we do not advocate the utilization of the normally distributed risk of ruin formula for tournaments. This, admittedly, was an extreme example because of the winner-take-all format, but the bias for skew right distributions is significant and present even for more graduated payout structures.

However, when we consider limit poker, we have a different situation. Here our initial distribution, though it is still somewhat skew right, is not nearly as distorted by this skew as the tournament example above.

As a starting point, we take the following distribution of hand outcomes, which, while simplified, is not too far away from what we might find for a typical mildly successful player.

Result in small bets (SB)	Probability	Weighted contribution
0	70%	0
-1	4.07%	-0.0407
-2	8.14%	-0.1627
-4	4.41%	-0.1763
-6	2.03%	-0.1220
-8	1.01%	-0.0814
-12	0.34%	-0.0407
+2	2.81%	0.0563
+4	3.44%	0.1375
+8	2.5%	0.2
+16	0.94%	0.15
+32	0.31%	0.1
Total	100%	0.02 SB/hand

If the reader disagrees with this characterization of the shape of the distribution of hand outcomes, the particular numbers chosen are not essential to the analysis that follows.

This distribution has a mean of 0.02 SB/hand and a variance of 10.80 SB²/hand. Using the normal approximation for a bankroll of 300 small bets, we find the risk of ruin to be:

$$e^{-(2)(300)(0.02)/10.80} = 32.89\%$$

We can set up Equation 22.1 for this distribution as well. This equation is fairly complex; we omit the details, although it is no different than our previous problems, only with more numerous terms added together.

Solving this with numerical methods, we find $a = 0.00378$. The risk of ruin of a 300-unit bankroll with this a value is:

$$e^{-300a} = 32.17\%$$

As you can see, these values are quite close; this property holds for most typical distributions of hand results. The point of this is that even though the skew of these distributions does exist, it is nowhere near as large as the skew of tournament outcomes (this remains the case even when we have graduated payout structures, as we shall see in a moment). Distributions of hand results and the normal distribution appear, in our experience, and from the data we've worked with, to be close enough to justify using the normal distribution risk of ruin as an accurate approximation to the full risk of ruin formula.

To review what we have covered so far. First we looked at tournaments. The primary characteristic of interest with regard to tournament distributions is that they are significantly skew right. When we apply the normal distribution risk of ruin function to tournament mean and variance, we overstate the risk of ruin significantly. This occurs because we lack large negatives; each tournament can only cost a single buyin total. Second, we considered limit ring games. Here we found that the distribution was sufficiently normal that the normal distribution risk of ruin formula was relatively accurate. So the characteristics of most importance in assessing risk of ruin for limit ring games are mean and variance.

A third type of distribution is common and increasing in popularity. This is the no-limit ring game distribution. Here we have large positives from hands where all the money goes in and we win, but also large negatives from hands where we lose our entire stacks. As a result of this, we find that the normal distribution risk of ruin again is a biased approximation. The characteristic of most importance here is *kurtosis*, which is the tendency of the tails of the distribution to be "fat," or the tendency of outcomes far from the mean to occur more often than the normal distribution would indicate. Recall that, for example, events three standard deviations away from the mean of the normal distribution occur only 0.3% of the time. However, in a game like no-limit, events several standard deviations away happen with a higher frequency. As it turns out, this effect is fairly minimal in its effect on risk of ruin for most distributions. Consider the following highly simplified distribution:

Result (in big blinds)	Probability	Weighted contribution
0	70%	0
+1	15%	0.15
-1	13%	-0.13
+50	1.1%	0.55
-50	0.9%	-0.45
Total	100%	0.12 BB/hand

The point here is that we can compare the risk of ruin values computed directly and from the normal distribution risk of ruin formula and see that the normal distribution understates the risk of ruin. Suppose our bankroll for this game is 500 units.

The mean of this distribution is 0.12 BB/hand, and the variance is 50.27 BB²/hand.

$$R(500) = e^{-2(0.12)(500)/50.27} = 9.19\%$$

Solving Equation 22.1 for this distribution, we find that

$$a = 0.004789$$

The risk from this value is:

$$e^{-0.004789(500)} = 9.12\%$$

So in this particular distribution, the kurtosis has very little effect. In fact, the risk of ruin is lower when calculated directly than by using the normal distribution. As a result of this and other calculations, it seems that the normal distribution provides an approximation with tolerable error that can be used in assessing risk of ruin in no-limit ring games as well.

However, it bears noting that no-limit ring game results can also be skew right in the same manner as tournaments, particularly for players who play very rightly in big pots. If a player loses many small pots but wins most of the big ones he enters due to playing style, the kurtosis has minimal effect, but the skew will tend to cause the risk of ruin formula to overestimate the chance of ruin.

We summarize the findings of this section:

Type of Game	Risk formula to use	Error of normal distribution
Limit ring games	Normal Distribution	Small
Tournaments	Direct Calculation	Large (using normal distribution results is values that are too high)
No-limit ring games	Normal Distribution	Small

Key Concepts

- The risk of ruin model can give us valuable insights into our chance of going broke from a particular bankroll level. This model assumes that we will play a single game with a known distribution of outcomes either indefinitely or until we lose the starting stake. The risk of ruin $R(b)$ is the chance that we will go broke starting from a bankroll b .
- Risk of ruin is 100% for games with negative expectation.
- Equation 22.2 provides us a method of calculating risk of ruin for any game with a fixed distribution of outcomes.
- Maintaining a bankroll at a given level will not provide insulation from risk of ruin; if we have a bankroll N and we take all the money above N out of the bankroll on a regular basis, our risk of ruin is 100%.
- Half-bankrolls are a useful shorthand for risk of ruin that can allow us to make off-the-cuff risk calculations by simply finding the number of half-bankrolls needed to assure a particular risk level.
- For normally distributed outcomes, Equation 22.4 gives the risk of ruin for a particular bankroll.
- For some distributions, such as tournaments, we cannot use the normal distribution formula because the outcome distributions are too skew. Instead, we can solve directly using Equation 22.2.

Chapter 23

Adding Uncertainty: Risk of Ruin with Uncertain Win Rates

One of the authors has a recurring vision of a world where tiny holographic projections appear above each player's head at the table, showing that player's precise "true" expectation in the current hand/game/session, including the players in the game, the player's state of mind, and so on. Unfortunately, such holograms do not exist, and so we must estimate and infer. In Part I, we discussed estimating parameters. In this section, we look at combining the risk of ruin formula we calculated in the previous chapter, which rely on a solid, known win rate, with an uncertainty distribution of win rates inferred from observed data. Throughout this section, we focus on limit poker, although these principles apply equally to no-limit ring games. However, since this analysis relies on the normal distribution risk of ruin formula, it cannot be used as is to solve problems involving tournaments.

Consider a player who has played a number of hours N with a win rate of w and a standard deviation of s . We can observe first of all that sample standard deviations converge on the population standard deviation quite quickly. So throughout this analysis, we will use s as the standard deviation σ of the game. However, we know from our experience that w does not converge on μ the player's "true" win rate, very quickly at all because the magnitude of s is high relative to w .

To see this, imagine that we knew $\mu=1$ bet/hr and $\sigma = 12$ bets/hr and we sampled 100 hours and called the result of that sampling w . The result w of this sample would be 68% likely to fall between 0 and 200 bets; 32% of the time, the result would be more than 1 bet/hr away from the mean. Let's say this player has a current bankroll of 200 bets. The difference between the risk of ruin at 0 bets/hr and that at 2 BB/hr is quite large!

0 bets:

$$R(200) = 1$$

2 bets:

$$R(200) = e^{-800/144} = 0.387\%$$

So this is a nearly 100% range of RoR values, all within one standard deviation of the population mean. There is even a fairly substantial difference between RoR at win rates of 0.8 BB/hr and 1 BB/hr:

0.8 bets:

$$R(200) = e^{-320/144} = 10.84\%$$

1 bet:

$$R(200) = e^{-400/144} = 6.22\%$$

However, the process of acquiring enough empirical data to tell the difference between the former and the latter is an arduous task. So we have a dilemma. We can use our win rate and standard deviation w and s and simply assume that they are accurate predictors; however, our true population mean may differ significantly from the observed values. A more robust approach is to account for the uncertainty of our win rate by using a distribution of win rates in lieu of

assigning a single win rate. We can do this by utilizing a normal distribution with a mean of our observed win rate and a standard deviation equal to the standard error of the observed win rate. Then we can utilize the normal distribution risk of ruin function as a value function and this hypothesized distribution of win rates as the corresponding density function. Then the expected value of the risk of ruin function will be the integral of the product of those two functions.

It seems valuable to take a moment to try to explain this process in somewhat less technical language. We have some observed win rate w . So we are saying that we can make a normal distribution curve around w . The closer we are to w , the higher the probability is that we are at our true win rate. At each w , we can use the risk of ruin function we've calculated all along to find the risk of ruin for that particular win rate. Then if we multiply all the risk of ruin values by the probability that the value is our true win rate and sum it all up, we get a "weighted" risk of ruin value that is more accurate than simply guessing at a particular win rate.

What's important to note here is that we can do this with only a little data about the game. If we only have a small number of hands, the dispersion of the distribution of win rates will be quite large. As we will see, the dispersion of the distribution of win rates is quite a critical value for finding this weighted risk of ruin.

The reader may also remember that we warned against using simply applying a normal distribution to observed results when we spoke of confidence intervals in Part I. This warning is still in effect; in fact, using this method doesn't incorporate information about the underlying distribution. This is particularly true where the observed mean is very high, when Bayesian analysis would dampen the observation relative to the population's win rate distribution. However, as this bias is restricted to the right tail of the normal distribution, its magnitude is fairly small compared to the gain in accuracy we get from applying the method. And attempting to correct this bias requires that we make attempts to characterize the shape of the *a priori* distribution from which this player is selected.

With that in mind, let us specify the terms of the model.

A player has played N hands with win rate w and standard deviation per hand s . The player has a bankroll b . His risk of ruin for any particular win rate is:

$$R(w, b) = \exp\left(\frac{-2wb}{s^2}\right)$$

$$f(x) = \exp\left(\frac{-2wb}{s^2}\right)$$

The standard error of his win rate observation is:

$$\sigma_w = s/\sqrt{N}$$

His observations about win rate create a normal density function around mean w and standard deviation σ_w .

$$p(x) = \left(\frac{1}{\sigma_w\sqrt{2\pi}}\right)\exp\left(\frac{-(x-w)^2}{2\sigma_w^2}\right)$$

Using Equation 22.3, we have:

$$\langle v(x) \rangle = \int_{-\infty}^{\infty} v(x)p(x)dx$$

This equation calls for us to integrate from $-\infty$ to $+\infty$; however, we know that the risk of ruin for win rates to the left of zero is always 1. Hence, we can split this integral into two parts: the integral from 0 to ∞ of the value function (risk of ruin) multiplied by the density function (normal for the distribution of win rates) plus the density function from $-\infty$ to 0. This last equals the probability that the win rate is equal to zero or less.

$$\langle v(x) \rangle = \left(\frac{1}{\sigma_w \sqrt{2\pi}} \right) \int_0^{\infty} \exp\left(\frac{-2xb}{\sigma^2}\right) \exp\left(\frac{-(x-w)^2}{2\sigma_w^2}\right) + p(w < 0)$$

The last term is the most important feature of this approach to the method. What it captures is that the uncertainty surrounding the win rate and the idea that having a bankroll of any size does not protect against the risk of ruin that comes from being a losing player. This probability, as we shall see shortly, is directly tied to the uncertainty σ_w ; the higher this dispersion is, the greater the chance that the win rate is below zero. Also, the risk of ruin for win rates that are positive but very low is also quite high itself.

We can also now point to a very practical reason for using the normal distribution for the distribution of win rates. Because both the normal distribution and the risk of ruin formulas are exponentials, the above expression simplifies!

$$\left(\frac{1}{\sigma_w \sqrt{2\pi}} \right) \int_0^{\infty} \exp\left(\frac{-2xb}{\sigma^2}\right) \exp\left(\frac{-(x-w)^2}{2\sigma_w^2}\right) dx + p(w < 0)$$

$$\left(\frac{1}{\sigma_w \sqrt{2\pi}} \right) \int_0^{\infty} \exp\left(\frac{-2xb}{\sigma^2} - \frac{x^2 - 2xw + w^2}{2\sigma_w^2}\right) dx + p(w < 0)$$

$$\left(\frac{1}{\sigma_w \sqrt{2\pi}} \right) \int_0^{\infty} \exp\left(-\frac{x^2 - 2x\left(2b\frac{\sigma_w^2}{\sigma^2} - w\right) + w^2}{2\sigma_w^2}\right) dx + p(w < 0)$$

Now call $u = 2b\frac{\sigma_w^2}{\sigma^2} - w$

$$\left(\frac{1}{\sigma_w \sqrt{2\pi}} \right) \int_0^{\infty} \exp\left(-\frac{(x+u)^2 + w^2 - u^2}{2\sigma_w^2}\right) dx + p(w < 0)$$

$$\left(\frac{1}{\sigma_w \sqrt{2\pi}} \right) \int_0^{\infty} \exp\left(-\frac{(x-(-u))^2}{2\sigma_w^2}\right) \exp\left(-\frac{w^2 - u^2}{2\sigma_w^2}\right) dx + p(w < 0)$$

That second exponential doesn't contain any x values, so we can pull it out of the integral, leaving us with a constant multiplied by a normal distribution function:

$$\exp\left(-\frac{w^2 - u^2}{2\sigma_w^2}\right) \left(\frac{1}{\sigma_w \sqrt{2\pi}} \right) \int_0^{\infty} \exp\left(-\frac{(x-(-u))^2}{2\sigma_w^2}\right) dx + p(w < 0)$$

Now we need to introduce a little bit of notation and discuss normal distributions a bit.

Recall that in Chapter 3 we defined:

$$\Phi(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^z \exp\left(-\frac{x^2}{2}\right)$$

where z was a normalized z -score.

$\Phi(t)$ represents the area under the normal distribution curve for a distribution with a mean of 0 and a standard deviation of 1 to the left of x , and we call it the **cumulative normal distribution function**.

If this normal distribution function had integrands of $-\infty$ and ∞ , then it would equal one. However, it is integrated from 0 to ∞ . What this means is that it equals the cumulative normal distribution function for a z -score of $-u/\sigma w$, or $(w - 2b \frac{\sigma_w^2}{\sigma^2})/\sigma w$. This type of value can be looked up in a table or retrieved using a spreadsheet function.

We are left with:

$$\exp\left(-\frac{w^2 - u^2}{2\sigma_w^2}\right) (\Phi(-u/\sigma w)) + p(w < 0)$$

$$-\frac{w^2 - u^2}{2\sigma_w^2} = -\frac{w^2 - (4b^2 \frac{\sigma_w^4}{\sigma^4} - 4bw \frac{\sigma_w^2}{\sigma^2} + w^2)}{2\sigma_w^2} = (2b^2 \frac{\sigma_w^2}{\sigma^4} - \frac{2bw}{\sigma^2})$$

$$\exp(2b^2 \frac{\sigma_w^2}{\sigma^4} - \frac{2bw}{\sigma^2}) (\Phi(-u/\sigma w)) + p(w < 0)$$

We have one additional simplification:

$$RoRU = R(w, b) \exp(2b^2 \frac{\sigma_w^2}{\sigma^4}) (\Phi((w - 2b \frac{\sigma_w^2}{\sigma^2})/\sigma w)) + (-w/\sigma w) \tag{23.1}$$

This is the formula for risk of ruin under uncertainty (RoRU). It contains what are essentially four separate items.

The first is $R(w, b)$, which is simply the normally distributed risk of ruin for the base observed win rate. This factor is then multiplied by the second element, an exponential involving the square of bankroll and the standard deviations of both the game and the win rate observations. All of this is then multiplied by a factor of the cumulative normal distribution for an additional expression involving w , the bankroll b , and the standard deviations of both distributions. This product (of the first three terms) is the weighted risk of ruin for the player for win rates that are positive. The last term is the cumulative standard normal distribution for win rates that are negative, multiplied by the risk of ruin for those win rates, which is 1.

Here's a table which looks at some risk of ruin calculations, using the normal distribution RoR and also this RoRU (risk of ruin under uncertainty) calculation:




Figure 23.1 RoRU vs win rate (bankroll -100, std. dev. -10 bets)

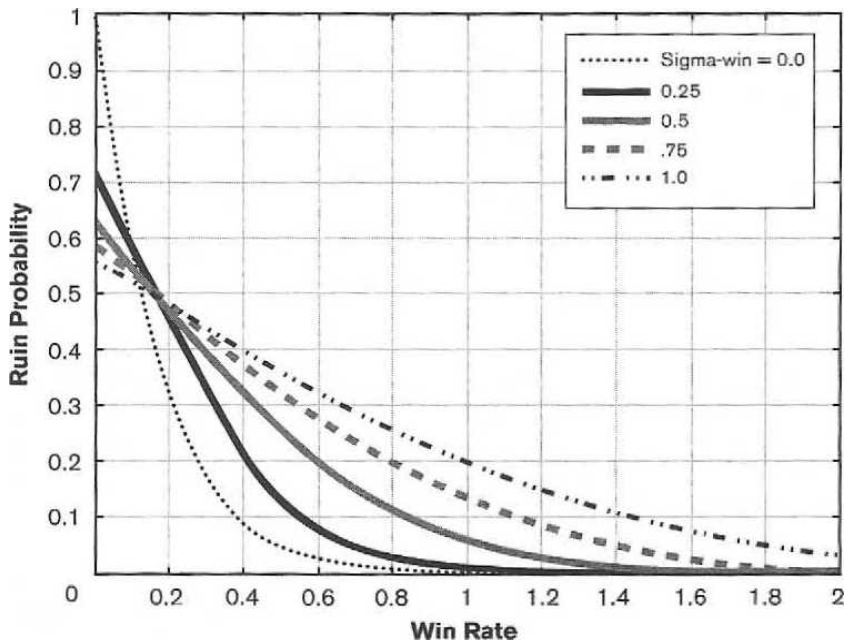


Figure 23.2 RoRU vs win rate (bankroll + 300 bets, std. dev. = 10 bets)

Of note in these figures is that when the uncertainty of win rate is high (the sample size is small), RoRU predicts a less than 100% chance of going broke, even if the win rate is negative up to this point. This is because the sample might just have been unlucky, and the true win rate might be above zero, in which case the RoR would be less than one,

RoRU is a powerful method of incorporating uncertainty about win rates into our risk of ruin calculations; we can adjust our expectations and understanding of the risk of going broke and include a quantitative adjustment for the possibility that we have just gotten somewhat lucky so far.

The risk of ruin under uncertainty formula we just derived is more complex, of course, than the previous formulas. The characteristic that makes it extremely useful, however, is that it attempts

to account in some way for the tendency for players to use win rates that are too high in estimating their risk of ruin. In terms of staying in action, it is difficult to overestimate the importance of this tendency and the resultant loss of bankroll principal that results.

Bankroll Considerations

We believe that most players overestimate their win rates for risk of ruin calculations for any of a number of reasons.

1) Initial Success - Players who begin their poker careers with a series of wins are more likely to continue playing. Players who begin their poker careers with a series of losses are far more likely to simply quit the game. This effect has perhaps a small selection effect toward players with higher EVs, but this effect is swamped by variance.

If we look at the set of all players who begin their careers by winning, we will find quite a number who are relatively poor players who have simply had positive variance events. These players will likely rely on their past (inflated) win rates as predictors of their future (expected) win rates and thus overestimate their prospects for the future. This effect is particularly pronounced for players who are considering playing professionally. The most likely time for a player to consider playing full-time is after some sort of upswing. Such upswings, while they may be real, are often positive statistical swings that may not reflect an underlying change in expectation.

2) Ego - Poker is an ego-driven game. Many players are convinced that they are much stronger players than they really are, that their results up to now have been due primarily to bad luck and not their poor play. This, of course, is sometimes true; but more often, it may be a combination of the two.

3) Misinformation - There is a tremendous amount of information floating around the poker community, both from public forums such as rec.gambling.poker, other Internet forums, magazines, and now even mainstream media sources. In these media, players can claim any (inflated) win rate they choose without any realistic chance of being exposed. This leads to a general inflation of "win rate stories," as individuals do not wish to feel inferior to other players. In addition, the poker literature has a disturbing tendency to dismiss questions of what constitutes realistic expectation in different types of games, preferring to rely on the "1 big bet per hour" platitude as a guide to what a typical winning player can expect.

4) Changing Game Conditions - Players try to project their win rates from now into the future. However, game conditions change - games dry up or shift to a different preferred game, as lowball games did with the introduction of holdem. Likewise, players may improve over time. Many things can change that make our estimates of the future inaccurate.

Very few poker players, especially those who have spent time and energy trying to improve, are willing to admit that there is a significant chance that they are losing players overall. In all honesty, overestimating one's win rate is very close to a universal property of poker players. There is one additional reason that win rates for risk of ruin calculations are often overstated, and it has to do with the fact that most poker players do not put all the winnings from poker into their bankroll. At this writing, one of your authors plays professionally. As a result, he spends money on things like mortgage payments, food, and so on. We speak about "win rate" and "bankroll" as it relates to risk of ruin using the same terms we do to describe the rate at which we win in a particular game and the total amount of money we have on hand to play poker with. But these are not at all the same.

First of all, in the United States, poker winnings are taxable. This can impact "win rate" by 30% or more, depending on how successful a player is. Second, "bankroll" should really be considered to be the amount of money we are willing to lose playing poker at a particular game and limit. For example, if a player plays \$20-40 and has \$12,000 set aside to play poker with, he might be said to have a bankroll of 300 bets. But if he will quit poker after losing \$6,000, for example, or if he will drop down to a lower limit (a topic we will take up later), his effective bankroll for \$20-40 is only 150 bets. He is not "ruined" in the sense of having no more capacity to play if he loses this amount, but it is illogical to consider the chance that he will lose \$12,000 playing \$20-40. In fact, he will never lose \$12,000 at this game. His "true" risk of ruin is some combination of his risk of losing at each game the amount that he will tolerate losing until he will step down yet again in limits, or stop playing poker altogether.

Third, professional players have living expenses. These have to be met whether the player wins or loses. This represents a steady outflow of money that must be considered in terms of bankroll and poker winnings. Even winning players who do not play professionally usually do not leave 100% of their poker winnings in the bankroll. They instead spend the money on vacations, jewelry, or plasma TVs. There are some different reasonable strategies for extracting money from poker winnings to meet expenses, whether luxury or necessity, but your authors prefer a particular method which enables us to continue to rationally assess risk while making steady outflows of money. We simply pay ourselves a salary denominated in some time unit and reduce our win rate by that amount. This is quite similar to starting a small business enterprise.

This may raise eyebrows with some of our readers. *Deduct the salary from win rate! What if the amount of money I need is almost as much as I make?* Herein lies the rub. If a player's adjusted win rate (after deducting the money that the player takes out for expenses) is close to zero (or even negative!), then **that player's risk of ruin is dramatically increased and will often be close to 100%**. Countless pros have ignored this principle, skimmed most of their profit off of the top of their bankroll, and eventually gone bust when a severe negative swing hit. Even if a player doesn't go broke, he often must forego income by being forced to drop down in limits (we will have more to say on this topic later.)

Every dollar removed from a bankroll reduces that bankroll's resiliency to negative swings, and so when calculating risk of ruin, it is imperative to use a win rate that accurately characterizes **not the amount of money won, but the amount of money put back into the bankroll as the game is played.**

For the aspiring professional, it gets worse; we have just shown a method by which we can assess the weighted risk of ruin, which is heavily impacted by the chance of a player's win rate being lower than his historical results would indicate. However, many players who are considering going professional after a somewhat longer time playing as an "amateur" have reached a threshold with their unadjusted win rate where the adjustment is not that large.

An example may help to elucidate this. Consider a player who has played 40,000 hands of his favorite game, and won 2.5 bets/100 hands with a standard deviation of 18 bets/100 hands. He has a bankroll of 300 bets. We can apply the RoRU function to find this player's weighted risk of ruin:

$$\text{RoRU} = 3.67\%$$

$$\text{RoR (simple)} = 0.98\%$$

The difference between these two numbers is just 2.69%; this is because the player has built up enough data that he is now almost three standard errors away from a zero win rate. As a result,

his weighted risk of ruin does not suffer much from being subjected to uncertainty.

The player is now considering going professional. He plans to take out from his monthly win approximately 1 bet/100 hands he plays to pay his modest living expenses, and to allow his bankroll to grow with the remainder. We can now recalculate his RoRU and RoR simple:

$$\text{RoRU} = 17.25\%$$

$$\text{RoR (simple)} = 6.22\%$$

What has happened here is that dropping from 2.5 bets/100 hands to 1.5 bets/100 hands (by taking out 1 bet/100 hands for expenses) has brought his risk of ruin higher, but also *significantly raised his chance of being a losing player after expenses*. Here, he is less than two standard deviations above zero. Remember, it is not enough to be a winner overall: a player must be a winner, and a significant one, in terms of the money that is put back into his bankroll in order to have a risk of ruin below one. In the case of professional players, this includes fiving expenses, business expenses, and so on. It also includes positive flows, such as a pro player's salary or income from investments in other players.

Virtual Bankroll

This even applies to non-professional players who only take money out of the bankroll after winning streaks; doing this decreases the capacity of the bankroll to weather future downswings and increases the risk of ruin of the bankroll. But this brings up another important consideration, which is the concept of *virtual bankroll*. After all, we do not play in a complete vacuum; for non-professionals and even some professionals, there may be money available to flow into the bankroll from other jobs' salaries, investments, or other sources of income. To do an exhaustive analysis of poker risk of ruin, all of these factors would need to be considered. But in some cases, an adjustment made in broad strokes is sufficient. For example, a serious player who has a job that pays \$300,000 a year has a virtual bankroll much larger than his present one, while a player who has no job at all and very few skills with which to get one might treat his poker bankroll a little more gingerly because it cannot be easily replaced.

The point of all this is that the impact on risk of ruin is not symmetric—slightly overestimating your win rate causes you underestimate your risk of ruin by a lot, while slightly underestimating it only causes you to overestimate it by a little.

One consideration that must color all of this analysis in a very important way is that game conditions are not static. Games ebb and flow, get better and worse. With the advent of televised poker, the world of poker players has grown immensely. It may be the case that this is a permanent change, but it's certainly not guaranteed. The same tiling is true of online poker. These types of things may also lead to an effect we discussed earlier. Upswings in profit are also correlated to favorable changes in conditions. Often when players are considering playing professionally, they have experienced some kind of winning streak. This may be due to a change in conditions in their preferred game which may or may not be permanent. The fact that conditions change makes it likely that most players ought to consider their overall risk of ruin to be slightly higher than the formulas presented here would suggest because it is more dramatically affected when conditions become worse than when they become better.

Key Concepts

- When we assess risk of ruin in the real world, we lack accuracy in our parameter estimates, such as win rate or standard deviation.
- We can partially correct for this by assuming that our true distribution of win rates is a normal distribution surrounding our point estimate.

- The risk of ruin under uncertainty model is the resultant risk of ruin formula, which includes an adjustment for the uncertainty surrounding the win rate point estimate.
- This method can be applied to situations where only a small sample has been recorded, and includes the idea that the effect of error in win rate estimation is not symmetric.
- When considering becoming a professional, a realistic understanding of risk is an important and often neglected consideration.

Chapter 24

Growing Bankrolls: The Kelly Criterion and Rational Game Selection

Even though this is a book on the mathematics of poker, we will take a short break from direct poker examples to consider a topic that is of critical importance in many fields. This is the so-called "Kelly criterion," which is of importance in stock portfolio theory, blackjack, sports betting, and any other field where the goal is to take a starting bankroll and some positive-EV wager, the player decides what size wager to make, and hopes to grow the bankroll as fast as possible.

We made reference in Part I to the concept of *utility*, which is a relatively broad concept that has many facets beyond the scope of this book. For our purposes, utility reflects a person's preference for money in an adjusted manner. We ignore in this book concepts like "feeling charitable," even though these are important to a full understanding of utility.

As an example of this concept, suppose we are offered a choice of two outcomes. We may receive \$20 immediately, or receive \$10 immediately. Clearly, we prefer the \$20. This is because we assign a higher utility to having \$20 than having \$10. Now consider the case where we can flip a fair coin for \$1,000,000, or receive \$500,000 immediately. In this case, we generally prefer the cash immediately. This is because the utility of a million dollars is not twice that of a half-million. 50% of the time we get the utility of a million, and 50% of the time we get the utility of zero.

We can construct a *utility function* which is a mapping of every bankroll x to its corresponding utility value $U(x)$. This function is unique to each individual and has the following properties:

- It is increasing for all values $x > 0$; that is, we never prefer a lower amount of money to the higher amount. If we did, if we had the higher amount we could simply give away the difference and increase our utility.
- It is usually concave downward. This simply means that the more money we have, the less utility an additional dollar will have. For example, the utility gained by a millionaire by finding a dollar in the street is much less than the utility gained by someone in poverty finding the same dollar.
- We generally treat it as though it were continuous (although since money is not strictly continuous, this is not precisely the case).

Suppose a bet B is proposed to us. There is some value of money X such that the utility $U(X)$ is equal to the utility of the bet, which is $\langle U(b) \rangle$, where b represents the outcomes of B . For very small amounts of money, X might be extremely close to $\langle B \rangle$, because utility is often close to linear for very small amounts. But on the other hand, if we consider the million dollar coin flip again, there is some value of money X which is the fair value of that bet to us in utility terms. This value X is called the *certainty equivalent* for B .

What the Kelly criterion attempts to achieve is to utilize a utility function such that the growth rate of a bankroll or portfolio is maximized. The reason for this is that we can maximize overall EV by simply betting the whole bankroll on any +EV bet. However, when we lose, we have less bankroll to grow from the next trial. To accomplish this, we maximize the Kelly utility function $U(x)$ for each bet instead of maximizing dollar EV. For reasons that are beyond the scope of this book, the utility function that maximizes the growth rate per trial is:

$$U(x) = \ln x$$

Interested readers can find a discussion of the derivation and selection of this utility function in any good finance textbook. Decisions that maximize this utility function are said to utilize the **Kelly criterion** or are simply called **Kelly betting**. Kelly betting has a number of important features that make it quite important in portfolio theory and in areas such as sports betting or blackjack. It:

- Maximizes the median bankroll after a given number of trials among all betting strategies.
- Minimizes the mean time required to reach a prespecified bankroll size among all betting strategies.
- Results in a distribution of bankrolls that is **log-normal**; that is, plotting the logarithm of the bankrolls results in a normal distribution. Because of the importance of this distribution in options trading and finance, it is quite well understood.

Example 24.1

As an example of how Kelly betting works, we consider the situation of a sports bettor. This bettor has some sort of model that allows him to pick games with a 3% edge; that is, his bets (including the bookmaker's juice) have a 3% advantage (they win 51.5% of the time and lose 48.5%), and are paid at even money. Now suppose this bettor has an initial bankroll of \$5000, and can bet in any amounts. What are the correct amounts to bet that maximize the rate of growth of his bankroll?

To calculate this, we seek to maximize the expected value of the Kelly utility function. So if our bettor bets x dollars, then there are two possible outcomes. 51.5% of the time, the resultant bankroll will be $(5000 + x)$, while 48.5% of the time, the resultant bankroll will be $(5000 - x)$.

$$\langle U(x) \rangle = 0.515(\ln(5000 + x)) + 0.485(\ln(5000 - x)).$$

To maximize this expression, we take the first derivative:

$$0.515/(5000 + x) - 0.485/(5000 - x) = 0$$

$$x = 150$$

So the amount that the bettor should bet on his 3% advantage is actually 3% of his bankroll; this is the case for all even-money bets. If he wins the first bet, he then bets 3% of \$5150, or \$154.50, while if he loses, then he bets 3% of \$4850, and so on. As this is not intended to be a text on gambling in general, we refer the reader to other sources for more information on this topic. The key concept we are trying to get across, which we will use in this section, is the substitution of the $[\ln x]$ utility function for the expected dollar value of a bankroll in making what we might call **metagame decisions**: those decisions which are not associated with a particular hand or a particular game, but rather, what games to play, what limits to play, or whether to play poker at all.

This **log utility** function leads us to a few comments about Kelly betting in general that we should keep in mind:

- Kelly betting has a risk of ruin of zero, assuming that we can subdivide our bankroll infinitely. The point is that when we lose, we scale our bet down immediately and when we win, we scale it up immediately.
- Kelly betting cannot apply directly to poker because there are not an infinite number of games at every limit. Although we might prefer it if our local cardrooms (or even online ones) would spread such limits as \$1.02-\$2.04 or \$12.46-\$24.92 in order to facilitate our

ability to maximize our bankroll growth, such a development seems unlikely.

- The logarithmic nature of Kelly betting implies some interesting properties. For example, doubling our bankroll is worth as much to the positive as losing half of our bankroll is worth to the negative. And as our bankroll becomes much smaller, it becomes a utility disaster to lose fractions of it, while winning large absolute amounts from a large bankroll is worth comparably less.
- Kelly betting lends itself to considerably more risk than many players are used to from the risk of ruin model; this is because the rapidly-scaling nature of the strategy makes it more resistant to risk of ruin. By contrast, most players do not like to move back down after moving up in limits. If a player is willing to aggressively move up and down as his bankroll fluctuates, he can effectively play at different limits on a significantly shorter bankroll than it would take to play at the highest limit all the time.

With regard to this last, we also feel that typical players are often much more risk-averse than is really warranted with regard to opening poker businesses. For example, many players want to have just a 5% chance of losing their starting stake. Most entrepreneurs, however, would be ecstatic to learn that they have just a 5% chance of failing. On the other hand, it is rather unlikely that one will be able to build up a poker business and then sell it for a large amount (in the manner of many entrepreneurs). But we should be rational in assessing the chance that we will go broke versus the rewards available when we do not. Often cautious players spend far too much time playing at lower limits than is justified by their particular situation and forego large amounts of profit by doing so.

Kelly betting provides a risk-based framework for assessing the utility of growing a bankroll; we can apply this framework to the problem of choosing between games at different limits. We have a finite number of game options; our local cardrooms, whether brick and mortar establishments or online poker rooms, carry only specific limits, and a limited number of games. Or there are not players to fill up empty tables at the limits we would prefer to play. As a result, we are often faced with a decision between two games. In some cases, the decision is totally clear. If we can have higher expectation and less variance in one game, then choosing that game is clearly indicated. But this is often not the case, and we must decide between playing in a game with higher return and higher variance (in absolute dollar terms) and playing in a game with lower return but lower variance.

We can use Kelly utility as a guide to making this decision; given that we are unsure what game to play, and we are willing to move up and down between the two. As a result, we can investigate the degree to which the two games affect our *utility* function and decide which game to play for a particular bankroll size.

Before we go further with this model, however, let us consider a different angle on the Kelly utility; here the question is the *optimal bankroll*. Previously we considered a fixed bankroll and found the optimal bet size; however, we can reverse the process for a game with a fixed bet and find the bankroll for which playing the game increases our Kelly utility maximally.

The Kelly utility for playing a game with a distribution of outcomes X is:

$$\langle \ln(B + x) \rangle - \ln B$$

where B is the bankroll and x is a random variable selected from the distribution X .

Hopefully, this is fairly intuitive. Essentially, the change in the Kelly utility from playing the game is the expected value of the Kelly utility after playing the game minus the Kelly utility

before playing it.

As a quick example, we return to a game we studied under risk of ruin. This is the game where we roll a die and win 1 unit when the die comes up 3-6, while losing 1 unit when the die comes up 1-2. This game effectively has a bet size of 1, since losing means a loss of 1 unit. If we have a bankroll B , then our Kelly utility from playing the game is:

$$U(B) = (1/3)(\ln(B - 1)) + 2/3(\ln(B + 1)) - \ln B$$

To maximize this, we take the derivative with respect to B :

$$\begin{aligned} U'(B) &= 1/(3(B - 1)) + 2/(3(B + 1)) - 1/B \\ 0 &= B(B + 1) + 2(B - 1)(B) - 3(B - 1)(B + 1) \\ 0 &= B^2 + B + 2B^2 - 2B - 3B^2 + 3 \\ B &= 3 \end{aligned}$$

So for this game, the growth of Kelly utility from playing the game is maximized when the bankroll is equal to 3.

With this concept of optimal bankroll in mind, we can make the following assumptions for this model:

- We have two games to be compared; we will call these games game 1 and game 2. Game 1 is the larger of the two games.
- We have known win rates μ_1 and μ_2 as well as standard deviations σ_1 and σ_2 , for the two games. Both games have positive win rates and finite variance.
- We have a bankroll B .

Obviously, if we were forced to play one of these two games and we had a very large bankroll, we would play the higher of the two, while if our bankroll was quite small, we would play the lower. Our model here is that there is a bankroll cutoff c such that above c playing the higher game will increase our Kelly utility by more and below c the lower game will increase out-Kelly utility by more. We seek to find this bankroll cutoff.

For each game, we have some distribution of outcomes \mathbf{X} . Our change in Kelly utility for each distribution is equal to:

$$\begin{aligned} U &= \langle \ln(B + x) \rangle - \ln B \\ U &= \langle \ln B(1 + x/B) \rangle - \ln B \\ U &= \langle \ln B \rangle + \langle \ln(1 + x/B) \rangle - \ln B \end{aligned}$$

$\langle \ln B \rangle$ is a constant, so we have:

$$U = \langle \ln(1 + x/B) \rangle$$

We can use the Taylor series expansion for $\ln(1 + t)$:

$$\ln(1 + t) = t - \frac{t^2}{2!} + \frac{t^3}{3!} - \frac{t^4}{4!} + [\dots]$$

to find:

$$\ln\left(1 + \frac{x}{B}\right) = \frac{x}{B} - \frac{x^2}{2!B^2} + \frac{x^3}{3!B^3} - \frac{x^4}{4!B^4} + [\dots]$$

If we use the assumption that B is substantially bigger than each outcome x, we can ignore the third and higher terms of this expansion, since they contain x^n in the numerator but $(n!)(B^n)$ in the denominator and thus approach zero extremely rapidly.

This leaves us with:

$$U \approx \left\langle \frac{x}{B} - \frac{x^2}{2B^2} \right\rangle$$

For any distribution, we have $\langle x \rangle = \mu$ and also the following:

$$\begin{aligned} \langle (x - \mu)^2 \rangle &= \sigma^2 \\ \langle (x^2 - 2x\mu + \mu^2) \rangle &= \sigma^2 \\ \langle x^2 \rangle - 2\langle x \rangle \mu + \mu^2 &= \sigma^2 \\ \langle x^2 \rangle - \mu^2 &= \sigma^2 \\ \langle x^2 \rangle &= \mu^2 + \sigma^2 \end{aligned}$$

Substituting into our formula for U, we have:

$$U \approx \frac{\mu}{B} - \frac{\mu^2 + \sigma^2}{2B^2}$$

The cutoff bankroll c that we are looking for is the bankroll at which the utility of Game 1 equals the utility of Game 2. Above this value, we will play Game 1; below it, we will play Game 2.

$$\frac{\mu_1}{c} - \frac{\mu_1^2 + \sigma_1^2}{2c^2} = \frac{\mu_2}{c} - \frac{\mu_2^2 + \sigma_2^2}{2c^2}$$

$$\begin{aligned} 2c\mu_1 - \mu_1^2 - \sigma_1^2 &= 2c\mu_2 - \mu_2^2 - \sigma_2^2 \\ 2c(\mu_1 - \mu_2) &= \mu_1^2 - \mu_2^2 + \sigma_1^2 - \sigma_2^2 \\ 2c &= \mu_1 + \mu_2 + (\sigma_1^2 - \sigma_2^2)/(\mu_1 - \mu_2) \end{aligned}$$

Because in poker, μ_1 and μ_2 are normally extremely small compared to the variances, we can abbreviate this formula to:

$$c \approx \frac{\sigma_1^2 - \sigma_2^2}{2(\mu_1 - \mu_2)} \tag{24.1}$$

Let's look at a more concrete example.

Example 24.2

Suppose there are two games we're considering playing in a cardroom; one is \$20-40 holdem, where we have a win rate of \$35/hr with a standard deviation of \$400/hr, and the other is \$40-80 holdem, where we have a win rate of \$60/hr, with a standard deviation of \$800/hr.

Then the cutoff bankroll is:

$$\begin{aligned} 2c &= 60 + 35 + (800^2 - 400^2)/(60 - 35) \\ c &= \$19,295 \end{aligned}$$

This means that with a bankroll of more than this amount, we should play the \$40-80 game, while with less than this amount we should prefer to play \$20-40 from a utility standpoint.

This framework provides a means for comparing the risk-adjusted utility of playing in two different games so that we can make the best decision about which game to play in a quantitative manner. The primary limitation on using this formula is that it is invalid for distributions of outcomes where some of the outcomes are of comparable size to the bankroll (in these cases, dropping the later terms of the Taylor expansion is not reasonable). This occurs most frequently in tournaments, especially those with large fields, where the top payouts may exceed 500 buyins. Nevertheless, for all types of ring games we can quantify the risk-adjusted value of playing in any two different games using this method.

Key Concepts

- When making real-world decisions, we cannot simply use expectation as a guide; instead we must consider utility. We can assign a utility function and evaluate outcomes in terms of utility.
- One common utility function which is studied very heavily in financial markets because of its special properties in portfolio theory is the Kelly criterion, or
- $U(x) = \ln x$.
- By substituting the Kelly criterion into two possible sets of game conditions, we can find the utility of playing in each game. When the games have appropriate conditions, we can identify a bankroll level above which we should play the larger game, and below which we should play the smaller.

Chapter 25

Poker Finance: Portfolio Theory and Backing Agreements

In addition to the Kelly criterion, many concepts from finance apply to the management of a poker bankroll and the selection of games, in this chapter, we will examine two concepts of particular importance. The first is examining the volatility and expectation of a portfolio of poker activities. As a parallel to securities investing, certain types of diversification are helpful, while others are not. Second, we provide a framework for evaluating the expectation of a backing agreement between a player and a backer.

Portfolio Theory

In portfolio management, one critical statistic that is used to evaluate performance is the *Sharpe ratio*. The Sharpe ratio of a portfolio is the ratio of its expectation to its standard deviation. When investments are infinitely scalable (for most individual investors, stocks are), maximizing the Sharpe ratio of a portfolio leads to the best risk-adjusted return.

$$S = \frac{\mu}{\sigma} \quad (25.1)$$

One of the main ways in which portfolio managers can maximize their Sharpe ratio is through the process of *diversification*; that is, by buying multiple securities that have a higher Sharpe ratio when bought together than separately. The following is a very short and simplified introduction. Those of you with experience in finance will recognize that we are leaving out many important facets of stock portfolio analysis; the point here is to explain in a rough way the principle without worrying about things that do not affect poker.

Consider two independent investments A and B. Both A and B have an expectation of \$10 and a standard deviation of \$100 (for the amount we have to invest). If we invest it all in A, we'll get \$10 expectation with a \$100 standard deviation. The same applies if we invest it in B. But if we invest half of it in A and half of it in B, then we'll still have an expectation of \$10, but our standard deviation will be $\$100/\sqrt{2}$, or about \$70. By doing this, we have increased our Sharpe ratio from 0.1 to about 0.14.

Now let's consider a portfolio made up of A and a third investment C, C has an expectation of \$5 and a standard deviation of \$100. We can see immediately that C is a worse investment than A, since we can get \$10 of expectation for \$100 in risk by buying A. To maximize our Sharpe ratio, then, should we simply buy 100% of A? Perhaps surprisingly, the answer is no.

Let's evaluate a portfolio made up of 80% A and 20% C. Such a portfolio would have expectation of $(0.8) (\$10) + (0.2) (\$5)$ or \$9. It would have a standard deviation of

$\sqrt{((80)^2 + (20)^2)} = \82.46 . So the Sharpe ratio for this portfolio would be 0.1091. This value is higher than the Sharpe ratio for a portfolio made up of only A, so it turns out to be better to mix the two investments, even though C is a worse investment in an absolute sense.

We can apply this reasoning to poker as well. Consider a situation outside of a pair of tournaments. One player (A) has an expectation of 1 buyin per tournament and a standard deviation of 9 buyins/tournament in his tournament. Another player has an expectation of 1 buyin per tournament and a standard deviation of 12 buyins/tournament in a different tournament. The first player wants to maximize his Sharpe ratio. What is the optimal percentage to swap with the second player?

Let $1 - \alpha$ be the fraction that these two players swap. Then A is investing in a portfolio P of a A and $(1 - \alpha)$ B. The win rate of that portfolio will still be 1 buyin per tournament, since all the components have that win rate.

We want to minimize the standard deviation of that portfolio, which we can also accomplish by minimizing the variance:

$$\sigma P^2 = (9a)^2 + ((1 - a)12)^2$$

Taking the derivative with respect to a and setting to 0:

$$0 = 162a - 288(1 - a)$$

$$0 = 450a - 288$$

$$a = 0.64$$

A should swap 36% of himself in order to maximize his Sharpe ratio.

We can derive this result more generally. Say we have two investments with expectation w_1 and w_2 , and standard deviations σ_1 , and σ_2 . We can normalize these bets to each other in terms of expectation by simply multiplying σ_1 by w_2/w_1 . We will call the scaled standard deviation s_1 . If we do this, then both investments will have the same win rate w_2 . We will mix these two investments in a proportion of a_1 of the first and a_2 of the second.

$$a_1 = 1 - a_2$$

The variance of the portfolio P will be:

$$\sigma P^2 = (a_1^2)(s_1^2) + (a_2^2)(\sigma_2^2)$$

We have omitted some steps here in deriving the variance. This equation is only valid for independent events, where there is no correlation. When analyzing securities, correlation is very important. Large-field tournament play, even in the same tournament, for players playing to maximize their own expectation is only correlated a very tiny bit, so we omit that for the purposes of this discussion.

Since $a_2 = 1 - a_1$ we have:

$$\sigma P^2 = (a_1^2)(s_1^2) + (1 - a_1)^2(\sigma_2^2)$$

Taking the derivative with respect to a_1 and setting to 0:

$$0 = 2a_1s_1^2 - 2(1 - a_1)^2(\sigma_2^2)$$

$$2\sigma_2^2 = 2a_1s_1^2 + 2a_1\sigma_2^2$$

$$a_1 = \frac{\sigma_2^2}{\sigma_2^2 + s_1^2} \tag{25.2}$$

We can also easily find a_2 :

$$a_2 = \frac{s_1^2}{\sigma_2^2 + s_1^2} \tag{25.3}$$

In order to maximize the Sharpe ratio of a portfolio, these are the percentages of each investment

of which the portfolio should be composed. In finance, this type of analysis is extended to an arbitrary number of different investments and optimal portfolios can be constructed. In the same way, poker players can reduce their volatility while preserving their Sharpe ratio by swapping action, either in tournaments or in ring games. This can even be true if one of the players has higher expectation: that player can still benefit Sharpe-wise by swapping with the other player.

Those familiar with stock analysis might be tempted to think that other types of diversification might help to improve portfolio performance: for example, playing different types of games (sit-and-go tournaments, multi-table tournaments, limit ring games, etc.). We believe that playing many games is valuable for a number of reasons, including the flexibility to take advantage of temporary situations (such as when a very weak and very rich player is playing his favorite game for high stakes), and developing a more general view of poker. But portfolio management is not one of the reasons. The reason that this type of diversification doesn't lower volatility is that it isn't really diversification at all.

When we own a stock, we own it 100% of the time. When we own two stocks, we own both of them 100% of the time, but only 50% as much. However, when we play two different types of poker, it is analogous to a situation where we buy a stock, hold it for half the time, and then sell it and buy another stock. As a result, playing different types of games doesn't reduce variance directly. However, one related concept can lower variance while retaining the same win rate: that is online multi-table play. If a player can concentrate adequately on multiple games, he can lower his variance by, for example, playing two games of \$10-\$20 instead of one game of \$20-40. Of course, he might be able to play two games of \$20-\$40, so he might be sacrificing equity. But by playing two games of smaller stakes, we can achieve similar expectation with lower variance, thus improving our Sharpe ratio.

Backing Agreements

The last topic for Part IV, a discussion of the mathematics of backing agreements, is a topic that is of some interest for those members of the poker community who are interested in backing players. Player backing agreements have existed as long as there have been players who lack sufficient bankroll to play the stakes they wish to. The usual form of these agreements is this: the player receives some sum of money (or a smaller amount, usually with a promise to advance more if needed up to a certain point) from the backer. The player plays some agreed-upon game or games, and then at certain points, the agreement settles and some prespecified division of the profits takes place. If there are losses, the backer usually absorbs them all, but the player must make up any previous losses before cashing out any profit. There may also be a salary component, in which case the player's share of profits is often reduced dramatically.

As our starting point, we take as a simple backing agreement the following:

- The player has a known win rate in a particular game of \$40/hr and standard deviation of \$600/hr but no bankroll.
- The backer and player agree that the duration of the agreement will be 100 hours of play and the backer will fund all play. At the end of this, the profits will be divided evenly, or if there are losses, the backer will take them all, and the agreement will be terminated (although both parties may agree to enter into another agreement).
- The player agrees to play to maximize play EV instead of his own personal EV

Over the past two chapters we borrowed heavily from the finance world. In evaluating the value of backing agreements, we do so again. This backing agreement, then, takes this form. "We will randomly select some value from a prespecified probability distribution. If the value selected is below zero, then the value to the player is zero. If the value selected is above zero, then the value

to the player is $1/2$ the value selected." This is equivalent to an **option**. In stock trading, an option is a contract that gives its holder the right to either buy a stock at a particular price, or sell the stock at a particular price before a certain time. To the player, this backing agreement is essentially a "call" option (where the player may buy half the result) at a strike price of zero.

We can use this to find the value of the backing agreement to the player. The distribution of outcomes for 100 hours of play is normally distributed around a mean of \$4,000, with a standard error of \$6,000. The value to the player of this agreement depends on the value of the player's result if he wins. To find this value, we can integrate the product of the player's result (x) and the normal distribution density function from zero to ∞ to get the value of the player's win when he wins.

$$\frac{1}{\sigma\sqrt{2\pi}} \int_0^{\infty} x e^{-\frac{(x-w)^2}{2\sigma^2}} dx \tag{25.4}$$

We can make the substitution $x = x + w - w$ to obtain:

$$\frac{1}{\sigma\sqrt{2\pi}} \int_0^{\infty} [(x - w)e^{-\frac{(x-w)^2}{2\sigma^2}} + we^{-\frac{(x-w)^2}{2\sigma^2}} dx]$$

We can split this into two integrals. The second integral can be simplified by factoring out the w and the remaining expression is just the cumulative normal distribution:

$$\frac{1}{\sigma\sqrt{2\pi}} \int_0^{\infty} (x - w)e^{-\frac{(x-w)^2}{2\sigma^2}} dx + w\Phi\left(\frac{w}{\sigma}\right) \tag{25.5}$$

To find the antiderivative of the expression $(x - w)e^{-\frac{(x-w)^2}{2\sigma^2}}$, we make the substitution $u = \frac{-(x-w)^2}{2}$, $du = -(x - w)$. This leaves us with:

$$\int -e^{\frac{u}{\sigma^2}} du = -\sigma^2 e^{-\frac{(x-w)^2}{2\sigma^2}} + C$$

We can now simplify Equation 25.5:

$$\begin{aligned} & \frac{1}{\sigma\sqrt{2\pi}} \int_0^{\infty} (x - w)e^{-\frac{(x-w)^2}{2\sigma^2}} dx + w\Phi\left(\frac{w}{\sigma}\right) \\ &= \frac{1}{\sigma\sqrt{2\pi}} - \sigma^2 e^{-\frac{(x-w)^2}{2\sigma^2}} \Big|_0^{\infty} \\ &< \text{wins} > = \frac{\sigma}{\sqrt{2\pi}} e^{-\frac{w^2}{2\sigma^2}} + w\Phi\left(\frac{w}{\sigma}\right) \end{aligned} \tag{25.6}$$

Equation 25.6 is the value of the entire option; that is, it is the value of the agreement if the player gets 100% of the win and takes none of the loss. To convert this value to the value of the backing agreement, we simply multiply by the player's share.

For our particular player, we can solve for the total value of the wins:

$$w = 4000$$

$$\sigma = 6000$$

$$z = (4000/6000) = 2/3$$

$$(6000/\sqrt{2\pi})(e^{-0.222222}) + (4000) (\Phi(2/3)) = \$4,906.72$$

Since we know that the total expectation for the player must be \$4,000, we can find that the value of the losses is -\$906.72.

For our example backing agreement, the player's expectation, then, is half of the win and none of the loss, or \$2,453.36. while the backer's share is all of the loss plus half of the win, or \$1,547.64. Thus the player gets about approximately 61% of the expectation of the backing agreement.

This methodology can be used to cover any backing agreement of this type. We can extend equation 25.6 by creating a threshold value a such that the player's share of the profits begins with wins above a by simply integrating from a to ∞ Instead of from zero.

Then we have:

$$\langle \text{wins} \rangle = \frac{\sigma}{\sqrt{2\pi}} e^{-\frac{(w-a)^2}{2\sigma^2}} + w\Phi\left(\frac{w-a}{\sigma}\right) \quad (25.7)$$

If we used the same backing agreement as before but the player's share only began at wins of \$1,000 or more, for example, then we would have:

$$(6000/\sqrt{2\pi})(e^{-.125}) + (4000) (\Phi(1/2)) = \$4,878.24$$

Obviously this value is smaller, since wins from \$0 to \$1,000 are excluded from its value. Looking at various values for the threshold, we can see a potential negotiating tactic for the player. Consider a player who proposes as follows:

"Instead of getting half the profits when I make more than zero, I think that I should only be cut in if I exceed the expectation that we've agreed upon. However, since I'm giving up some money for those times that I win less than my expectation, I should get a larger piece of the profits above expectation."

This line of reasoning is sound. But how much higher a percentage of the profits should this player receive? We invite the reader to test his intuition here before continuing to see the answer.

The value of the option with a threshold of \$4,000 is:

$$(6000/\sqrt{2\pi})(e^0) + (4000) (\Phi(0)) = 4,393.65$$

If the player was to receive 50% of the value of the wins at threshold \$0, he would receive \$2,453.36. To reach parity with that value, he "deserves" 55.8% of the win above \$4,000.

Another technique that can be applied here is creating a tiered structure, where the player receives different percentages of amounts won, usually increasing as the player's win grows. Consider the following backing agreement:

- The player has a known win rate in a particular game of \$20/hr and standard deviation of

\$240/hr.

- The backer and player agree that the duration of the agreement will be 100 hours of play and the backer will fund all play. At the end of this time, any losses will be absorbed by the backer, If there are profits, they will be divided according to the following schedule:

Profit Tier	Backer Share	Player Share
0 - \$1000	75%	25%
\$1000 - \$2000	60%	40%
\$2000 - \$3000	50%	50%
\$3000 - \$4000	40%	60%
\$4000 - \$6000	30%	70%
\$6000+	25%	75%

In the above table, the backer/player shares are what the player earns on that amount of profit. For example, if the player's result is \$1,500, he earns 25% on the first \$1,000 and 40% on the remainder, or \$450.

We can calculate the value of such an agreement using Equation 25.7,

$$w = \$2,000$$

$$\sigma = \$2,400$$

The value of the total option for the first tier (\$0-\$1,000) is simply the value of the total option with threshold 0 minus the value of the total option with threshold 1000:

$$(2400/\sqrt{2\pi})(e^{-.34722}) + (2000)(Z(5/6)) = \$2,271.93$$

$$(2400/\sqrt{2\pi})(e^{-.08681}) + (2000)(Z(5/12)) = \$2,200.93$$

The difference is \$71. This is the value of the total option for the first tier. We can do the same sorts of calculations to find the values of each tier:

Profit Tier	Backer Share	Player Share	Tier Value	Value to Player
\$0 - \$1000	75%	25%	\$71.00	\$17.75
\$1000 - \$2000	60%	40%	\$243.47	\$97.39
\$2000 - \$3000	50%	50%	\$402.69	\$201.34
\$3000 - \$4000	40%	60%	\$473.53	\$284.12
\$4000 - \$6000	30%	70%	\$746.02	\$522.84
\$6000+	25%	75%	\$334.33	\$250.74
Total			\$2,271.93	\$1,374.19

This is roughly equivalent to a 60.4% share in the profits above zero.

This last example shows an alternative method for structuring backing agreements that might be slightly more effective at solving one of the two big issues facing backing agreements in general. These two issues are at the core of deciding how to structure a backing agreement. The first issue is the problem of *aligning the interests of the player with the interests of the backer*.

One common problem that occurs quite frequently in backing agreements of a fixed term is that when the option expiry date (the end of the backing agreement) nears, if the player is at least

moderately far away from the backing agreement threshold, the player has a direct incentive to take on a lot of variance. To take an extreme example, assume this is the last hand before the backing agreement expires. The player is \$2,000 away from the threshold value where he begins to get shares of the profits. The player is offered an opportunity to put \$3,000 in the pot with QQ against AA preflop (assume the AA player accidentally exposed his hand or the like.) This would be a terrible play for the backer and under normal circumstances, and the player as well. However, the player's EV from gambling here is actually about 18% of his split of \$1,000, while his EV from folding is zero. Here the player's EV function is dramatically different and in this extreme case effectively opposed to the backer's.

One way of addressing this problem that has been used in some circumstances involves paying some sort of salary to the player that he receives win or lose, while reducing the player's share of the outcomes. However, this has the added effect of making it less valuable for the player to play well and try to win. This problem is familiar to managers in many sales industries: some automobile dealerships have experimented with commissioned employees and salaried salespeople and so on.

The second issue is that of *fairly compensating both parties*. Both sides want to make money, and presumably the backing agreement should benefit both parties. But what elements should be included in this "fair" division of profit? Some principles that merit consideration are:

- A risk premium for the backer.
The backer is assuming all the risk in these agreements and therefore ought to receive a premium for that risk-taking.
- The player may not always need backing.
Often players who are being backed could instead play at some lower limit for some reduced rate. The player may feel that his gain from the backing agreement should reflect this. Viewed another way, this is like payment for the player's time, which could be spent in some other way. Since the backer is effectively making money passively on the deal, the player may demand a larger share of the profits.
- Uncertainty about the player's win rate.
Often in creating backing agreements, some implicit or explicit win rate and standard deviation will be used. However, if the player's win rate is actually *less* than the parties believe, then the player's share of the overall value of the option increases. For example, assume 100 hours of play with $w = \$20$ and $\sigma = \$200$, compared to 100 hours of player with $w = \$15$ and $\sigma = \$200$. The total value of the player's play is now three-fourths as much. However, the value of the 0-threshold option is 81.3% as much with $w = \$15$ as with $w = \$20$. The player retains a higher percentage of his share of the win than the backer if his actual win rate turns out to be lower than the parties believe *a priori*.

None of these considerations include other items such as the risk of credit default or fraud, where the player absconds with the backer's money or fails to accurately report results. Of course, backing is a high-risk activity; however, the returns from a strong, stable backing agreement can far outstrip returns from other types of investments. Another problem with backing is that players who win often don't need to be backed for long; by contrast, players who are break-even or win slightly often require backing for longer periods of time in order to be profitable.

Key Concepts

- A poker bankroll shares many characteristics with a portfolio, and so elements of portfolio theory can be brought to bear on it.
- Sharpe's ratio provides a metric for risk-adjusted return. The higher the ratio, the higher the return relative to risk of the investment.

- Players can benefit utility-wise from swapping with each other. This is even true if one player is stronger than the other but the swap takes place at par.
- Backing agreements share many characteristics with options, and can be valued similarly in deciding on a fair deal.

Part V: Other Topics

*There are more things in heaven
and earth, Horatio, Than are
dreamt of in your philosophy.*

William Shakespeare, *Hamlet*, I.v

Chapter 26

Doubling Up: Tournaments, Part I

The largest growth area for poker over the past few years has undoubtedly been in the proliferation of tournaments. From the \$10,000 World Poker Tour™ events to \$1 or freeroll online megatournaments with thousands of players, tournaments have become the most visible and popular form of the game. There are a number of different structures tournaments can take, but they all fundamentally reduce to a poker contest with two key principles:

Tournament chips are irreplaceable - once all a player's tournament chips are gone, the player is eliminated from the tournament.

The first of these principles results in strategic adjustments related to skill - that is, a concept of passing up small edges that put one's whole stack at risk because of presumed situations that will occur later where larger edges will be available. Players who have an edge over the field should be less willing to commit all their chips for a thin edge because if they lose, they cannot replace those chips with additional cash as they could in a ring game. Conversely, players who have negative EV in the tournament should be more willing to gamble all-in on situations with slightly negative chip expectation.

This "skill adjustment," however, has been far overdone by some players; there is a limit to the type of situation that one can legitimately pass up and still actually have such an edge in the future. After all, as the blinds escalate in a tournament, it becomes more and more difficult to find hugely profitable situations. In addition, occasionally opponents will have big hands at the same time as we do. But this skill adjustment is real, and we can seek to quantify it approximately. We model this skill edge as a recursive effect and later in this chapter present a model that attempts to quantify the effect of a skill edge and provide insight into actual situations.

Rebuy tournaments obviously provide a temporary exception to this rule, although all rebuy tournaments eventually reduce to a tournament that satisfies it. We will revisit the topic of taking rebuys later in this chapter.

The prize pool of the tournament is distributed according to a payout structure that is coordinated with the order of bustout and not the raw accumulation of chips.

The second of these principles provides the main mathematically distinguishing characteristic of tournament play in the late stages when compared with cash games. In cash games, seeking to maximize cash EV means seeking to maximize chip EV; chips in cash games are cash. But in tournaments, this is not always the case. Obviously there is a strong correlation between chip position and prize equity; however, the decoupling of what we will call **chip equity** (the number of tournament chips a player holds) and **prize equity** (the EV of those chips in terms of the prize pool) can create significant adjustments.

A simple example will help to clarify the divergence of chip equity and prize equity in a tournament due to the payout structure.

Consider a 200-player limit holdem tournament of equally skilled players that has been reduced to 19 players. The tables are now going hand-for-hand. 10th-18th place pay \$250, and the payouts then escalate to \$6,000 for 1st place. The total prize pool is \$20,000. There are 300,000 tournament chips in play, and the blinds are 300-600.

Three players are extremely short-stacked. Player A has just 1 T100 chip remaining, but has the button at his table. Player B has T400 left and is under the gun. Player C has T500 left and is in also under the gun.

Early on in the tournament, the chips that the players hold are worth a share of the prize pool; A's 100 is worth $\frac{1}{3000}$ of the \$20,000 prize pool or about \$6.67. But in this situation, his T100 chip has been significantly increased in value. Consider that if he folds his next hand, both of his shortstacked opponents will be all-in for the blind. Assuming that each of them has a 40% chance of winning the hand against (presumably) the strongest hand in the field against them, A has an 84% chance of surviving into the money and getting \$250. So even discounting his chances of actually doubling up and building his chip into a stack, his T100 is worth a minimum of \$210. This is a far cry from the \$6.67 that his chip is worth *a priori*. The tournament payout structure has given his chip additional value. This value is of course taken away from the values of chips in big stacks; recall that collecting all the chips in the tournament is worth only 20% of the prize pool.

This is perhaps a rather extreme example but situations that are of the same character occur in just about every tournament. In these cases, survival is actually worth something - real money. This is generally the case for very small stacks near the "bubble". Their chair is actually worth more than their chips. In addition, at steps in the money, there is incentive for small stacks to play more tightly and attempt to move up, though this is only valid to a point. Some authors have (wrongly) taken this principle and attempted to extend it into the early-game, arguing that strong players should attempt to survive at almost all costs, avoiding confrontations in the early going and attempting to make the money and then making a push toward first.

We believe that this view sacrifices equity because skill-independent chip equity and prize equity are closely tied together in the early going. There simply aren't enough chances to get one's money in as a substantial favorite to justify the loss of tournament chip expectation from the resulting timid, exploitable play. This view also tails because even with the flatter structures now favored by many tournament directors (at this writing 20-25% to first place), a large percentage of the prize pool is concentrated in the top three places and the number of tournament chips held by a player is a strong predictor of that player's chance of finishing in the top three places.

We advocate, then, a strategy which is mindful of the necessary adjustments brought on by the payout structure along with those that a skilled player should make to ensure his repeated ability to apply his skill, but which recognizes that these strategy adjustments are fairly minor while the money is still reasonably far away. From an exploitive standpoint, as well, we seek to capitalize on our opponents' misunderstandings of these concepts, particularly near the money. This results in a lot more "just out of the money" finishes, but it also results in more high finishes and higher expectation overall.

Tournament models are generally most useful in attempting to express the *mapping* between chip equity and prize equity in tournaments. If we can map chip equity- to prize equity, then we can return to the familiar idea of maximizing EV from decisions. In this case, however, we do not maximize chip EV, but instead the corresponding cash EV. The first model we consider attempts to quantify the first effect (skill) on tournament decision-making in the early and middle stages of a tournament.

The Theory of Doubling Up

Many players will pass up even large edges in the early or middle stages of tournaments, believing (almost certainly wrongly) that even better edges will present themselves. But if this were actually true, then these players would have gigantic edges over the field in terms of equity.

Observation does not bear this out. Nonetheless there is an effect - we seek to measure and quantify the effect, and use it as a guide in making accurate decisions at the table.

To assess the value of different tournament stacks, we can use a variety of different models. These models are useful in different stages of the tournament. Some of these models, such as the ones we will introduce in the next section, are primarily useful when the players are near or in the money and the problem is allocating the prize payouts such that each player gets a fair share. The Theory of Doubling Up, by contrast, is intended to be applied in the opening and middle stages of the tournament as a guide to decision-making.

First, let us consider a winner-take-all tournament. In winner-take-all tournaments, excepting skill considerations, a player's equity in the tournament is proportional to his chip stack. We then further make an important assumption: *The chance of doubling a certain player's chip stack before that player busts is constant throughout the tournament.*

There are a number of effects that work on a player's chance of doubling up throughout the tournament. First, there is a mild selection effect. Stronger players tend to survive longer, and the longer a tournament goes, the stronger the overall field will be (on average). This causes all players' chance of doubling up to fall as the tournament progresses.

Secondly, the increasing blinds tend to compress all players' chances of doubling up toward 50% by reducing their latitude to apply their skill (or alternatively, their lack thereof). In the early going, skilled players will have plenty of opportunities to gain chips. Hence, doubling up early is better than doubling up later.

However, neither of these effects is particularly strong, and so we are comfortable taking the above assumption to be true in modeling tournament equities.

Based on this assumption, we can construct a model of a winner-take-all tournament with X players of equal skill who begin with equal stacks. Consider one single player. Let E be his chance of winning the tournament outright, C be his (constant) chance of doubling his stack before busting, and N be the number of times he must double his stack in order to win.

Then we have the simple relationship

$$E = C^N \tag{26.1}$$

For example, there are 128 players, he must double his stack seven times 1-2-4-8-16-32-64-128. Since each of these doubles is roughly independent, the chance of doubling seven times is C^7 .

With this formula, if we have any two of these quantities we can calculate the third. Let's consider the situation at the beginning of the tournament. Since all the players are equally skilled, we know that our player has an equity $E = 1/X$. We also know that in order to win the tournament, our player must increase this stack from its starting value to X times that amount. The number of doubles that this requires can be expressed as:

$$X = 2^N$$

So, for a 2-player tournament, N would obviously be 1, as one double would be required to win, while for a 128-player tournament, N is 7. More generally:

$$N = \log_2 X$$

Substituting:

$$1/X = C^{\log_2 X}$$

$$\log_2 1/X = (\log_2 X) \log_2 C$$

$$\log_2 C = -1$$

$$C = 0.5$$

Of course, this is what we would expect for a winner-take-all tournament where all players are equally skilled; that each player would have a 50% chance of doubling before busting from any stack size. It is easily verified by plugging in other values of A" (representing different stack sizes) that chip values are exactly linear for players of equal skill.

This methodology also generalizes to situations where the players are of unequal skill. For example, consider a hypothetical player A playing in a 100-player tournament where he has EV of 1 buyin net (2 buyins per tournament gross). We can calculate his C by plugging in his initial values:

$$E = C^N$$

$$2/100 = C^{\log_2 100}$$

$$0.02 = C^{6.643856}$$

$$C = 0.5550$$

So A's chance of doubling his current stack before busting is 55.5%. Now we can calculate his chance of winning the tournament from any stack size S (where S is the number of starting stacks he has) by plugging into the equity formula $E = C^N$. N, of course, is the number of doubles it will take to reach 100.

So for a stack size of 2 (double the initial stack), A has equity of $E = C^{\log_2 50}$, or 0.0360. Note that this value is more than the *a priori* value of the chips (which would of course be 0.02), but less than double A's chance of winning the tournament from the start.

What this model does, in effect, is provide us a means for evaluating marginal tournament equity decisions ***including the skill of the player making the decisions***. The qualitative judgment of "I have an edge over the field, so I should protect my chips" is quantified into an adjusted equity for any stack size that utilizes the assumptions that underlie the model. By doing this, this model provides a mapping of chip equity to tournament equity. To see how the model might work in practice, consider the following marginal situation from a no-limit holdem tournament.

The blinds are 75-150. Player B, with a stack of 3000 (250 players played with starting stacks of 1500), raises to 450 on the button with Qs Ts. The small blind folds and the big blind calls. The flop comes Ks 8c 2s. The blind checks and player B bets 500. The blind check-raises all-in (and has B covered). B estimates his equity' at 36% and is faced with the decision of calling 2050 to win 4025.

An analysis based purely on tournament chip values would indicate that calling is clear - (0.36) (6075) -2050 = 137 chips, for nearly a tenth of a buyin worth of profit. But let's assume that Player B had an edge of $3/4$ of a buyin over the field at the beginning of the tournament. Calculating C for Player B:

$$1.75/250 = C^{\log_2 250}$$

$$C = 0.5364$$

Now we can calculate the equity E of the three scenarios:

If B calls and wins, he will have 6075 chips, or 4.05 starting stacks.

$$E = 0.5364^{\log_2(250/4.05)} = 0.024603 \text{ buyins.}$$

If B calls and loses, he will have 0 chips and have 0 equity.

If B folds, he will have 2050 chips, or 1.36667 starting stacks.

$$E = 0.5364^{\log_2(250/1.36667)} = 0.009269 \text{ buyins.}$$

So B's equity from calling is $(0.024603)(0.36)$, or 0.008857 buyins. Comparing this to B's equity from folding, we find that calling is actually incorrect given B's skill edge over the field.

Another result that can be easily derived from the Theory of Doubling Up includes the "coinflip" problem - what chance of winning do you need in order to accept a coinflip for all your chips on the first hand? Since $E_0 = C^N$ at the start of the tournament, if you call and double up, your equity becomes $E_1 = C^{N-1}$. If you take a coinflip with chance of winning W , then your equity becomes WE_1 . If you decline, then your equity is E_0 . Setting these things equal, we get:

$$C^N = WC^{N-1}, \text{ or } W=C.$$

So you are indifferent to accepting a coinflip if your chance of winning the coinflip is equal to C . Yet consider that many players have a stated preference that they would decline a QQ vs AK (57-43) confrontation in the early tournament. By applying the Theory of Doubling Up, we can find out something about what these players believe their equity in, for example, a 250-player tournament to be.

$$E = C^N = (0.57)^8 = 0.01114$$

...or about 2.85 buyins per tournament. In order for it to be correct to decline a 57-43 confrontation with no dead money in the pot, one has to have nearly three times the average equity in the tournament. Our observations lead us to believe that having this win rate in a typical tournament field would place such a player among the best players in the world.

Another interesting effect of the Theory of Doubling Up is what it indicates about players with negative expectation. Of course our readers are all winning players, but nevertheless, the mean result of all players in a tournament is to lose the entry fee, so someone must be losing. The Theory of Doubling Up, then, indicates that losing players should be willing to commit their chips even in marginally bad situations. This is a result of the idea that losing players should encourage variance, as it is their best chance to win. The more that other players get a chance to apply their skill, the worse off the losing player will be. Hence, they should be willing to commit all their chips in zero-EV situations or even slightly bad ones.

The Theory of Doubling Up also has some important implications for multi-way confrontations, in poker, in general, the "third man in" to an all-in confrontation needs quite a strong hand. This effect is exacerbated by the tournament structure. Consider a situation on the flop where two players are all-in for a pot-sized bet and a player must consider calling all-in for the third bet. In a situation where chips are linear, he would need just a 1/4 chance of having the winner to make the call breakeven. In the three-way confrontation, however, the player needs a much stronger hand. If he wins, his chips will be worth four times as much as if he folds; that is, he will have

reduced his necessary doubles by 2. If his original chance of doubling up was, for example, 55%, he needs 30.25% of the pot just to break even equitywise, compared to 25% chipwise. This is because so often he will be broke and have no future opportunity to apply his skill.

In addition to the form we presented previously, we can show the Theory of Doubling Up formula in a different manner. We have the following from the previous discussion:

$$E = C^N$$

We can transform this using the rules of logarithms. Suppose we have a stack size S (which is the fraction of total chips in the tournament). Then the number of doubles we need is:

$$N = \log_2 (1/S) = -\log_2 S$$

$$E = C^{-\log_2 S}$$

$$\ln E = (-\log_2 S)(\ln C)$$

$$\ln E = (\ln S) (-\log_2 C)$$

$$\log_S E = -\log_2 C$$

$$E = S^{-\log_2 C}$$

We denote the quantity $-\log_2 C$ as β . making this formula:

$$E = S^\beta \tag{26.2}$$

Here we have a restatement of Equation 26.1; only now it is our stack raised to a constant power. We can go further with the ideas of the Theory of Doubling Up and simplify the calculations as well as incorporate ideas that affect our equity estimates. Previously, we have dealt only with winner-take-all tournaments. However, we can additionally incorporate a structured payout into the Theory of Doubling Up.

To understand this methodology, consider a winner-take-all match play tournament with 256 players. Then a single player's equity in the tournament is simply equal to $(256)(C^8)$, or his chance of doubling up seven times (that is, winning matches seven times) multiplied by the number of buyins he wins. We can rewrite this expression as:

$$(256)(C^8) = (2C)^8$$

In fact, in general, we can modify the $E = C^N$ equation to express equity in the total prize pool of X buyins as follows:

$$E_X = XC^N$$

$$E = (2C)^N$$

(since N is the number of doubles required)

Now consider the same tournament, except that the total prize pool is divided between first and second place equally - that is, 1st and 2nd are both 128 buyins. His equity here is actually $(128)(C^7)$, or $(2C)^7$.

How does this relate to the previous equity?

$$(2C)^8 > (2C)^7$$

$$C > 1/2$$

By splitting the first prize equally among the first two players, our player's equity is reduced if his chance of doubling is greater than $1/2$, while it is increased otherwise. This makes sense - basically, what's happening here is that instead of getting his chance of winning the last match times 256, he gets 128 instead. So his chance of winning the last match is simply replaced by $1/2$.

What this reflects is that in tournaments where not all the money is paid to first place, the opportunity to apply skill is reduced, because a player who acquires all the chips only acquires a fraction of the prize pool. So in the second case discussed above (where the prize pool is split between first and second place), our skillful player has less equity by one factor of $2C$. According to our equity estimates from the Theory of Doubling Up (which of course applies directly and not as a model in this case), the skill multiplier for this tournament is the same as if the tournament contained 128 players instead of 256.

We call this the *effective tournament size* - essentially what we're doing is quantifying the degree to which the structured payouts reduce the opportunity to apply skill. We can say that as a close approximation, a tournament of 256 players where the first and second prizes are equal has the same latitude for skill as a 128-player tournament where it's winner-take-all. We denote the effective tournament size (generally) as k , and define k more formally to be a value such that the following equation is satisfied for a given prize distribution:

$$(2C)^{\log_2 k} = E_N \tag{26.3}$$

where C is the chance of doubling up and E_N is the number of buyins the player wins in the tournament with the prize structure intact. Then k will be the size of the equivalent tournament (for the Theory of Doubling Up) with a winner-take-all structure.

Of course, tournaments don't often split the prize pool equally between first and second place, preferring to instead give a larger piece of the pool to first. Suppose that our hypothetical match play tournament instead has payouts of f of the prize pool to first place and $(1 - f)$ to second. Now our player's equity becomes:

$$\begin{aligned} &(\text{total prize pool})[(1^{\text{st}} \text{ money})(p(1^{\text{st}})) + (2^{\text{nd}} \text{ money})(p(2^{\text{nd}}))] \\ &(256)(f(C^8) + (1 - f)(C^7 - C^8)) \\ &(256)C^7(2fC + 1 - C - f) \end{aligned}$$

We know that for our hypothetical whining player, C will be a little bit higher than 0.5; for example 0.53. We'll use the symbol δ to denote the difference between C and 0.5,

$$\delta = C - 0.5$$

Then our equity formula simplifies to:

$$(256) C^7 (\delta (2f - 1) + 1/2)$$

This formula is fundamentally made of three parts:

- $(256) C^7$ is our hypothetical player's chance of making it to the finals times the entire prize pool.
- $\delta (2f - 1)$ is the "extra" equity he wins because of skill.
- $1/2$ is his base share of the prize pool.

Using this equity formula, we can find the effective tournament size for this tournament.

Let d equal $\log_2 k$, the number of doubles that would be required in a winner-take-all tournament of size k .

There exists some winner-take-all tournament such that:

$$(2C)^d = E_N$$

In this case, E_N is equal to our equity expression, so:

$$(2C)^d = (256)C^7 (\delta (2f - 1) + 1/2)$$

$$(2C)^{d-7} = 2 (\delta (2f - 1) + 1/2)$$

$$(d - 7)(\ln 2C) = \ln (2 (\delta (2f - 1) + 1/2))$$

$$(d - 7) = \ln (2 (\delta (2f - 1) + 1/2)) / \ln (2C)$$

$$d = \ln (2 (\delta (2f - 1) + 1/2)) / \ln (2C) + 7$$

$$k = 2^d$$

Suppose that f is $154/256$ and C is 0.52 . Then

$$d = \frac{\ln (2(0.02(2(\frac{154}{256} - 1) + 0.5))}{\ln 2C} + 7 \approx 7.206$$

$$k = 2^d = 147.68$$

Notice that k here is dependent on the value of C that we use. However, it turns out that the value of k is relatively insensitive to the chosen value of C as long as C falls within the usual parameters. We will utilize this fact in extending this discussion to payout structures with more paid places.

So our effective tournament field size here is actually about a hundred players smaller than the actual field size, simply as a result of the prize structure being divided. Note that this effective field size is used as the field size in the Theory of Doubling Up equations. Also, the E term must be multiplied by this number of buyins to find the proper equity.

For example, to find the equity of a player with $C = 0.53$ who has won his first four matches (thus needing 4 more doubles to win), we can use the following:

$$kE = C^{\log_2 (k/S)}$$

$$(147.68)(E) = (0.53)^{\log_2 (147.68/16)}$$

$$E = 19.287 \text{ buyins}$$

We can also calculate this directly as a check:

$$\text{Chance of making the finals: } 0.53^3 = 0.14888$$

$$\text{Chance of winning in the finals: } 0.53$$

$$(0.14888)[(0.53)(154) + (0.47)(102)] = 19.288 \text{ buyins}$$

Most tournaments are not match play, and most are not limited to two places paying. However, accepting some assumptions about the likelihood of brushing in different places, we can actually generalize this result to find the effective tournament field size (and by doing so, be able to generate the proper skill adjustments) for any payout structure.

Assumption:

The chance of a player's finishing in place N is equal to his chance of doubling up to a stack size of 1/N of the chips in play, less his chance of finishing in a higher place.

So a player's chance of finishing first is his chance of doubling up to all of the chips in the tournament. His chance of finishing second, then, is his chance of doubling up to half the chips in the tournament less his chance of finishing first. His chance of finishing third is his chance of doubling up to a third of the chips in the tournament, less this chance of finishing first or second, and so on. This series does hi fact sum to one: even though it's true that it's not required that a player double up to 1/N chips in order to finish Nth, this approximation is a reasonable one.

We can manipulate the Theory of Doubling Up and this assumption to find an effective tournament field size for any tournament. The following discussion is a little challenging, so we mark it off and you can feel free to skip it

A player's equity (in buyins) is given by the following formula:

$$E_N = \sum_{i=1}^x p_i v_i$$

where p_i is the probability of placing in i^{th} place, v_i is the prize payout for finishing in i^{th} place, and X is the total number of players in the tournament.

The probability of finishing in each place is equal to the probability of doubling up to the appropriate chips minus the chance of doubling to the appropriate chips for the next higher spot:

$$p_i = C^{\log_2(\frac{X}{i})} - C^{\log_2(\frac{X}{i-1})}$$

We define $N = \log_2 X$. This is the number of doubles it would take if this were a winner-take-all tournament. Substituting into the equation for the probabilities of finishing in each place:

$$p_i = C^{N - \log_2 i} - C^{N - \log_2 (i-1)}$$

This last equation only holds true if $i > 1$; at $i = 1$, $p_i = C^N$.

We now define a separate series of terms, related to p_i .

$$q_i = \sum_{i=1}^x p_i$$

So q_i is the "cumulative" probability that we will finish in the top i players; q_1 is the chance that we will finish first, q_2 is the chance that we will finish first or second, and q_N is the chance that we will finish in N^{th} place or higher.

Plugging in pi values, we have:

$$q_i = C^{N - \log_2 i}$$

Now consider the prize structure of our tournament. Calling the total prize pool P , where v_i is the payout for finishing i^{th} we have:

$$P = \sum_{i=1}^X v_i$$

Then w_i is the incremental increase in guaranteed money for finishing in i^{th} place instead of $[i + 1]^{\text{th}}$ place such that:

$$w_i = \sum_{i=1}^X v_i - v_{i+1}$$

with the special rule that $w_X = v_X$

We also have the following identity:

$$P = \sum_{i=1}^X i w_i$$

because

$$\begin{aligned} \sum (i w_i) &= \sum (i(v_i - v_{i+1})) \\ &= \sum (i v_i) - \sum (i v_{i+1}) \\ &= \sum (i v_i) - \sum ((i-1) v_i) \\ &= \sum (v_i) = P \end{aligned}$$

The total prize pool is equal to the money guaranteed for finishing i^{th} or higher times the number of players who finish there, summed for all i values.

We saw previously that

$$E_N = \sum_{i=1}^x p_i v_i$$

A player's equity is equal to the value of each finish times the probability of that finish, summed over all finishes.

However, we can also say the following:

$$E_N = \sum_{i=1}^x q_i w_i$$

A player's equity is equal to his chance of finishing in each place or higher multiplied by the

incremental money he is guaranteed by getting to that spot.

To show this, suppose that we have just three payouts (this approach generalizes to more payouts), v_1 , v_2 , and v_3 and the probabilities of finishing each of these spots are p_1 , p_2 , and p_3 , respectively.

We then have:

$$\begin{aligned} q_1 &= p_1 \\ q_2 &= p_1 + p_2 \\ q_3 &= p_1 + p_2 + p_3 \end{aligned}$$

and

$$\begin{aligned} g_1 &= v_1 - v_2 \\ g_2 &= v_2 - v_3 \\ g_3 &= v_3 \end{aligned}$$

We can now show that $q_1g_1 + q_2g_2 + q_3g_3$ is equivalent to our equity:

$$\begin{aligned} & q_1g_1 + q_2g_2 + q_3g_3 \\ & p_1(v_1 - v_2) + (p_1 + p_2)(v_2 - v_3) + (p_1 + p_2 + p_3)v_3 \\ & p_1v_1 + p_2v_2 + p_3v_3 \end{aligned}$$

With this, we have our equity for the game. What we are looking for is the k value such that the following equation is satisfied:

$$2N \sum_{i=1}^x q_i w_i = (2C)^D$$

where D is the number of doubles necessary to reach k players, or $\log_2 k$. This presumes that w_i is denominated in buyins and not fractions of the prize pool.

$$\begin{aligned} 2N \sum_{i=1}^x C^{N-\log_2 i} w_i &= (2C)^D \\ \sum_{i=1}^x (2C)^N C^{-\log_2 i} w_i &= (2C)^D \end{aligned}$$

$$\sum_{i=1}^x C^{-\log_2 i} w_i = (2C)^{D-N}$$

$$\sum_{i=1}^x i(2C)^{-\log_2 i} w_i = (2C)^{D-N}$$

Next we make use of the idea that k is insensitive to the value of C chosen. Suppose we choose C very close to 0.5. In keeping with convention, we use ε to denote twice the distance between 0.5 and C (The value 0.5 exactly is a discontinuity because it involves dividing by zero).

$$\sum_{i=1}^x i(1 + \varepsilon)^{-\log_2 i} w_i = (2C)^{D-N}$$

Assuming ε is close enough to 0.5, we can approximate the exponential of expansion as follows:

$$(1 + \varepsilon)^x = 1 + x\varepsilon.$$

$$\sum_{i=1}^x iw_i(1 - (\log_2 i)\varepsilon) = (1 - (N - D)\varepsilon)$$

The iw_i term simply sums to 1, and we are left with:

$$\sum_{i=1}^x iw_i(\log_2 i) = (N - D)$$

$N - D$ can be thought of as the number of doubles that are "missing" from the tournament when the prize structure is incorporated. To find the actual effective tournament field size, we can divide X by 2^{N-D}

This much simpler formula can be used to find the effective prize structure for any tournament. We can test it on our previously examined tournament.

$$g_1 = 154/256$$

$$g_2 = 102/256$$

$$204/256 = N - D$$

$$k = 256 / 2^{204/256} = 147.35 \text{ where } D = \log_2 k$$

The slight variance here is caused by the fact that when we first looked at this tournament, we chose a C of 0.52.

Key Concepts

- The Theory of Doubling Up allows us to evaluate equity decisions in the early to middle portions of a tournament by using a constant "chance of doubling up" and finding the number of doubles required to get all the chips. Using these two measures, we can effectively model tournament skill and get a sense for the marginal values of chips.
- The $E=C^N$ form of the Theory of Doubling Up can be further refined to understand the effect that a flat payout structure has on skill. When money is given to a place other than 1st, the effect of skill is diminished.
- We can solve for an effective tournament size (see Equation 26.3), which is the size of a winner-take-all tournament where the skill adjustments are the same as our structured payout tournament.

Chapter 27

Chips Aren't Cash: Tournaments, Part II

A Survey of Equity Formulas

The Theory of Doubling Up is not designed to capture situations near the money or at the final table. Yet we commonly want to have an explicit mapping of our tournament chip stack to its cash value is near the conclusion of a tournament, either because there is the possibility of dealmaking (when this is allowed) or to account for the payout structure's distortion of the chip to cash mapping so that we can make more accurate decisions.

In fact, it is worth noting that the two main sources of this distortion (skill and the payout structure) each apply at different moments throughout the tournament. At the beginning of the tournament, the payout structure is of relatively small importance, but skill is at its maximum effect. By contrast, at the end of a tournament, the blinds are often high and so skill is fairly reduced (although the specific skill of playing shorthanded might be emphasized), while the payout structure's effect on the mapping is maximized.

In any event, there is no easy answer to the question of mapping chip stacks to prize equity, even for players of equal skill. Various approaches have been proposed to generally solve this problem. In this section, we present a number of different models that have been proposed and make some comments about the biases which are present in each. We begin by defining some terms.

Each of the i players remaining in a tournament has a chip stack S_i . The **total number of chips remaining** in the tournament \mathcal{T} is then:

$$\mathcal{T} = \sum S_i$$

A player's fraction of the remaining chips is:

$$s_i = \frac{S_i}{\mathcal{T}}$$

Obviously the sum of all s_i values will be 1.

Additionally, we define v_j as the prize payout for finishing in j^{th} place. The sum of all of the prizes is the total prize pool V . The **prize pool remaining to be paid out** when there are n players left V_n is as follows:

$$V_n = \sum_{k=1}^n v_k$$

We can also consider the **incremental prize pool**, the amount left to be decided by players finishing in places other than n^{th} place: Letting w_j be the incremental prize pool for j^{th} place, and W_n the cumulative prize pool through n^{th} place:

$$w_j = v_j - v_n$$

$$W_n = \sum_1^n w_j = V_n - nv_n$$

Since the v_n value is a constant, the decision processes within the tournament for the **gross prize pool payouts** and the **incremental prize pool payouts** are identical.

A tournament is considered **in the money** if $v_n > 0$; that is, the next player to bust receives a payout. This occurs at the precise n value where $v_n = 0$ but $v_n + I > 0$, where the tournament is said to be **on the bubble**.

Now we define $p_{i,j}$ as the probability that player i finishes in place j . Then we have a basic prize equity formula:

$$\langle X_i \rangle = \sum_{j=1,N} p_{i,j} v_j$$

Most equity formulas that have been presented assume that the chance of finishing in first place $p_{i,1}$ is approximately equal to the fraction of chips s_i . This is true if the players are of equal skill; however, in situations where the players are of unequal skill, methods such as the Theory of Doubling Up could be used to estimate the chance of finishing first. Instead of using chip stacks as a proxy for the chance of finishing first, we will simply define x_i as the probability of player i finishing first. This can be substituted with any probability-generating method you wish.

Equity formulas typically take the chance of finishing in first place as a given, and provide a simplified methodology for assigning probabilities to the other places.

The Proportional Chip Count Formula

The first model for assigning probabilities to the other places that we present is the frequently employed "chip count" method, which can be expressed as follows:

$$\langle X_i \rangle = v_n + s_i W_n \tag{27.1}$$

This model consists of two parts:

- First, all players receive the guaranteed value of the finish that they are guaranteed.
- Then, the remaining prize pool is divided amongst the players according to their chip count.

For example, suppose the stacks are 25,000, 20,000, 10,000, and 5,000, and the payouts are \$2,000, \$1,400, \$900, and \$700. Then the "chip count" method using Equation 27.1 would allocate the remaining prize pool of \$5,000 as follows:

\$2,800 is distributed equally to each player as his guaranteed payout.

The remaining \$2,600 divided to each player according to his chip count.

Player 1: $\langle X_1 \rangle = 700 + (5/12)(2200) = \$1,616.67$

Player 2: $\langle X_2 \rangle = 700 + (1/3)(2200) = \$1,433.33$

Player 3: $\langle X_3 \rangle = 700 + (1/6)(2200) = \$1,066.67$

Player 4: $\langle X_4 \rangle = 700 + (1/12)(2200) = \883.33

At a glance, perhaps these numbers look reasonable. But the flaw of this method can quickly be seen if we change the stacks so that Player 1 has 48,000 and the other players 4,000 each:

$$\text{Player 1: } \langle X_1 \rangle = 700 + (4/5)(2200) = \$2,460$$

$$\text{Players 2-4: } \langle X_i \rangle = 700 + (1/15)(2200) = \$846.67$$

Clearly this isn't right. Player 1 seems to be entitled to receive more than first place money! This is a systemic flaw in this method; it tends to distribute too much of the prize pool to players on the basis of their chips, and too little to players on the basis of their chairs. The more equal the stacks are, the more accurate this formula is. However, this is true of all models - any model that does not tend to equalize the payouts as the x_i values equalize cannot be right!

The reason for this flaw is that this model is equivalent to saying that $p_{i,j} = s_i$ for all values of j ! Well, we can see how this is wrong. Suppose our player is the chip leader (equally skilled). Then his chance of 1st place is larger than that of all the other players. But his chance of 2nd, 3rd, 4th and so on down to n^{th} is the same. So his total chance of finishing in any place in the tournament is ns_i . If s_i is larger than $1/n$, then his total chance of finishing in some place in the tournament will be higher than 1. This is obviously incorrect.

One potential modification we can make to this formula is to constrain the finish data to sum to one by making the following corrections:

$$m = \left\lfloor \frac{1}{x_i} \right\rfloor$$

$$p_{i,j \leq m} = x_i$$

$$p_{i,m+1} = 1 - m x_i$$

Here, the brackets in the expression for m are the **floor function**, which returns the largest integer less than or equal to the bracketed expression.

Hence, if $x_i > 0.5$, then $m = 1$, and we have

$$p_{i,1} = x_i$$

$$p_{i,2} = 1 - x_i$$

This approximation is a more reasonable one when the stacks are very large. Consider the revised probabilities for the example we just considered:

For Player 1, who has $5/12$ of the chips, $m = 2$.

Then

$$p_{i,1} = 5/12$$

$$p_{i,2} = 5/12$$

$$p_{i,3} = 1/16$$

and l's revised equity is:

$$\langle X_1 \rangle = 700 + (5/12)(1300) + (5/12)(700) + (1/6)(200) = \$1,566.67$$

Unfortunately, this modification causes the formula to no longer sum to the total prize pool. Therefore, it is primarily useful to establish a slightly lower upper bound for the equities of the chip leaders. It is still fairly biased on behalf of the tall stacks, however.

The Landrum-Burns Formula

The second model we will consider was originally proposed by Stephen Landrum and Jazbo Burns on rec.gambling.poker. This model uses the following methodology:

We calculate each player's chance of finishing first.

We then assume that when he does not finish first, he finishes equally often in the remaining places.

$$\langle X_i \rangle = p_{i,1}v_1 + \sum_{k=2}^n \frac{(1-p_{i,1})}{n-1} v_k \quad (27.2)$$

Using this formula, we can recalculate our previous example where the stacks are 25,000, 20,000, 10,000, and 5,000, and the payouts are \$2,000, \$1,400, \$900, and \$700. Then the "Landrum-Burns" method calculates the equities as follows:

$$\text{Player 1: } (\frac{5}{12})(2200) + (\frac{7}{36})(1400) + (\frac{7}{36})(900) + (\frac{7}{36})(700) = \$1,416.67$$

$$\text{Player 2: } (\frac{1}{3})(2200) + (\frac{2}{9})(1400) + (\frac{2}{9})(900) + (\frac{2}{9})(700) = \$1,333.33$$

$$\text{Player 3: } (\frac{1}{6})(2200) + (\frac{5}{18})(1400) + (\frac{5}{18})(900) + (\frac{5}{18})(700) = \$1,066.67$$

$$\text{Player 4: } (\frac{1}{12})(2200) + (\frac{11}{36})(1400) + (\frac{11}{36})(900) + (\frac{11}{36})(700) = \$1,083.33$$

But we can equally expose the error of this method by considering the second example, where Player 1 has 48,000 and the other players 4,000:

$$\text{Player 1: } (\frac{4}{5})(2000) + (\frac{1}{15})(1400) + (\frac{1}{15})(900) + (\frac{1}{15})(700) = \$1,800$$

$$\text{Players 2-4: } (\frac{1}{15})(2000) + (\frac{14}{45})(1400) + (\frac{14}{45})(900) + (\frac{14}{45})(700) = \$1,066.67$$

The \$1,800 number is at least "reasonable," since it doesn't exceed the payout for first (as occurred with the "chip count" method). However, calculating it directly exposes a difficulty with the method. Player 1 may protest that he is not as likely to finish fourth as second, should he fail to win. In fact, he will finish fourth very infrequently, if the blinds are fairly high. So the "fair" price should be adjusted up from the Landrum-Burns value. Likewise, suppose one player has a tiny stack compared to his opponents. He is not as likely to finish second as he is to finish fourth.

This is the bias of the Landrum-Burns method - it overvalues small stacks at the expense of big stacks. This occurs because large stacks will not bust out immediately as often as they finish second. The proportional chip count method and Landrum-Burns together provide us a lower and upper bound to the plausible values; the "true" fair value lies somewhere between them.

In *Gambling Theory and Other Topics*, Mason Malmuth introduced two additional equity formulas.

The Malmuth-Harville Formula

This formula is parallel to the Harville horse-racing formula.

Essentially, what this formula does is the following:

- The chance of finishing first for each player is x_i ,
- We can use the following formula to calculate the chance of finishing in 2nd place given that player k has finished first:

$$p(X_{i,2} | X_{k,1}) = \frac{x_i}{1 - x_k}$$

Then our probability of finishing second is:

$$p_{i,2} = p(X_{i,2})$$

$$p_{i,2} = \sum_{k \neq i} p(X_{i,2} | X_{k,1}) p(X_{k,1})$$

$$p_{i,2} = \sum_{i \neq k} \frac{x_i x_k}{1 - x_k}$$

This formula was originally developed for betting favorites in show or place pools in horse racing. However, this formula actually applies more strongly to tournament poker because of an idea pointed out to us by Jeff Yass:

In horse racing, if a horse is a prohibitive favorite, where x is closer to one, a substantial amount of his non-wins will be due to a problem or event which will also prevent him from finishing 2nd, such as an injury or disqualification. It is unreasonable, for example, to believe that a horse with an 80% of winning will still have a higher than 80% of finishing second if he fails to win, and even more unreasonable to believe he has a >99% of finishing in the top three.

But in poker tournaments, this effect is minimized. If a tall stack does not win he will still be very likely to finish second, and very- unlikely to finish in the last remaining spot.

We can generalize this rule as well:

$$p(X_{i,j} | X_{1,1} \wedge X_{2,2} \wedge \dots \wedge X_{j-1,j-1}) = \frac{x_i}{1 - x_1 - \dots - x_{j-1}} \quad (27.3)$$

That is, the probability of finishing in the j^{th} place is simply the ratio of winning probability for the player to the win probabilities of the remaining players. In most cases, this formula is a more accurate representation of equity than either the proportional chip formula or Landrum-Burns.

One limitation of this method is that if there are many players remaining it is fairly difficult to calculate at the table because the number of permutations of finishes; as such, it is primarily useful at understanding situations away from the table or for analyzing situations that recur frequently, such as sit-and-go tournament equities or the like.

The Malmuth-Weitzman Formula

A second model proposed by Malmuth in *Gambling Theory and Other Topics* was developed by Mark Weitzman, which we will call the Malmuth-Weitzman formula.

In that text, Malmuth seems to prefer this formula over the Malmuth-Harville formula, although he claims that this formula requires the solution of a system of linear equations for each situation. We believe that it is slightly less accurate than the other formula but the solution to the linear equations is unnecessary, as we shall see.

The main idea of this formula is that when a player busts out of the tournament, his chips are distributed equally (on average) among his opponents.

$$x_i = P(X_{i,1}) = \sum_{k \neq i} P(X_{i,1} | X_{k,n}) P(X_{k,n}) = \sum_{k \neq i} (x_i + \frac{x_k}{n-1}) p_{k,n}$$

This results in a series of equations that we have to solve. For each player we have a value $p_{i,n}$ which is the probability of busting out next and a summation for all players of the equities from distributing their chips equally among all players.

For example, if we have four players, then we have the following equations (we use p_k here as shorthanded for $p_{k,n}$). All the probabilities shown are of busting out next.

$$\begin{aligned}x_1 &= \left(x_1 + \frac{x_2}{3}\right)p_2 + \left(x_1 + \frac{x_3}{3}\right)p_3 + \left(x_1 + \frac{x_4}{3}\right)p_4 \\x_2 &= \left(x_2 + \frac{x_1}{3}\right)p_1 + \left(x_2 + \frac{x_3}{3}\right)p_3 + \left(x_2 + \frac{x_4}{3}\right)p_4 \\x_3 &= \left(x_3 + \frac{x_1}{3}\right)p_1 + \left(x_3 + \frac{x_2}{3}\right)p_2 + \left(x_3 + \frac{x_4}{3}\right)p_4 \\x_4 &= \left(x_1 + \frac{x_1}{3}\right)p_1 + \left(x_4 + \frac{x_2}{3}\right)p_2 + \left(x_4 + \frac{x_3}{3}\right)p_3\end{aligned}$$

With four equations and four unknowns, we can solve this straightforwardly using algebra. However, we can do better than this. These equations are part of a family of equations called *Vandermonde* equations. We will simply provide the solution to these equations and show that it is correct. We claim that the above matrix implies that the probability of busting out next is inversely proportional to the number of chips (or win probability) left. If we define b as the sum of the inverses of all win probabilities:

$$b = \sum \frac{1}{x_i}$$

Then

$$p_{k,n} = \frac{1/x_k}{b} \tag{27.4}$$

We can substitute this back into the above formula:

$$x_i = \sum_{k \neq i} \left(x_i + \frac{x_k}{n-1}\right)p_{k,n}$$

$$x_i = \frac{1}{b} \sum_{k \neq i} \left(\frac{x_i}{x_k} + \frac{1}{n-1}\right)$$

$$x_i = \frac{1}{b} \sum_{k \neq i} \left(\frac{x_i}{x_k} + 1\right)$$

$$x_i = \frac{1}{b} \left[x_i \sum_{k \neq i} \frac{1}{x_k} + 1\right]$$

The summation of $1/x_k$, for all k not equal to i , rather it is $b - 1/x_i$.

So we have:

$$x_i = \frac{1}{b} (bx_i - 1 + 1) = x_i$$

The system of linear equations generated by the assumption of this model implies that each player's chance of going bust next is proportional to the inverse of his chip count or win

probability.

The Malmuth-Weitzman is a reasonable formula as well and is probably near the edge of some players' ability to calculate at the table, which increases its usefulness. One objection to this formula is that it's a bit sticky when it comes to situations where multiple tables are left in play. If two players are quite short stacked at different tables, it seems counter-intuitive that the player with twice the chips has half the chance of busting. Also, it's a little unclear why players at different tables receive an equal share of the busting chips. However, by and large, this is our favorite formula for assessing tournaments that are in or near the money. When assessing deals at the final table, we use the Malmuth-Weitzman as our baseline equal skill formula.

The important thing to understand from all this is that there are a number of different ways to map chip counts to tournament equities. Because these are all models, they inherently have some sort of bias. Understanding the common methods and the biases they have can lead to an enhanced ability to make in-tournament decisions as well as to make more favorable deals, when the circumstances arise.

Tournament Results and Confidence Intervals

In Part I, we discussed confidence intervals as they applied to win rates. For ring games, confidence intervals can be calculated in a typical fashion. This is because ring games are essentially an aggregation of hand results, which are in turn subject to the law of large numbers. However, for tournaments, constructing confidence intervals is not so easy. Indeed, constructing confidence intervals in the usual way doesn't yield the most accurate analysis.

Consider the case of a winner-take-all tournament with 1000 players. Suppose that we assume that the population's win rates will all fall between -1 and +1 buyin per tournament. Say we play the tournament once. The player who wins the tournament will have experienced a 30 standard deviation event. However, each time the tournament is played, someone is going to win. Even if we play the tournament ten times, it doesn't particularly help matters.

Another difference is that unlike ring games, where estimates of variance quickly converge because of the distribution they follow (chi-squared), we don't normally have good estimates of variance for tournaments. Variance is also highly correlated to win rate because, in tournaments, all large variance events are positive, while negative events are smaller variance (-1 buyin, typically).

The real problem is that for tournaments, it takes many more samples before the distribution converges to a normal distribution.

One problem is skew, which we briefly discussed in Chapter 22. This is defined as:

$$S = \frac{\langle (X - \mu)^3 \rangle}{\sigma^3}$$

Skew for a n-player winner-take-all tourney is around \sqrt{n} , while the skew of a normal distribution is zero. We know the skew of the sum of n identical independent events is the skew of one event divided by \sqrt{n} . As we take a larger and larger sample in trials, the skew of the sample will eventually converge on zero. This is one of the ways in which sampling distributions eventually converge on the normal distribution. But here we need more than n samples for the skew to be within 1.

But we can work on the problem nonetheless. The expectation of our gross profit for the

tournament is:

$$\langle X \rangle = \sum p_i V_i$$

Where V_i is the prize for finishing in i^{th} place. Let V be the total prize pool and we can define v_i to be the fraction of the prize pool awarded to the i^{th} place finisher.

$$v_i = \frac{V_i}{V}$$

Assuming no rebuys and add-ons, we define the *size* of the tournament as the prize pool divided by the number of players n .

$$s = \frac{V}{n}$$

We additionally have the *cost* of the tournament b_c , which is the total amount each player must pay to play the tournament. The difference between these two is the *entry fee*, which can be either positive (when the players pay a fee to the house for spreading the tournament) or negative (such as when the house adds money to the tournament for a guarantee or other promotion). Letting ε_b denote the entry fee:

$$\varepsilon_b = b_c - s$$

For example, suppose we had a tournament that was \$100 + \$10, with a \$50,000 guarantee from the casino. So \$100 from each player's \$110 entry goes to the prize pool; the remainder goes to the house. If the prize pool is short of \$50,000, the casino makes up the difference. If there are 300 players, then we have:

$$s = \$166.67$$

$$b_c = \$110$$

$$\varepsilon_b = \$56.67$$

Our net expectation w for entering a tournament is then our expectation of gross profit $\langle X \rangle$ less the cost of the tournament b_c , or:

$$w = \langle X \rangle - b_c$$

$$w = \sum p_i V_i - s - \varepsilon_b$$

This expression effectively has two parts. The first part,

$$w_0 = \sum p_i V_i - s$$

should be related to the player's poker and tournament skill, while the second part, ε_b , is be independent of poker skills, except possibly game selection.

We define a *normalized win rate*, ω , as:

$$\omega \equiv \frac{\omega_0}{s}$$

$$\begin{aligned}
 &= \sum p_i \frac{V_i}{s} - 1 \\
 &= \sum (np_i) \frac{V_i}{V} - 1 \\
 &= \sum (np_i - 1)v_i
 \end{aligned}$$

This normalized win rate is basically the number of buyins (where buyins is normalized to tournament size) that we can expect to net per tournament.

The statistical variable we want to look at is the *normalized win*, Z , such that:

$$Z = \frac{X - s}{s}$$

This has the special property that:

$$\langle Z \rangle = \omega$$

It turns out Z is a statistical variable we are interested in for a number of reasons. The idea here is to get a handle on the behavior of the Z variable, and use that information to construct appropriate confidence intervals.

Note that typically, the mode value for Z is -1 .

One possible assumption that would yield an appropriate Z variable is:

$$p_i = \frac{1 + \omega}{n}$$

This assumes that the player's win rate is an effective multiplier on his chance of finishing in each paid place. We begin with the player finishing in each position with equal likelihood. Then we move ω/n finishes from "out of the money" finishes to each payoff spot. This is at odds with the results of the Theory of Doubling Up, which suggests that a skilled player's chance of finishing first is higher than his chance of finishing second, third, and so on. However, we feel this approximation is fairly reasonable to use anyway. Recall from Chapter 25 that:

$$\begin{aligned}
 \langle x^2 \rangle &= \mu^2 + \sigma^2 \\
 \sigma^2 &= \langle x^2 \rangle - \mu^2
 \end{aligned}$$

For this distribution,

$$\begin{aligned}
 x &= X - b_c \\
 \mu &= \langle X - b_c \rangle
 \end{aligned}$$

The variance of a single tournament, then, is:

$$\begin{aligned}
 \sigma_x^2 &= \langle (X - b_c)^2 \rangle - \langle X - b_c \rangle^2 \\
 \sigma_x^2 &= \langle X^2 \rangle - \langle X \rangle^2
 \end{aligned}$$

$$\sigma_x^2 = \sum p_i V_i^2 - (\omega + b)^2$$

$$\sigma_x^2 = \sum \left(\frac{1 + \omega}{n} \right) (V v_i)^2 - (1 + \omega)^2 b^2$$

$$\sigma_x^2 = \frac{(1 + \omega)}{n} (nb)^2 \sum v_i^2 - (1 + \omega)^2 b^2$$

$$\sigma_x^2 = \left[n \left(\sum v_i^2 \right) (1 + \omega) - (1 + \omega)^2 \right] b^2$$

Then we have the variance of Z as:

$$\sigma_z^2 = \langle Z^2 \rangle - \langle Z \rangle^2 = \frac{\sigma_x^2}{b^2} = \left(n \sum v_i^2 - 1 - \omega \right) (1 + \omega)$$

Hence, the variance of Z is roughly proportional to $1 + \omega$ for reasonable values of ω . If we examine the value of:

$$k = n \sum v_i^2 - 1 \tag{27.5}$$

We will find that this expression is close to the effective number of players that we found when we examined the Theory of Doubling Up. We will call k the ***tournament variance multiple***. Using the variance formula above for the variance of Z is a better estimate of the variance of a tournament than the observed variance of a sample for any sample size that isn't enormous.

Now consider a series of independent tournaments. Suppose that we have a series of normalized outcomes (Z-values) for m tournaments.

The average of these normalized outcomes is \bar{Z} such that:

$$\bar{Z} = \frac{\sum_{j=1}^m Z_j}{m}$$

Since each Z_j value is independent of the number of tournaments played, we have the following formula for the variance of a sample of m Z-values:

$$\sigma^2 = \frac{\sum \sigma_{z_j}^2}{m^2} = \frac{\sum (k_j - \omega)(1 + \omega)}{m^2} = \frac{(\bar{k} - \omega)(1 + \omega)}{m} \tag{27.6}$$

In all normal circumstances, ω will be much smaller than k . So we can very roughly approximate the above expression with k/m . Then to make the variance of a sample of size m equal to 1, we must play about the same amount of tournaments as the average variance multiplier for the tournaments. If the variance of the sample is equal to 1, then we have an approximately 68% chance of our observed win rate ω being within 1 buyin of the population mean (neglecting Bayesian implications).

We can call Equation 27.6 the ***win-rate-based variance*** formula, because it uses our observed win rate and the observed structure of the tournament to calculate the variance of the tournament.

This formula has the flaw that in truth, having a win rate of ω does not normally imply that you enhance your chance of finishing in every paying spot by a factor of ω ; in fact, some players'

finishes are more top-heavy than others, with more firsts than seconds and thirds, while others are more weighted toward lower finishes. Another method of calculating this would be to use the Theory of Doubling Up to predict the rate of finish in each spot and calculate the variance that way.

The reason that we do not use the observed variance is that unless the sample size is very large, these formulas will give us a better estimate. If m were large however, we could utilize another way of calculating variance, which would be to use the observed win rate $\widehat{\omega}$ and observed variance $\widehat{\sigma}^2$ to calculate k .

$$k = \widehat{\omega} + \frac{m\widehat{\sigma}^2}{1 + \widehat{\omega}}$$

We then use this value of k in the win-rate-based variance formula. This technique is likely more accurate when m is very large, because the convergence of the observed variance to the true variance occurs more rapidly than the convergence of the sampling distribution to the normal distribution.

The question we are trying to answer with confidence intervals is the following:

If our normalized win rate is ω , how lucky or unlucky is our observed result?

We cannot strictly use the normal distribution as our guide here because the sampling distribution has not yet converged on the normal distribution. However, we can use the normal distribution to map a *z-score* (see Equation 2.6) to a confidence level (if, for example, being more than 2 standard deviations away fails the 95% confidence test) if $\omega > \widehat{\omega}$.

$$z = \frac{\widehat{\omega} - \omega}{\sigma(\omega)} < -2$$

If this expression is true, we can conclude that ω fails the test, because we know that the distribution is skew-positive and the probability of having a negative event (having an outcome more than two sigmas below expectation) is actually less than the normal distribution would indicate. So we can exclude some very high win rates based on observed data.

However, for $\omega > \widehat{\omega}$, we have to be more careful. The normal distribution is inappropriate for use here because in our sampling distribution, due to the skew, it is more likely that our observed outcome is a positive event than the normal distribution would indicate.

We can, however, use a statistical idea called Tchebyshev's (or Chebyshev's) Rule:

$$P\left(\frac{\widehat{\omega} - \omega}{\sigma(\omega)} = z > a\right) < \frac{1}{a^2}$$

This says that the probability that our outcome is actually in excess of four standard deviations on one side of the mean is always less than $1/16$. This may seem like a particularly blunt tool, after all, the normal distribution has a similar value of 0.00003. However, this rule has the excellent property of applying to any distribution, no matter how skew.

Let's take a moment now to review all of these ideas in the context of a hypothetical player, player T. Player T is a tournament player, who plays four \$50 + \$5 tournaments per day. Two of

the tournaments attract about 100 players (these are the morning tournaments) and the other tournaments attract about 400 players (these are the evening tournaments).

The morning tournaments pay out according to the following schedule:

Place	1	2	3	4	5	6	7	8	9
Pct	30%	20%	12%	10%	8%	6%	5.5%	4.5%	4%

The evening tournaments pay out as follows:

Place	1	2	3	4	5	6	7	8	9	10-18	19-27
Pct	27.5%	17.5%	11.6%	8%	6%	4.5%	3.5%	2.6%	1.7%	1.2%	0.7%

Player T has played these tournaments for 200 days. As a result, he has 400 morning tournament data points, and 400 evening tournament data points.

The following is a list of his finishes in each type of tournament over that span:

Morning tournaments:

Place	Frequency	Gross Buyins	Net Buyins	Total net
1st	11	30	29	319
2nd	7	20	19	133
3rd	5	12	11	55
4th	3	10	9	27
5th	8	8	7	56
6th	4	6	5	20
7th	7	5.5	4.5	31.5
8th	2	4.5	3.5	7
9th	5	4	3	15
No cash	348	0	-1	-348
Total	400			315.5

So Player T's overall cash rate in morning tournaments is $\frac{52}{400}$, or 13%. His overall win is 315.5 buyins, a rate of about 0.79 buyins per tournament.

Evening tournaments:

Place	Frequency	Gross Buyins	Net Buyins	Total net
1st	4	110	109	436
2nd	2	70	69	138
3rd	2	46.4	45.4	90.8
4th	0	32	31	0
5th	2	24	23	46
6th	4	18	17	68
7th	1	14	13	13
8th	0	10.4	9.4	0
9th	2	6.8	5.8	11.6
10th - 18th	14	4.8	3.8	53.2

19th - 27 th	20	2.8	1.8	36
No cash	349	0	-1	-349
Total	400			543.6

And in the evening tournaments, Player T has cashed about 13% of the time again, with an overall win of 543.6 buyins, a rate of about 1.36 buyins per tournament.

Taken together, T has played a total of 800 tournaments, with an aggregate win rate of 1.07 buyins per tournament. We can make an inventory of the statistics we've discussed here:

$$\begin{aligned} \omega &= 1.07 \\ m &= 800 \\ \bar{z} &= 1.07 \\ k &= n \sum v_i^2 - 1 \end{aligned}$$

For the morning tournaments, we have:

$$k = (0.30)^2 + (0.20)^2 + \dots + (0.04)^2 = 16.105$$

And for the evening tournaments, we have:

$$k = (0.275)^2 + (0.175)^2 + \dots + (0.007)^2 = 53.6456$$

Then our average k value is:

$$\bar{k} = 34.875$$

and our variance for an 800 tournament sample consisting of half morning and half evening tournaments is approximately:

$$\sigma^2 = \frac{(\bar{k} - \omega)(1 + \omega)}{m} = \frac{(34.875 - 1.07)(1 + 1.75)}{800} = 0.0876$$

or a standard deviation of 0.296 buyins per sample (of 800 tournaments).

The observed sample variance (from the sample actually observed) is about 0.3468 buyins. This difference occurs because the sample is slightly more concentrated in first place finishes. You might object that Player T's apparent tendency to finish first more often than other places skews his underlying distribution. This might be true in practice, although the effect would not be terribly important. However, since we wrote the program that randomly generated the above results, we know that Player T's chance of finishing first was the same as his chance of finishing in any other money position. This is yet another case of the possibility of being fooled by random noise.

Applying the normal distribution to the observed win rate and standard deviation would result in a 95% confidence interval of [0.3802, 1.7675]. This is definitely inaccurate on the high side. Because the distribution is skew-positive, the chance that the true win rate is higher than 1.7675 and T had a very negative event is actually lower than the normal distribution would indicate. On the low side, it is entirely inappropriate to use the normal distribution, because the chance that the observation is actually a positive event from a lower win rate is much higher.

Utilizing the win rate-based variance calculation, T actually has smaller variance. We find that two standard deviations above would be 1.666 buyins per tournament. This is still slightly inaccurate. In fact, at this point, the sample would fail a hypothesis test for 95% accuracy. The true bound is actually lower than this, but only modestly. On the other side, however, we can use Tchebyshev's Rule and conclude that the chance that the true distribution is lower than 2 standard deviations below the observed mean (or below 0.4818) is no greater than $\frac{1}{4}$.

Key Concepts

- A number of equity formulas have been proposed to map chip equities to tournament equities in the late stages of tournaments.
- The Landrum-Burns and Proportional Chip Count methods have clear flaws but are easy to calculate at the table. Since they are flawed in opposite directions, evaluating both can be useful to identify a range of values for the true equity.
- The Malmuth-Harville and Malmuth-Weitznan formulas are more accurate, but harder to implement at the table.
- Tournament equity confidence intervals are an interesting topic because of their significant skew and also because the variance per tournament is highly correlated to the win rate. We can use these facts to obtain more accurate confidence intervals even without strongly significant data (which is often difficult to gather because it takes so many trials).

Chapter 28

Poker's Still Poker: Tournaments, Part III

Game Theoretic Tournament Concepts

From a game theory standpoint, tournaments have one important quality that affects play-beyond the usual poker considerations, even among equally skilled players. In a tournament, every player has some stake in every hand, because every player gains equity when other players are eliminated. Because of this, hands played headsup are not constant-sum as they are in other forms of poker. When an all-in confrontation occurs, both players lose some equity. This amount can range from negligible (as would be the case on the first hand of the tournament) to extremely significant (as on the bubble in a supersatellite). One effect of this is to create an advantage to the first player to enter a pot, but particularly to the first player to enter a pot for all his chips. Because of the escalating blinds of tournaments, these opportunities occur rather frequently in no-limit holdem.

In a normal full cash game situation, players behind an all-in raiser must hold a hand that has enough equity to justify calling against the raiser's distribution of hands. This means the caller must generally have a tighter distribution than the raiser. For example, if the raiser's distribution is {55+, A9+, KJ+}, a potential caller with AJo has just 42.8% equity against this distribution, despite the fact that his hand is well above the minimum of the initial jammer.

In a tournament, this effect is exaggerated. Players behind the first jammer must not only overcome the effect created by playing against a distribution, they must also overcome the loss of equity caused by engaging in a confrontation. If the fraction of their stack that must be committed is small, then this effect is less important, but if they too must call all-in, there can be a significant cost to that as well.

There is a game theory toy game that has some parallels to this situation; it is called the "game of chicken." Two drivers accelerate towards each other at high speed down a long road. Both drivers may swerve to avoid the accident. If both swerve, then both break even, while if both drive on, there is a terrible accident and both lose large amounts of equity. If one drives on and the other swerves, however, the swerver is labeled "chicken," and he loses a small amount of equity for his loss of "face."

Example 28.1 - Game of Chicken

The payoff matrix for this game is as follows:

	Player B	
Player A	Drive on	Swerve
Drive on	(-10, -10)	(1, -1)
Swerve	(-1, 1)	(0, 0)

Note that this game is not zero-sum, so there is not a unique optimal solution. In fact there are three equilibria (strategy pairs where neither player can unilaterally improve). The equilibrium that is most often discussed in game theory is where both players use a mixed strategy of swerving and driving on designed to make the opponent indifferent to either choice. But there are two other equilibria, which are of some interest in poker. The first is where Player A drives on and Player B swerves. Note that Player B cannot improve his payoff; he can only cause mutual destruction. And likewise Player B driving on and Player A swerving is an equilibrium.

In this subgame, poker is played sequentially. So when it is folded around to the blinds and they are at jamming stack sizes, Player A (the first player) can simply declare "drive on," (jam) and B's best response is to swerve (fold). In poker, of course, the penalty for a confrontation is not ten times as large as the penalty for folding one's blinds, and there are common situations where the second player's equity is so high as to ignore this additional penalty altogether. But the idea is that inasmuch as there is a game of chicken at work in this situation, the first player should attempt to establish the favorable equilibrium of jam-fold if possible. This means playing slightly looser than jam or fold, since the confrontation equity penalty makes it more difficult for the opponent to exploit you.

Play On the Bubble

The bubble of tournaments, especially larger multi-table tournaments, is often a tremendous opportunity to exploit tight play. Many players attach an unreasonable amount of psychological value to "finishing in the money," causing them to play extremely tightly just as the step into the money occurs, or when a step within the payout structure occurs (as from three tables to two tables, when the third table is paid out a constant amount). We exploit this phenomenon by liberally raising the blinds of these players and using this period to accumulate chips rapidly. This is one reason that our distribution of tournament finishes is somewhat top-heavy, with a disproportionately large amount of "just out of the money" finishes, and a similarly disproportionate amount of high finishes.

However, there are a fair number of situations where relatively tight play is correct - most obviously in supersatellites, but also in situations where the step into the money is significant. One important example of this is the common sit-and-go tournament structure. Often, these tournaments have ten players, and the payouts are structured 5-3-2. The large jump from zero (4th place) to two units (3rd place) actually makes it correct to play significantly tighter (especially when calling all-in) because the bonus for making it into the money is so large.

Example 28.2

To illustrate this, consider the following situation in a sit-and-go tournament:

4 players remain; all have 2,500 m chips. The payouts are 5-3-2. The blinds are 200-400. The under-the-gun player and the button fold. The small blind jams 2500. What equity does the big blind need in order to call? We can approximate this by looking at the cash values of the possible outcomes. The values here are approximations - with large blinds it's often difficult to accurately estimate the chance of others going broke, etc.

Case 1: He calls and wins and has 5000 in chips. Then his prize equity is at least 50% of 5 plus 50% of something between 2 and 3; let's just say 2.5 for simplicity. Then his prize equity is 3.75.

Case 2: He calls and loses and gets zero.

Case 3: He folds and now has 2100 and the small blind. His prize equity is 21% of 5 plus about 29% of zero plus 50% of something between 2 and 3. So his total prize equity is about 2.3

In order for it to be correct to call, he needs to have $X\%$ equity in the pot such that $3.75X > 2.3$, or 61.33%.

Comparing this to the equity he would need to make it a +chip EV play:

$$(X)5000 - 2100 = 0$$

$$X = 42\%$$

We can see that he should play much tighter on the bubble here.

Adjusting for playing on the bubble at the table can be tricky. One mnemonic that we often use in lieu of attempting to calculate equities is to assess the rough value in chips of moving up one additional place, and utilize that value in our chip equity calculations. For example, suppose we are on the bubble of a 500-player \$100 + 9 tournament where there are 500,000 chips in play. 50th place pays \$120. Our equity from moving up one spot is approximately 1200 chips. When we jam, we might consider the additional penalty when we are called and eliminated to be 1200 chips (the \$120 we forewent by jamming instead of folding). This often doesn't amount to much, as it involves a parlay of being called and then losing (perhaps 25% for the first and 65% for the second, or a total parlay of just 16%, or 200 chips). However, if only 10 places paid, for example, and 10th was \$1200, while 11th was zero, then the penalty would be substantially higher and proper play would be tighter.

Shootouts and Single-Table Satellites

We group these two types of tournaments together because they are similar- in structure; a single table of players play until one player has all the chips. The winner of the table either wins all the buyins or advances to another round where there is a prize structure. In either case, the key strategic point of these single-table tournaments is that there is no step into the money, so chip preservation strategies should only be employed in order to maintain the opportunity to apply skill.

Occasionally in these types of tournaments, the second place finisher gets some sort of prize, which is often only a small fraction of the first place prize. As a result, this prize often has only a small effect on strategy. Normally most efforts should be expended toward winning the table outright.

One benefit of these types of tournaments is that the skilled player gets the opportunity to play shorthanded much more often. This is normally favorable, as many opponents will not have a great deal of experience playing shorthanded. Also, in headsup or shorthanded situations, opportunities to use the jam-or-fold strategies arise much more often. In single-table satellites, deals are often struck once the action is headsup or three-handed. Ordinarily, the important idea is that the percentage of chips a player holds is approximately equal to the percentage of the prize pool he deserves, skill being equal.

Match Play

One format gaining in popularity online is the "match play" format, where the tournament is structured in some fashion as a series of headsup matches. It should be evident that goals of chip preservation are of little value here; the goal is simply to win each match.

One concept that may be of use in playing headsup matches is the knowledge that (excepting small skill adjustments) playing a hand in a headsup match where one player has a larger stack than the other is equivalent equity-wise to playing a hand in a headsup match where both players have stacks equal to the smaller of the two stacks. The result of this concept is that the player with a bigger stack should simply ignore his additional chips (except for skill considerations).

Another thing that bears mentioning (which David Sklansky points out in his book about tournament play), is that in a headsup no-limit match with even moderately sized blinds, neither player can have a tremendous edge. In fact, a great deal of the edge can be neutralized by the weaker player moving all-in preflop. Limit headsup matches often give a larger advantage.

Single-Table Sit-and-Go Tournaments

Another format popularized online but which is now frequently seen in brick and mortar casinos is the "sit-and-go" single table tournament. Generally these are no-limit tournaments (although other games and formats are available, they are less popular), where there is one table, and the prize structure is something like 50% to first place, 30% to second, and 20% to third.

The sit-and-go format has a number of positive elements for the skilled player. First, the tournaments go quickly (often limits escalate every ten hands or ten minutes). Normally, fast escalation of blinds is less favorable for the skilled player, but in this case, many players offset this effect by playing many of these tournaments simultaneously. Second, these tournaments have quite low variance compared to most other types of tournaments since 30% of the field gets paid. A half-bankroll for an expert player even at higher buyin levels will be only 6 to 7 buyins. Third, many parts of the tournament can be played relatively by rote with just a little preparation. For example, when these tournaments get headsup, the blinds are almost always high enough that jam-or-fold play is indicated.

Additionally, there is wide latitude for the application of skill because of the distortion caused by the payout structure. When the tournament becomes four-handed, the difference for a small stack between busting fourth or surviving into the money is quite large. As a result, play on the bubble of a sit-and-go is quite interesting, with the high blinds putting pressure on the small stacks, while the big stacks steal from them. Several different interesting dynamics can develop, all of which require different adjustments. For all of these, we will provide the stacks with four players remaining and make some comments about the ensuing situation.

One Big Stack: (5500/1500/1500/1500, 150-300 blinds)

In this situation, the big stack has a very large advantage as long as these relative stacks are maintained and the tournament remains four-handed. For example, if the big stack simply jammed every hand, the other players would have a difficult problem to solve as they need quite a strong hand to justify calling all-in. For example, if the big stack jams and the big blind is left to potentially defend, we can approximate the equity of the big stack:

From calling and losing: 0.

From calling and winning: $(30\%)(5) + (40\%)(3) + (25\%)(2) + (5\%)(0) = 3.2$ buyins.

From folding: $(10\%)(5) + (30\%)(3) + (30\%)(2) + (30\%)(0) = 2$ buyins.

Hence, he needs $2/3.2$ buyins, or about 62.5% equity against a random hand. The hands that have this are: {66+, A9s+, AT+, KJs+}, or about 7% of hands. The advantage gained by this is so great for the big stack that he can actually consider such counter-intuitive plays as folding the small blind to a very small stack in the big blind in order to prolong this situation.

Four Equal Stacks: (2500/2500/2500/2500, 200-400 blinds)

A like situation occurs here; this is linked to the idea that it is very dangerous for players to call all-in when losing will cause them to go out fourth. In fact, all the adjustments on the bubble revolve around this idea.

Suppose the under-the-gun player jams blind here.

Then each of the other players has a similar decision as the last one; to call all-in, or to fold.

From calling and losing: 0.

From calling and winning: $(50\%)(5) + (30\%)(3) + (20\%)(2) = 3.8$ buyins.

From folding: $(25\%)(5) + (25\%)(3) + (25\%)(2) + (25\%)(0) = 2.5$ buyins.

Here the required equity is even higher: 65.8%, which limits the range of calling hands to {77+, AJs+}. This is only about 5% of hands. Note that even AKo is a fold.

We can see that if the three remaining players will call only 5% of the time each (neglecting coincident overcalls), that the first player must profit from picking up 600 in blinds when not called, even if he has an average of 25% equity when he is called.

So in this type of situation, we see the result of a sequential game of chicken. The first player declares that he's driving straight, and everyone else can only pick him off with powerful hands.

Abuser and Abusee: (6000/2500/750/750, 200-400 blinds)

A third type of situation is as above, where there is one big stack and one smaller stack that is still significantly bigger than the other two stacks. In this situation, the big stack can repeatedly jam at the 2500 stack, and the 2500 is restricted in his ability to call. The 2500 stack desperately wants one of the smaller stacks to bust so that he can play reasonably against the large stack. As a result, the big stack can play passively against the very small stacks in order to maintain this situation where he can steal with impunity.

These are just a few of the situations that can occur near the bubble of a sit-and-go tournament, but the key concept is that the large step into 3rd place (from zero to two buyins) is extremely important in sit-and-go strategy.

Supersatellites

Supersatellites are another very popular structure, particularly as gateways for players with limited bankrolls to enter larger buyin tournaments. The distinguishing characteristic of a supersatellite is the division of the prize pool into a certain number of equal prizes (occasionally with a little cash left over). These prizes are often denominated in entries to a larger tournament. Then all of the final N players receive this prize.

This structure creates extreme distortions in play strategy, especially as the money approaches. For example, in a supersatellite that pays eleven places, any chips a player accumulates beyond perhaps a quarter of the chips in play have virtually no value. It is often correct for a player who has accumulated a large stack to simply fold all hands, effectively walking away from the table.

A story was related to us about a supersatellite to the main event of the World Series of Poker. The "hero" of the story was one of two large stacks left in the supersatellite. There were ten players left, and eight players won seats. Between the two big stacks, they had approximately 75% of the total chips in play and the two big stacks' chips were roughly equal. In any event, one of the big stacks pushed all-in for about 50 times the blind preflop. The other big stack looked down to find AA.

At this point in the story, we interrupted with the question. "Why did the second big stack even look at his cards?" The story went on that the second big stack called and was against JJ and lost. This is a fairly unlucky result, but even with a very large advantage, such as the AA had here, it would still be incorrect to call.

The reason for this is that the AA hand was not wagering chips: he was wagering his equity in a seat. The fact is that his chance of winning a seat by simply folding away all hands was close to 100%. By taking on the other big stack, he could only improve his chances to a number closer to 100%, but he risked being busted out of the tournament. Even though he had an 81% probability of winning the pot, the gain from winning was extremely small compared to the cost of losing.

This distortion manifests itself throughout the later stages of the tournament; the mapping of chips to equity is decidedly non-linear, and the principle that additional chips are worth less is magnified, particularly as one's stack moves near certain threshold levels. In typical tournament structures, there is somewhat of a bias toward being the player to push all-in. In supersatellites, this effect is magnified, and often a player with a medium or large stack is justified in folding to another player's all-in bet even with a reasonably strong hand, because the cost of losing the showdown is so high and the reward for winning it so meager due to the devaluation of the additional chips.

Multi-table Tournaments

Multi-table tournaments are the most common structure for large buyin tournaments, as well as one of the most popular when the buyin is smaller. Players start at a large number of tables, and as seats open up (from players being busted), other tables are broken and the players from the broken tables filled into the seats.

Multi-table tournaments place a higher premium on full-table play than most other tournament structures. However, shorthanded play is still extremely important here, because the situations where the play is shorthanded have a disproportionate impact on total win rate, since they occur with two tables to go and at the final table.

The play in multi-table tournaments is in most cases quite similar to play in a ring game: the payout structure of most multi-table tournaments is quite gradual and lacks a very large step in to the money or from place to place that would cause tournament adjustments. Near or in the money, there are sometimes adjustments to be made, but generally these are much smaller than they are in some of the other types of tournaments discussed here. The fundamental idea is that the game is still poker, and that for all but the most marginal circumstances, gaining chip equity is directly tied to gaining cash equity in the tournament.

This is not to say that a fixed ring game strategy that is appropriate for situations where the money is fairly deep will be effective when the average stack is 10 blinds, however. Nor is it correct from an exploitive standpoint to play a typical ring game strategy if the opponents in the tournament play vastly differently from opponents in a ring game. What is constant is the idea that aggressively accumulating chips should not give way to passively surviving simply because chips are irreplaceable. There are some specific adjustments in play that come about directly because of the tournaments structure, which we will discuss in a moment.

Rebuy Tournaments

Rebuy and add-on tournaments provide the opportunity to sweeten the prize pool beyond the initial buy-in amount by allowing players to buy more chips (subject to certain restrictions) for some specified time period. Often, rebuy tournaments are structured something like this:

There is a buy-in amount $\$X$. with some additional entry fee (such as $\$100 + 9$), for which players receive some fixed number of chips N . Players may "rebuy" as long as their stacks are less than or equal to N . So at the start of the tournament they may buy N chips for $\$X$. They may continue to do this throughout the first three limits of the tournament. After the third limit, the rebuy period is over, and the players may take any last rebuys as well as an "add-on," on which there are no restrictions except that each player make take only one. Often the addon chips are at a discount to the original buyin. For example, if each player started with 1000 chips, the add-on might be 1500 chips for $\$100$.

This naturally leads to the question of whether to rebuy. We can look at a number of situations and assess whether a rebuy is indicated or not. For all of these cases, we ignore risk and variance

issues and assume that the players are adequately bankrolled for the tournaments they are entering and their goal is to maximize equity on their entries. In particular, then-decision to enter the tournament was rational.

Case 1:

It's the beginning of the rebuy period. Player A loses his starting chips on the first hand. Should he rebuy?

Almost certainly yes. Player A can simply view this as an opportunity to buy-in to a completely different rebuy tournament, except that the second tournament has no entry fee. Since he was willing to pay the entry fee to this tournament because he believes his edge is large enough to overcome it, he must certainly have sufficient edge to buy into the second tournament. Hence he should rebuy.

Case 2:

It's the end of the rebuy period- Player A has increased his starting chips of 1000 to 6000. There are 300 players, and Player A thinks his chance of doubling up overall is 54%. The total prize pool is likely to be about \$90,000 (200 players will take the add-on, making a total of one million chips in play). He can now take an add-on for \$100 and get 1500 in chips. Should he do so? How large does his stack have to be before he should decline the add-on?

Player A's equity can be estimated using the Theory of Doubling Up. If A has 6000 chips, then the number of doubles he needs in order to have all the chips is:

$$N(6000) = \log_2 1,000,000/6,000 = 7.38$$

$$N(7500) = \log_2 1,000,000/7,500 = 7.06$$

$$E(6000) = 0.54^{7.38} = (0.0105) (90,000) = \$953$$

$$E(7500) = 0.54^{7.06} = (0.0129) (90,000) = \$1,162$$

Since the value in equity of the additional chips is more than \$200 but they cost only \$100. A should definitely take the add-on.

For the second part of the question, we need to find a stack size x such that A is indifferent to taking the rebuy. We know that:

$$N(\text{no add-on}) = \log^2 1,000,000/x$$

$$N(\text{add-on}) = \log^2 1,000,000/(x+ 1500)$$

Additionally, we know that:

$$E(x) = E(x+ 1500) - \$100$$

$$E(x) = (90,000) (0.54^{N(\text{noadd-on})})$$

$$E(x+ 1500) = (90,000)(0.54^{N(\text{add-on})})$$

Solving these equations simultaneously, we can find that the value of the add-on is greater than \$100 all the way to stack sizes of half the chips in play (well above the level at which the Theory of Doubling Up would break down). This indicates that for this skill level, buying the discounted add-on chips is basically always correct.

Even at a C of 0.51, it would be correct to take this discounted add-on at all stack sizes. However, at a C of 0.51. the breakeven point for buying 1500 additional chips at par (for \$150)

is near 8,325.

What we can glean from this is that assuming you were a winning enough player to enter the tournament in the first place, taking add-ons where the chips are at a premium is virtually always correct. Rebuying at par, however, should be done with care. The Theory of Doubling Up can be a guide to the skill-adjusted value of chips here.

Playing with Shortened Stacks

One of the distinguishing characteristics of tournament play is that generally, some players in the field will not bust until they are forced to by the tournament structure. As a result, tournament structures almost always include escalating blinds. As a result of this, chip stacks are normally reduced to levels (relative to the blinds) that are rarely found in cash games. For example, it is not unusual for the late stages of a no-limit holdem tournament to have an average stack of just ten to fifteen blinds. However, no-limit holdem ring games are rarely-played in situations where the average stack is this small. By "shortened" stacks, we do not exclusively mean small stacks relative to the other players; we also refer to the situation where the stacks are shortened relative to the blinds because the blinds have escalated to the point where all the stacks are short.

Two major types of adjustments are made. The first primarily affects no-limit play. At smaller stack-to-blind ratios, there comes a point where in lieu of raising a small amount, we simply move all-in. This threshold generally occurs at a stack size where making an appropriate smaller raise commits us to the pot anyway, because we will be getting at least 2 to 1 on calling an all-in reraise. *A general rule of thumb for this is that if our stack is six times the preflop pot (including antes and blinds) or more, then we should prefer to raise to an appropriate smaller amount, while below this level, we should prefer to jam all-in.* This is close to the recommendation we made in the jam-or-fold section, where approximately ten blinds ($6 \times 1.5 = 9$) was the threshold for playing jam-or-fold.

Second, hand selection in shortened-stack play is adjusted from larger-stack play. Hands that have high showdown equity are preferred over hands that have high implied odds potential, and hands that contain "blockers" are often valuable because they make it more likely that opponents will not have strong enough hands to call. We saw the first of these effects in the jam-or-fold discussion, where hands such as T9s were jams up to very high blind levels. This is because when called, these types of hands have high equity against the distribution of calling hands. The second of these effects makes it more worthwhile to raise with hands such as Ax, because holding an ace removes one from the opponent's potential holdings. In no-limit, this second effect is very strong when raising in late position because there are fewer players behind who can call.

Playing with Antes

In stud, and also in the middle and later stages of many no-limit holdem tournaments, antes are present. Antes affect play with short stacks dramatically because they create additional overlay for players with very short stacks. Consider a stud situation where a player has just three or four antes left. When this player commits his chips, he will often play against just one or two opponents but will get eight to one on his ante chips, in addition to even money action on his remaining chips. The taller the small stack is, the less pronounced this overlay is. Contrast this with, for example, limit holdem, where frequently a short stack with just two blinds will play against the big blind only when he decides to commit, with only the overlay of the small blind preventing him from getting even money. This means that in stud tournaments, players should be a little more willing to pass up terrible hands with very small stacks because as their stacks shrink, their potential overlay grows.

This also occasionally allows for play adjustments in mixed-game tournaments; occasionally it will be correct to pass up marginal hands in flop games with a small stack if the stud round is coming up in order to get this overlay.

In no-limit holdem, of course, the antes create a situation where there is more money in the pot without increasing the size of the blinds. This makes raising first in more attractive, as picking up the blinds for a small raise can often be profitable immediately. Exploitively, it is frequently the case that players play so tightly in certain ante situations, especially when the antes are large, that raising every hand is profitable. For example, suppose the limits are 200-400 with a 50 ante in a 9-handed game, and by raising to 1000, you expect to pick up the blinds 50% of the time. Then you can simply raise with any two cards, since you immediately win 1050 when everyone folds. In addition, when players do play with you, you will occasionally have a strong hand. These situations occur frequently, and from an exploitive standpoint it is definitely worth it to be aware of the players who will fold too often to raises when there are antes in play.

Tournament Backing Agreements

In Part IV, we provided some discussion about the value of a backing agreement to the backer and the player. In that section, we considered only ring game backing agreements. However, tournament backing agreements are much more popular and common, and so we would be remiss to not provide a short discussion of evaluating such things.

The most common structure (in our experience) of a tournament backing agreement is some sort of open-ended agreement with the following conditions:

- The backer provides some amount of bankroll.
- The player plays tournaments according to the agreement.
- There is some division of profits (often split evenly between the backer and the player).
- Most agreements allow the player to take a distribution whenever the player is ahead -this will typically be after a moderate to large sized win.

One thing to consider in this type of arrangement is that since the fixed bankroll is static, an agreement with the above conditions will have a risk of "ruin" - that is, that the fixed bankroll will be lost - of 100%. This is because previous profit will have been distributed to the backer and the backee. This does not necessarily mean that the backing agreement is doomed to lose money; only that if the backer is unwilling to increase the size of the bankroll, at some point, the player will lose the entire bankroll amount.

We can attempt to evaluate a closed-ended tournament agreement, however, using the idea of effective tournament size. Suppose that we have a hypothetical player A with an estimated win rate of about 1 buyin per tournament, and that A is going to enter into a backing agreement with a backer B to play 70 tournaments. At the end they will settle up with B taking any losses and the wins being split equally between A and B. The tournaments are of size approximately 400 players.

We can approximate the outcomes of these tournaments by considering them to be winner-take-all tournaments of a smaller size. We'll use the tournament variance multiple k from the previous section (Equation 27.4), which for a tournament of 400 players with some typical payout structure is about 72. We can use the binomial distribution to assess our player's chance of winning such a tournament. Since his edge is approximately 1 buyin, his chance of winning the tournament is approximately 1 in 36.

His chances of winning the virtual tournament N times in 70 are:

<i>N</i>	% chance
0	13.92%
1	27.84%
2	27.44%
3	17.77%
4	8.50%
5	3.21%
6	0.99%
7	0.26%
8	0.06%
9	0.01%
10	0.00%

These are not his chances of winning the real tournament; those are of course much lower. However, his outcomes from playing the real tournament can be roughly- approximated by his outcomes of playing the virtual tournament (of a much smaller size).

We can convert the above table to a form that more accurately reflects the backing arrangement by adding some columns related to the settlement.

<i>N</i>	<i>p</i> (<i>N</i>)	Outcome	Backer's share	Weighted share
0	13.92%	- 70	(70.00)	(9.74)
1	27.84%	2	1.00	0.28
2	27.44%	74	37.00	10.15
3	17.77%	146	73.00	12.97
4	8.50%	218	109.00	9.27
5	3.21%	290	145.00	4.65
6	0.99%	362	181.00	1.80
7	0.26%	434	217.00	0.56
8	0.06%	506	253.00	0.15
9	0.01%	578	289.00	0.03
10	0.00%	650	325.00	0.01
Total				30.13

The backer and player together will win 70 buyins total (because of the player's 1 buyin per tournament win rate). The backer's share is only 30.13 because of the backing agreement, instead of 35. So the backer loses about 4.9 buyins to the player because of the nature of this backing agreement. This is a significantly better case for the backer than the usual agreement, because settlement occurs only at the end of the seventy tournaments. We can account for this settlement effect using, for example, computer simulation.

We simulated 10,000 series of 70 tournaments with a player who finished in each money place at twice the normal rate and seeded at the end of 70 tournaments. The average return to the backer here was approximately 32 buyins. The variance between these two numbers is likely a systematic bias. This method (using the variance multiplier) slightly underestimates the return to the backer. We then ran the same simulations, but allowed the player to settle whenever he was ahead. The average return to the backer here was approximately 27 buyins. So "settling when ahead" here is worth about five buyins (out of 70, remember) to the player.

"Fairness" in backing agreements is a relative term. In a very real sense, backing agreements exist in a marketplace and the rules of the marketplace dictate what the "price" for backing a player is. Backers should be aware; however, that the shorter the duration of the backing agreement and the more often the player may settle, the expectation of the backer is diminished by the skew of tournament results.

Key Concepts

- In treating tournaments (with multiple payouts) using game theory, we must consider that the game is not strictly a two-player game. When two players enter a confrontation, all the other players gain. Hence, tournament play is heavily influenced by the "game of chicken" concept - it is often far better to be the first one into the pot for all the chips than to be the second.
- Many different tournament structures have been proposed and have become popular. Each of these structures introduces different specific play adjustments in order to maximize equity.
- Tournament backing agreements can be evaluated using some of the same tools as we used for evaluating ring game backing agreements. The severe skew of tournament results makes these backing agreements often very favorable for the player, particularly if settlement occurs when the player is ahead.
- One specific recommendation for no-limit tournaments is six that if our stack is six times the preflop pot (including antes and blinds) or more, then we should prefer to raise to an appropriate smaller amount, while below this level, we should prefer to jam all-in.

Chapter 29

Three's a Crowd: Multiplayer Games

Throughout the preceding material we have primarily focused on two-player games. One reason for this is practical. Especially at higher limits, many pots are contested heads up, and the degree to which a player can make money in the game is tied very strongly to his ability to play well in heads-up situations. Another reason that we focus primarily on heads-up play is that it yields more easily to analysis than multiplayer play.

Multiplayer games, as we shall see presently, frequently include, in addition to their "unexploitable" elements, strong exploitive elements that make them much more complex than two-player games. What frequently happens is that there is an equilibrium that creates a situation where no single player can unilaterally improve his equity. However, individual players can disturb the equilibrium and create situations where other players have incentive to change their play, implicitly colluding against the other players. These situations are difficult to analyze and have no "solution" per se; instead, we can only learn about how to cope with the various types of exploitation we may find.

However, since most poker games are in fact multiplayer games, we would be remiss in not at least introducing the field of multiplayer game theory. Recall the idea of equilibrium from two-player zero-sum games (see Chapter 10). This situation occurs when neither player can unilaterally improve his equity by changing strategy. We stated previously in Chapter 10 that all zero-sum two-player games have at least one such equilibrium. In the 1950s, mathematician John Nash proved that all games with a finite number of players, zero-sum or not, possess at least one such equilibrium. We call these *Nash equilibria*. Of course, the equilibria for two-player zero-sum games are also Nash equilibria, but we normally apply the Nash modifier only when dealing with multiplayer games.

A Nash equilibrium is a set of a strategies for all players such that no single player can improve his expectation from the game by unilaterally changing his strategy.

In two-player zero-sum games, it was typical that if there were multiple equilibria, they were all similar except that some contained dominated strategy choices that were Irrelevant against an optimal opponent. However, in multiplayer games, there may be several Nash equilibria with quite different characters.

To see this, we can revisit the two-player non-zero-sum "game of chicken" introduced in Chapter 28 (see Example 28, 1).

	Player B	
Player A	Drive on	Swerve
Drive on	(-10, -10)	(1, -1)
Swerve	(-1, 1)	(0, 0)

We can convert this game to a zero-sum game by adding a third passive player G, who has no strategy decisions, but simply has an outcome equal to the negative of the other two players' outcomes:

	Player B	
Player A	Drive on	Swerve
Drive on	(-10, -10, +20)	(1, -1, 0)
Swerve	(-1, 1, 0)	(0, 0, 0)

This type of conversion can be done to any non-zero-sum game in order to convert it into a zero-sum game with one additional player. In any event, recall that we identified three equilibria for this game:

- A and B both swerve 90% of the time and drive straight on 10% of the time.
- A drives on 100% of the time and B swerves 100% of the time.
- B drives on 100% of the time and A swerves 100% of the time.

If you check, you will see that in each of these cases, neither player can improve his own outcome unilaterally by changing strategy.

Finding a Nash equilibrium for a game is an important first step in analyzing it. In just about all cases, the Nash equilibrium does represent strong play for each player. In two-player zero-sum games, strategies that form a Nash equilibrium are optimal. But Nash equilibria alone do not suffice to describe the best play in multiplayer games. There are a number of reasons for this, but one of the most important is the presence of *alliances*.

Just as in strategy board games, players may form alliances within a multiplayer game to their (presumed) mutual benefit. This causes the analysis of multiplayer games a serious problem, which is that multiple players can "collude" to enhance their joint EV at the expense of other players. Often there are many possible alliances that could be formed, and choosing from among them is not a matter that can be solved directly by analysis, whether using game theoretic or simply exploitive analysis.

We can subdivide collusive alliances into "hard" and "soft" varieties. These are distinguished by whether or not players are sharing secret information about the game and whether or not they are acting to maximize their joint or their individual expectation. "Hard" collusion, which occurs when players share secret information about the game or when they act against their own self-interest to the joint interest of the alliance, is prohibited directly by the rules of the game. For example, a cheating team might signal their confederates about their hole card holdings and manipulate the betting to their benefit. Or, players might agree to share tournament equity in a tournament and then play softly against each other or work to even out their chips, thus maximizing their joint expectation in the tournament.

Such behavior subverts the game because individuals are no longer acting in their own individual self-interest (within the game). Two or more competent colluders would have little trouble extracting a substantial amount of additional expectation over and above what they would make from playing individually. In terms of analysis, we often simply ignore "hard" collusion as a matter of course, deeming it to be outside of the rules of the game. However, players should be aware that hard collusion does exist and is often very difficult to detect.

"Soft" collusion, on the other hand, does not normally violate the rules and is an integral part of multiplayer poker. The canonical example of this concept is the extremely loose and passive play in low-limit games. Often what occurs in these games is that most if not all of the players have an informal, unspoken agreement that they will play very loosely and passively on early streets. As a result of this unspoken agreement, plays that would be equity disasters in tight, aggressive, games (such as limping under the gun in limit holdem with 75o) become only mild mistakes. This concept has been labeled "implicit collusion" by other authors, and this is a reasonable term with which to describe it.

It may seem to the reader that this implicit collusion is relatively unimportant from an exploitive point of view. After all, it would seem that if we can accurately read our opponents' hands and

strategies, we can make money regardless of what they do as long as the game is not biased against us. However, this turns out to be false in some circumstances.

Example 29.1

Consider the following toy game:

Three players play and they ante \$10 each, making a pot of \$30. Each player receives one card from a standard deck. One player has the button, which rotates after each hand. There is a round of betting. Each player may check or bet \$30. Once a player has bet, then the remaining players act in turn to either call or fold. At the showdown, all players who have not folded reveal their cards. All players holding an ace or king (a winner) split the pot, unless there are no aces or kings, in which case all remaining players (who hold losers) split the pot. This game is simple and symmetrical. It would seem that on each hand, the button player would have some sort of advantage, but since the button rotates, the game remains fair.

Now presume that we are invited to play this game by two players. The two players agree to do two things: to state their strategies at the outset and to agree to show whatever information is necessary (folded hands, etc.) to verify that they are in fact playing those strategies. On the face of it, this would likely seem like a good opportunity. So we sit down.

The player to our left states his strategy as:

"Check all hands if possible, and only call bets when I have a winner."

The player to our right states his strategy as:

"Bet all hands, and call all bets."

It's plain to see that each of these strategies individually is fairly exploitable, and that we can determine the proper exploitive response and our expectation against these strategies. We can simplify our strategy to simply deciding whether to call the bet of the player to our right (call him the *maniac*). Betting any hands is simply equivalent to calling his bet. Obviously, if we have an ace or king, we will call the bets. The question is what to do if we have a losing card?

There are three cases of action:

Case 1: We have the button.

In this case, the player to our left (call him the *rock*) checks and the maniac bets.

Case 2: The rock has the button.

In this case, the maniac bets immediately.

Case 3: The maniac has the button.

In this case, both we and the rock check, and the maniac bets.

In all three of these cases, we have gained no information about either hand and it is our turn.

When we do not hold an ace or king, neither of the opponents will hold a winner ($\frac{43}{51}$) ($\frac{42}{50}$) of the time, or about 70.8% of the time. The remaining 29.2% of the time, we will have zero equity in the pot. When neither of the opponents hold a winner, if we call, the total pot will be \$90, and we will get half of it, or \$45.

$$\langle \text{call} \rangle = (0.708)(45) + (0.292)(0) - \$30 = \$1.87$$

$$\langle \text{folds} \rangle = 0$$

So it seems that our best approach is to call with a loser, as this has higher expectation than folding. With this data in hand, we can calculate our total EV for the game:

Our's hand	Opp's hand	Result (inc. antes)	Probability	EV
Winner	Both winners	\$0	0.25%	0\$
Winner	R-W, M-L	+ \$20	1.86%	+\$0.372
Winner	R-L, M-W	+ \$5	1.86%	+\$0.093
Winner	Both losers	+ \$50	11.41%	+\$5.707
Loser	Both losers	+ \$5	59.93%	+\$2.996
Loser	1 or 2 W	- \$40	24.69%	-\$9.875
Total			100%	-\$0.707

What a strange result! Here we have a game where we know both players' strategies, each of those strategies seems weak and exploitable, the game is symmetrical and otherwise fair...and yet we cannot win.

This is because these two players' strategies form an implicit alliance. By playing with these strategies they actually collude against the player in the middle, without exchanging information or even playing reactively to the situation within the game. It is probable, although we have not as yet proved it, that similar situations exist in real-life three-handed poker games, such that two players could specify their entire strategies before the game started and make it impossible for the third player to have positive expectation.

Also note that the third player here has the opportunity to decide which of the opponents will win. If he plays like the rock, then the maniac will be the only net winner in the game. If he plays like the maniac, the rock will be the only one to come out ahead. In a sense, the third player in has the opportunity to join an alliance with either of the other two players, by playing the opposite of their strategy. However, none of the alliances he can join result in him personally gaining equity.

There are other types of alliances possible; in fact, sometimes there may be multiple possible alliances on the same hand. In this case, which alliances should be chosen by the various players is not always indicated by a simple EV analysis.

Consider the following three-handed situation in holdem (cards face-up and a moderately-sized pot):

Player A: A♠ A♣

Player B: A♥ Q♥

Player C: 8♥ 7♥

Board: 9♥ 6♥ 2♦

Now let's consider what collusive possibilities there are.

Suppose that the made hand and the straight flush draw decide to collude. So the straight flush draw bets and the aces raise. Now the nut flush draw is in a difficult spot. He has just five outs on the turn, and doesn't have odds to call. When the nut flush draw folds, the colluding hands have won the entire pot.

Or suppose that the made hand and the nut flush draw decide to collude. Now the straight flush draw has eight outs and is a little less than 2-1 to make his hand. He does, then, have odds to call. So the flush draw and the made hand (who control together 79% of the equity in the pot) are able to extract four bets from the straight flush draw.

Or suppose that the two flush draws decide to collude. Then the nut flush draw folds, because the nut flush draw adds no additional equity to his alliance. The straight flush draw assumes all the outs the nut flush draw had, and the alliance pays just one bet to draw to all their outs.

So it's clear from examining this that one difficulty with studying multiplayer games using game theory is that it's unclear how the play should go even when the cards are face-up, due to the possible collusive alliances. An important and emerging branch of mathematics, *cooperative* game theory, is the study of games where the players may cooperate to improve their joint equity in the game.

We will only briefly mention this topic. A more complete treatment is beyond the scope of this text. Cooperative game theory is often geared toward games where *side payments* are allowed; that is, players may offer each other inducements outside the game to create alliances. In poker, this might take the form of chip passing or playing out of the same bankroll, both of which are prohibited or at least frowned upon by the rules of the game. Additionally, cooperative game theory often deals with situations of *repeated play*, where the players play the game over and over, with the players responding to new information revealed with each iteration. Both of these conditions can apply, but especially the last is useful in actual poker play. A concept of considerable theoretical interest is the *Shapley value*, which is a sort of weighted average of the value of all the possible alliances into which a player may enter. For more discussion on this and related topics, we refer you to texts on cooperative game theory.

In the second part of this chapter, we will discuss (note that the term "solve" doesn't have a clear meaning here) three multiplayer scenarios and find some lessons. The first two of these situations were originally posed to us by Chris Ferguson; both deal with situations where one player is all-in.

The conventional wisdom in situations where one player is all-in with no side pot, particularly in tournaments, is that the two remaining players should a) be more inclined to check the hand down in an attempt to eliminate the all-in player or b) at least refrain from bluffing with weak hands as there is no additional money to be won by bluffing out the other player who still has chips. But it isn't so clear as that...

Example 29.2 - Dry Side Pot Game #1

Three players, X, Y, and Z play.

The game is stud, and it's the river.

We solve only a one-street game here for illustrative purposes. Previous betting would influence the action, of course in real poker.

X: (??) A♥ A♦ A♣ A♥ (?)

Y: (??) 5♦ 6♦ 7♦ 8♦ (?)

Z: (??) K♥ K♣ K♦ K♠ (?)

Z is all-in for the pot of *Pbets*.

X and Y remain with chips.

One full street.

Limit betting.

Neither X nor Z can improve on the hands that they're showing; hence we have here what is essentially a situation where Y is clairvoyant - either he has a straight flush or not. X should never bet, and should call Y's bets enough to make Y indifferent to bluffing. But Y can never be

indifferent to bluffing; if he does bluff, even if X folds, he cannot beat Z's hand and never gains by bluffing X out of the pot. So X needn't call at all.

The Nash equilibrium for this game is that Y value bets his straight flushes, and X always folds to a bet. None of the players can unilaterally improve their equity.

But suppose Y does bluff, just a little bit. Every time that Y bluffs and X folds, Z wins the pot. Y gains nothing by doing this, of course, except to shift equity from X to Z. This will likely cause X aggravation, as Y is taking money away from X and giving it to Z. To counteract this, X might call Y's bets some fraction of the time so that he can retain his equity in the pot.

But if he does this, Y gains equity immediately because X will now be calling Y's value bets as well. In fact, Y doesn't even need to bluff in an appropriate ratio; he can gain in this way by bluffing just a few hands. Of course, X could just play the equilibrium strategy - in that case, Y risks nothing. In fact, in this situation, Y cannot be hurt by X's actions.

This is a simple example that shows how the player betting can sometimes benefit from deviating from equilibrium in order to tempt other players away from equilibrium. Next we consider a more complex example.

Example 29.2 - Dry Side Pot Game #2

Three players, X, Y, and Z play.

Each has a uniform hand from the $[0, 1]$ distribution.

Z is all-in for the antes - the pot is 9 bets.

We solve this game for a particular pot size because the general case is a bit complicated and the complication gets in the way of understanding the lesson.

X and Y remain with chips.

One half street.

Limit betting.

After X checks dark, let us consider Y's strategy. Clearly Y wants to bet his very best hands for value, hoping to get called by X with a mediocre hand. In the analogous game without an all-in player, Y would also bluff with his very worst hands. However, here if he does successfully bluff out X, he still has to beat Z.

Is he then limited to simply betting for value? If he does this, it seems that X can exploit him by folding mediocre hands instead of calling. Y should attempt to prevent this in some way. One thing is clear, however; X will have one threshold for calling Y's bets.

Call X's threshold for calling x , and Y's threshold for value betting y_1 . Suppose that Y is going to bet some hands worse than y_1 (separated by a checking region). We'll call these bets "bluffs," although they're not the same as the very worst hands Y bet in previous games. In fact they are semi-bluffs, which hope that X will fold a hand better than Y's but than Y's hand will still beat Z's. X will call whenever he has positive equity by doing so. Clearly, X will never call with a hand that cannot beat Y's "bluffs," so the threshold x will be between y_1 and Y's bluffing region. For now we'll just call the bluffing region YB, and its call its width w . It is not true that in this game the size of the bluffing region is a times the size of the value betting region, because the all-in player reduces the equity of a successful bluff.

X's equity of calling Y's bet at the threshold x is:

Y's hand	Z's hand	Probability	<X, call>
$[0, y_1]$	$[0, 1]$	y_1	- 1
Y_B	$[0, x]$	$w(x)$	+ 1
Y_B	$[x, 1]$	$w(1 - x)$	+ 10

X's total equity from calling at X must equal zero (since he must be indifferent to calling or folding).

$$wx + 10(w(1 - x)) = y_1$$

$$wx + 10w - 10wx = y_1$$

$$w(10 - 9x) = y_1$$

$$w = y_1 / (10 - 9x)$$

At the upper bound of Y's bluffing region (call this y for a moment), we have:

X's hand	Z's hand	Probability	<Y, bluff>
$[0, x]$	$[0, 1]$	x	- 1
$[x, 1]$	$[0, y]$	$y(1 - x)$	0
$[x, y]$	$[y, 1]$	$(y - x)(1 - y)$	+ 9
$[y, 1]$	$[y, 1]$	$(1 - y)^2$	0

$$x = 9(y - x)(1 - y)$$

$$x = 9y - 9x - 9y^2 + 9xy$$

$$10x = 9(1 + x)y - 9y^2$$

$$9y^2 - 9(1 + x)y + 10x = 0$$

$$y^2 - (1 + x)y + 10x/9 = 0$$

$$y = (1 + x) / 2 \pm (\sqrt{(1 + x)^2 - 40x/9}) / 2$$

This indifference equation yields a quadratic. This is because there are two points at which Y is indifferent between bluffing and checking; the upper bound of the bluffing region as well as the lower bound of that region.

Notice that this equation actually specifies the entire bluffing region. We can see that the center of the region will be $(1 + x) / 2$, and the width of the region will be distance between the two

bounds, or $(\sqrt{(1 + x)^2 - 40x/9})$.

$$w = (\sqrt{(1 + x)^2 - 40x/9})$$

And at y_1 , we have:

X's hand	Z's hand	Probability	<Y, value bet>
$[0, y_1]$	$[0, 1]$	y_1	- 1
$[y_1, x]$	$[0, 1]$	$x - y_1$	+ 1
$[x, 1]$	$[0, 1]$	$(1 - x)$	0

This, by far the easiest of the indifference equations, is the familiar idea that Y should bet with half the hands with which X will call.

$$y_1 = x - y_1$$

$$y_1 = \frac{1}{2}x$$

We now have three indifference equations:

$$w = y_1 / (10 - 9x)$$

$$y_1 = \frac{1}{2}x$$

$$w = (\sqrt{1+x})^2 - 40x/9$$

Solving these simultaneously, we get:

$$w = x / (20 - 18x)$$

$$x / (20 - 18x) = (\sqrt{1+x})^2 - 40x/9$$

$$x \approx 0.2848$$

Using this, we can identify the other thresholds. Y's bluffing region will be centered on 0.6424, and have a width of 0.0190 units, and y_1 will be approximately 0.1424.

So we can see that the "a" value for this game is slightly higher than in the games with no all-in player. Y bets about 14% of his hands for value, and bluffs about 1.9%. α in this case would be $\frac{1}{10}$. There is a dramatic effect, however, on X's calling strategy. Instead of calling 9/10 of the time as he would do if Z were not present, he calls just 28.5% of the time.

These values comprise the Nash equilibrium for this game - neither X nor Y can unilaterally improve his equity by changing strategies.

But now we can illustrate the multiplayer difficulties we discussed at the beginning of the chapter. Suppose Y decides to bluff more often. He decides to widen w to be 3% instead of 1.9%. Recall that we showed in the analysis immediately preceding that Y was indifferent to bluffing at the endpoints of his bluffing region in the Nash equilibrium. Since he's bluffing hands outside this range, his expectation will be lower.

X will gain a small amount of this value when he calls because Y more frequently has a bluff. But the player who benefits most of all is Z, because Y bluffs out X's hands that are between X's calling threshold and Y's bluffing threshold much more often. Whenever the situation was such that X's hand was better than Z's hand, which in turn was better than Y's hand, and Y bluffs X out, Z gains the entire pot. Overall, Y takes a small loss from overbluffing and being called by X, However. X loses a larger amount on the hands where he is bluffed out; the "extra" equity goes to Z.

X can modify his strategy to try to increase his equity (remember, the Nash equilibrium has been disturbed by Y's strategy change). Suppose X calls with more of his in-between hands -that is, he increases x , his calling threshold. Now X gets more of the pots that he's "entitled" to when $X < Z < Y$, but he also winds up paying off Y more often when Y has a value bet. Looking at the situation from Y's point of view, X calling more is actually quite welcome. Y gains a bet (for value) about 14.2% of the time, loses a bet about 3% of the time, and loses about a third of the 9-

unit pot as well. Sometimes Y bluffs X out when Z has a weak hand and wins the entire pot. Adding all these up results in a net gain to Y of about 2% of a bet when X calls with additional hands. By doing this, X reclaims most of his "lost" equity from Z - **but by doing so, he must redistribute some of it to T.**

The counter-intuitive idea here is that when Y overbluffs, and X plays to maximize his expectation, Y gains. In fact, it is in Y's best interest to make sure that X knows he will overbluff. Otherwise, X might just play his Nash equilibrium strategy and Y would lose equity by overbluffing. So Y is much better off to announce his strategy to X **in the hopes that X will exploit it.** By overbluffing, Y unilaterally enters into an alliance with Z. The alliance will bluff X out of the pot and give Z equity. However, by announcing his strategy, Y offers X the opportunity to enter into a different alliance - this time X and Y against Z. The only catch is that X has to give Y equity in order to enter into this alliance. The alternative is that X can ally with Z (at personal cost to X) to punish Y for his manipulative behavior.

As an important aside, bluffing into dry side pots violates convention, and there are strong reasons not to engage in this behavior because of the overtly collusive aspect of it. Players should consider carefully whether the tiny hit of potential equity they may gain by making these plays is sufficiently compensated by potential retaliation in future hands or tournaments. These issues are fairly unclear. We present this example, though, to show an example of a wholly counter-intuitive situation that exists only because of the all-in player's presence and the opportunity for Y to transfer equity from X to Z.

Example 29.4 - The Clairvoyant Maniac Game

Three players (X, Y, and Z) play.

Z is clairvoyant, while X and Y have equivalent hands.

One half street

X and Y may only check and call/fold, Z may bet one unit.

Pot of 9 units.

All players have chips left.

You may notice that we often utilize a pot size of 9 units; this is because a (from Equation 11.1) is a convenient value of $\frac{1}{10}$ in this case.

First, we can find the Nash equilibrium for this game. Our experience with half-street games tells us that Z will value bet his winning hands and some fraction of his losing hands. He will do this such that the two players remaining will be indifferent to calling. It turns out here that there are many Nash equilibria. In fact, any strategy that X and Y follow that satisfies the following two conditions is a Nash equilibrium for this game.

- The total fraction of the time that Z gets called is $\frac{9}{10}$.
- Y never overcalls when X calls.

The first of these is the standard "make the bluffer indifferent to bluffing" idea: the second is simply a matter of pot odds for the overcaller. If X calls first, then Y stands to split the pot with X by calling. When Z value bets, Y loses one bet by calling. When Z bluffs, Y gains a net of five bets by overcalling. Since Z will only bluff $\frac{1}{10}$ as often as he will value bet, Y can never overall. Suppose that we simply take as our candidate equilibrium that X calls 50% of the time, and Y 80% of the time when X folds. Then we have a Nash equilibrium and the value of the game to Y is $\frac{9}{10}$ of a bet when he has a value betting hand.

As we saw previously, however, Z can attempt to disturb the equilibrium by bluffing more often.

Suppose that Z announces that he is going to bluff $\frac{1}{8}$ as much as he value bets instead of $\frac{1}{10}$. Now X's exploitive response is to call all the time, picking off Z's bluffs. However, Y is still unable to overcall because he lacks pot odds. This is actually bad for Z; he loses additional equity on his failed bluffs and does not gain enough back from additional called value bets to offset this.

Nevertheless, suppose that Z decides that he must not be bluffing enough. So instead he announces that he will begin bluffing three times as often as he would in equilibrium, or $\frac{3}{10}$ as often as he value bets. X will of course keep calling all the time. But consider Y's situation now. $\frac{3}{13}$ of the time, Z is bluffing, and Y can win five bets net by overcalling. The remaining $\frac{10}{13}$ of the time, Z is value betting and Y loses a bet by overcalling. Now Y has positive expectation overcalling Z's bet and X's call.

Who gains from this? The answer is, of course, Z! Now the $\frac{10}{13}$ of the time he is value betting, he gains two bets from the call and overcall, while losing only one bet when he bluffs. So he gains a total of $\frac{17}{10}$ bets per value bet instead of only $\frac{9}{10}$.

The upshot of all this is that if Z plays the equilibrium strategy, no alliances are formed. If he bluffs just a little more than the equilibrium amount, then X and Z forms an alliance against Y. If he bluffs a lot more often than the equilibrium amount, then Y and Z form an alliance against X. The key to this game is how Z can increase his equity substantially by bluffing much more often than the equilibrium strategy, assuming his opponents will move to maximize their equity.

We have considered just a few multiplayer games here: the pattern should hopefully be clear. We can always find a Nash equilibrium for a game, where all the players cannot unilaterally improve their equity. Often, however, one or more players can disturb the equilibrium by changing strategies. When the other players move to maximize their equity by exploiting the disturbance, alliances are formed, and often the disturber can gain equity from the shift, without counter-exploiting at all. This is of course impossible in two-player zero-sum games, as any equity that one player gains must be lost by the other player. It is this idea that leads us to refrain from using the term "optimal" in regard to multiplayer strategies.

Key Concepts

- Multiplayer games are much more complicated than two-player games because players can collude with one another, entering into implicit alliances to maximize their joint equity. These "soft" collusive alliances are not prohibited by the rules of the game - in fact; they are an integral part of it.
- A Nash equilibrium is a set of a strategies for all players such that no single player can improve his expectation from the game by unilaterally changing his strategy.
- In many situations, there are several alliances possible. Each player must decide which alliance to enter and how to induce others to ally with him.
- We saw several examples of situations where one player was able to disturb the equilibrium in such a way that he benefited unless another player took a negative EV action to punish him for disturbing the equilibrium.
- It is often the case that the best way to disturb the equilibrium is to bluff much more often.

Chapter 30

Putting It All Together: Using Math to Improve Play

Over the course of this book, we've considered a wide variety of situations. We've looked at a number of different toy games and tried to distill the salient lessons from their solutions. We've looked at a number of real poker examples and even solved one common real-life poker situation (the jam-or-fold game in Chapter 12). There's a lot of theory in this book and a lot of material that may be unfamiliar even to those who have some experience with game theory and other branches of mathematics.

But ultimately, this book is about playing poker; in home games, cardrooms and online. It's true that poker does hold some intrinsic intellectual value, and it's not implausible that we might study it even if it were not profitable to do so. The fact that poker, unlike most games, is played for large sums of money makes its study not only worthwhile from an intellectual standpoint (as might also be the case for chess, backgammon, or bridge), but from a financial standpoint as well. So the work we have done that has been presented in this book is not aimed at satisfying our urge to solve mathematical puzzles, but to enhance our (and now your) ability to earn money at poker.

With that in mind, it seems relevant to articulate our philosophy of playing poker, which finds many of its roots in the topics of this book, in the hopes that it will aid the reader in what might well be the most difficult part of taking on a game theoretical approach to the game -attempting to apply the principles of game theory to actual play in practice. Applying these principles in many cases requires retraining the mind to think in new and different ways.

The problem that arises when trying to play using game theory is that we don't have optimal strategies to the common and popular poker games. In fact, in multiplayer scenarios, no such strategies exist. So using a game theoretic approach means attempting to approximate. It is this process of attempting to approximate optimal play that we will discuss here.

We have identified four major areas where our strategic and tactical choices are informed by quantitative or game theoretic concepts, which we will consider in turn.

- Aggression.
- Playing with balance.
- Playing strategically.
- Relying on analysis.

Each of these areas is important in it's own right, but there is significant crossover between categories.

Aggression

It has surely been the mantra of poker authors for years that aggressive play is correct, aggressive play gets the money, and so on. Yet we find that many players, even players who play fairly well, seem to settle in at a level of aggression with which they are comfortable. They characterize play more passive than theirs as "weak." and play more aggressive than theirs as "maniacal." It is surely possible to play too passively (this is extremely common at lower limits), and it is surely possible to play too aggressively (more common at higher limits, although not unheard of elsewhere).

Both of the authors have been the subjects of a fair amount of criticism for their "maniacal" play.

Yet we believe that the rather high level of aggression with which we play is supported by analysis. When we study toy games (particularly games that are quite similar to real poker, such as the jam-or-fold game), we are often surprised at how aggressive these strategies are. When the stacks are 10 blinds, for example, we jam with 58,3% of hands. Many players consider this to be reckless, but we have demonstrated to our satisfaction that such a strategy cannot be exploited and extracts the maximum value against strong play. As a result, we have consistently outperformed our chip-value expectation when headsup.

This theme recurs as we look at other types of games. When we looked at the one-and-a-half street clairvoyance game (see Example 20.5), we saw that the concealed hand frequently semi-bluffed with hands that were at a significant disadvantage. In the high-low [0,1] game (Example 18.1), we saw that the second player bets more than 75% of his hands if checked to, and so on. What we find consistently throughout our study is that optimal play tends to be very aggressive. It is occasionally true that there are games where strong play is a little more passive. This frequently occurs when the antes are small relative to bet sizes. But in general, we find throughout our work that aggression is a central component of optimal strategies.

We can also consider the cost of mistakes. The cost of being a little too aggressive is often small compared to the cost of folding a little too much. Folding too often surrenders the entire pot to the aggressor, while putting in the occasional extra bet costs only the extra bet. In the jam-or-fold game, calling too conservatively as the defender can be very costly against an exploitive opponent.

This is in fact a rule that has wide application across many forms of poker, but especially limit. Players who fold too often are often making larger and more exploitable mistakes than players who are too aggressive. We also believe that in many common situations, our opponents make these mistakes frequently. Players who are obsessed with making "big laydowns" are all too common at the middle and higher limits. As a result, playing aggressively, while in and of itself a strong strategy even against strong opponents, simultaneously serves to automatically exploit the most common errors made by opponents. This is one of the important features of optimal or near-optimal play in poker; it is not simply a matter of seizing zero equity (as optimal play in Roshambo does, for example). Near-optimal play also does extremely well in extracting value from an opponent's mistakes.

For these reasons, the conscious effort to play aggressively and extract value from marginal situations is one of the important principles of our poker play. We almost always raise when we are the first to enter a pot. We defend our blinds (in flop games) liberally and play the later streets with vigor. When we are faced with unclear decisions, or when we feel that two plays are EV-neutral, we often adopt the more aggressive play as a matter of principle. We are willing to splash our chips around in many situations where others might not. In the tournament setting, particularly, this often creates the image of a player who will seek confrontation and as a result, some players will avoid contesting pots with us. In this way, this is exploitive, but we are careful not to overextend our aggression into purely exploitive play because our opponents may counter-exploit. What analysis tells us, however, is that optimal play seems to contain a much higher level of aggression than is commonly thought, and so we act accordingly.

Balance

Some suboptimal strategies give up value because they are exploitable. Other suboptimal strategies give up value because they are too passive or too aggressive or too tight or too loose, or whatever. As an example, consider the half-street [0,1] game. If both sides are playing optimally, the first player checks dark and the second player bets and bluffs with appropriate fractions of his hands. We now consider two ways in which the second player can change his

strategy to make it suboptimal.

One way he could change his strategy is to bluff less but value bet just as often. The first player's proper exploitive response is to tighten up his calling standards, since he isn't getting as much value from picking off bluffs, but is still losing the same amount when he calls a value bet. By doing this, he exploits the second player's strategy and gains value.

A second way in which the second player could change his strategy to make it suboptimal is to change his frequencies of betting and bluffing so that they are still in an appropriate ratio (a bluffs for each value bet), but he bets and bluffs a little less often. Then the first player has no exploitive response; he is still indifferent to calling with his medium-strength hands. So in this sense, this type of strategy change is unexploitable. However, the first player passively gains value because the second player fails to extract as much value as he can. We call this type of strategy balanced, because it doesn't lose value because of the opponent's exploitation.

Note that balance applies only to strategies, not to hands. This is because we rarely consider the play of the actual hand we hold in isolation. Instead, at each decision point, we construct a strategy. Each possible hand we can be holding (based on the previous action) is assigned an action. Ideally, these assignments meet all the following criteria:

- For each action, the set of hands that will take that action contains some hands that benefit substantially from a particular opponent's response. Consider, for example, that we are considering raising our opponent's bet on the turn in holdem. There are three possible responses to our raise: the opponent could reraise, call, or fold. Our strongest made hands benefit the most from the opponent reraising, so they should be represented in our raising range. Likewise, our moderately strong made hands benefit from the opponent calling, and our weaker semi-bluffs benefit the most from our opponent folding. So we might construct our raising range such that each of these hand types is appropriately represented. If we do this properly, the opponent will not be able to exploit us by taking one action more frequently than others in response to our action. Any value we lose from one type of hand will be made up by another.
- The set of hands that we play in a particular manner is sufficiently diverse to prevent exploitation. This satisfies the principle of information hiding.
- We assign hands to actions in such a way that hands are played in their most profitable manner as often as possible. Recall that when we previously discussed playing before the flop against a raiser with AA and AK in holdem in Chapter 9, the differential between playing aces one way or the other in isolation was small, but with AK it was large. As a result, we chose to jam with both hands because that strategy benefited AK disproportionately. Likewise, we construct distributions of hands with an eye to maximizing the value of the primary hands in that distribution.
- In cases where there is significant action to come, as occurs frequently in big bet poker, we try whenever possible to have a credible threat of having the nuts on a later street. If your previous play in a particular situation marks you as not having nut hands in your distribution, then this is ripe for exploitation.

As you can see, playing with balance is primarily concerned with not being exploited. In combination with an appropriate level of aggression, balance is the core of what we might call near-optimal play.

Optimal play is perfectly balanced and has the precisely appropriate level of aggression. Near-optimal play is the attempt to create strategies that are as close as possible to optimal, given the limitations of our ability to measure this.

A player who plays with perfect balance is unreadable from a strategy standpoint. In practice, it is impossible to determine what "perfect" balance is. However, one approach that is useful is to simply pretend that you must announce your strategy to the opponents before each action. To typical players, who are used to drinking that a description of their strategies is immensely valuable information, this may seem ludicrous. But when you consider that this is the nature of optimal strategies - the optimal player can simply announce his strategy in its entirety and the opponent cannot make use of this information except to avoid making mistakes - this approach makes a lot more sense. Evaluating and exposing an entire strategy forces us to consider those elements of the strategy that are exploitable, and in our quest for balance, this type of thinking is quite valuable. This type of thinking leads naturally into the next of the four major areas, which is playing strategically.

Playing Strategically

Both playing with balance and with an appropriate level of aggression are intimately related to the third major area, which is the broad heading of playing strategically. Essentially, what we mean by this term is widening our view onto the game of poker, both inside of a particular game or a particular hand, and outside of the game, from a risk management standpoint.

Within the game, there are a number of different ways in which we can think strategically. One of the most important is in changing our view of the game from a game with many-decision points to a game with fairly few. In other words, we try to understand that our strategy carries through from street to street or decision to decision and that before making a decision, we must realize that our decision directly affects all our future decisions. As an example, take the situation in a no-limit holdem tournament where a player has T3600 left, with blinds of 200-400. If this player open-raises to 1200 on the button, he will be getting at least 2-1 to call an all-in reraise from one of the blinds. As a result, he will likely be forced to call with any hand he initially raised. By raising to this lesser amount, he has committed himself to the pot, while giving his opponents the option to flat call or put all the money in preflop. Our choice in this situation is to raise all-in preflop, because it is a sounder play from a strategic sense. We have the same number of strategic options, while our opponents have fewer.

In the same way, thinking strategically sees the several streets of the game as one connected entity as much as possible. We frequently make decisions for at least one and sometimes two streets in advance, often with conditions related to the card that falls. For example, we might make a play such as call on the turn with the intention of "walling" the river if a blank hits, and folding to a bet or checking behind if the flush draw comes in. The "wall" strategy is a term that we frequently use in describing play on the river. The player playing the "wall" calls a bet but bets if checked to, thus ensuring that 1-2 bets go in on the river. One strength of thinking of the hand in this way is that it ensures that we accurately account for implied odds and the potential future cost of actions.

In the same way, when we discussed the topic of balance, we pointed out that this strategic element was present there as well. We seek to play distributions of hands whenever possible, in preference to individual hands. In the same way, we often play our distributions in a way that is analogous to multi-street play in the clairvoyant game; betting and bluffing in a ratio on each street, and giving up on some bluffs on each succeeding street. Of course, in real poker, the bluffs are in fact semi-bluffs, and the non-static nature of the game changes the value of our value betting hands from street to street. But the principle remains the same. Doing this accurately requires that we know our complete distribution at each point in the hand that is the set of hands is consistent with our previous action.

Our fellow poker theorists Paul R. Pudaite and Chris Ferguson call this process "reading your

own hand," a delightful phrase that we wish we had coined. Poker books (including this one!) discuss to a significant extent the process of finding the distribution of hands your opponent can hold, but from a game theoretic view, the distribution of hands your opponent holds (except at the beginning of the game) is irrelevant. Optimal strategy is playing your hands in a perfectly balanced and appropriately aggressive manner. If your opponent can exploit you by playing a different set of hands than you were expecting, then your play is suboptimal.

There is a difficulty to this, and it is a difficulty in all strategic thinking about poker that prevents many players from taking this approach to the game. The difficulty is that poker is complicated, and there are a very large number of game nodes. It is generally beyond the limits of the human mind to treat the entire game as a single strategic entity. Instead, we compromise, and search for insight through solving games which are simpler and also through attempting to exploit proposed strategies ourselves in order to find their weaknesses. In addition, we try as much as possible to unify the various aspects of the game in our minds - but that being said, there are many practical moments where we simply reduce a poker situation to an EV calculation. Being able to do this at the table (accurately estimate our opponent's distribution and our equity against that distribution) is a valuable skill.

Reliance on Analysis

The fourth area in which we utilize the mathematics of poker is in how we learn about the game, both away from the table and at it. We are constantly made aware of how misleading empirical evidence is in poker. Few things that we perceive are true in a probabilistic sense - we have already covered the difficulty of obtaining sample sizes for strategy reading, and so on. It is not at all unusual for strong players to go on horrible runs, and often even players with strong emotional makeup can be shaken by particularly negative results. Even when things are not going badly, a few self-selected unlucky or unusual data points may cause players to question their play and make adjustments.

Even more importantly, many players gain knowledge about the game by reading books or asking then talented or successful friends. While the advice given by these sources is likely honest, it is less likely to be correct, and even more so, less likely to be correct as understood by the recipient. Incorrectly applying advice to situations where it does not apply is a common error. Even when the message is correctly understood, it has significant danger of being misinformed or having a bias from results-oriented thinking.

We find that relying on analysis as our basis for poker knowledge is a much more reliable and impartial way of learning about the game. Instead of asking advice from "better players," we construct games to solve and situations to discuss with others with an eye to demonstrating the nature of the situation mathematically or through some model.

This is helpful in two major ways. First, a properly completed analysis is correct, given a set of assumptions. If you start with assumptions, and you analyze the situation without making errors, the solution will be correct. Properly applying it to situations takes care and practice. But fundamentally we believe that this method is superior to accepting someone else's word for it, not least because you gain a much deeper understanding of a situation by performing the analysis yourself.

Second, using mathematics in poker gives us a rock-steady source of information. If we jam on the button with T8s three tournaments in a row in a jam-or-fold situation and get called and lose each time, there is no temptation to back off and play tighter. Because of the analysis we have done, we know that this is the right play. It is easier to ward off tilt or self-doubt when you are confident that your play is correct. Using mathematics, we can obtain proven solutions that generate that confidence and help us play better as a result.

A Final Word

Over the course of this book, we have attempted to provide an introduction to the application of mathematics to poker. We have presented results and formulas, some of which will be of direct use to you immediately. Just the jam-or-fold tables in Chapter 13 might be worth the price of this book per day to a professional sit-and-go player who does not currently have this information. In addition, however, we have tried to present a mode of thinking about poker that eschews the guessing game about what the opponent's hand is for a more sophisticated approach that uses distributions and balanced play to extract value. This volume is not exhaustive; in fact, it is but an early harbinger of what we expect to be a quantitative revolution in the way poker is studied and played at the highest levels. We look forward to the exciting work that will be done in this field.

At the outset, we stated that our ideas could be distilled to one simple piece of play advice: Maximize average profit. Hopefully over the course of this book, we have been able to present new ways of looking at what this statement means and given you new methods to employ at the table so that you can in fact make more money at poker.

Recommended Reading

Books on Game Theory:

- Luce, R. Duncan and Raffa, Howard. *Games and Decisions*. Dover. (1957)
- von Neumann, John and Morgenstern, Oskar. *Theory of Games and Economic Behavior*. Princeton University Press. (1944)
- Williams, J. D. *The Compleat Strategist*. Dover. (1966)

Books on poker and gambling in general:

- Epstein, Richard A. *Theory of Gambling and Statistical Logic*. Academic Press. (1967)
- Sklansky, David. *Theory of Poker*. Two Plus Two Publishing (1987)

Books on related topics:

- Bernstein, Peter. *Against the Gods: The Remarkable Story of Risk*. John Wiley & Sons. (1996)
- Taleb, Nassim Nicholas. *Foiled by Randomness: The Hidden Role of Chance in Life and in the Markets*. Thomson. (2004)

About the Authors

Bill Chen was born in Williamsburg, Virginia in 1970 and grew up in Alabama. He received a PhD in Mathematics from the University of California at Berkeley in 1999. He started thinking about poker while a grad student when he augmented his National Science Foundation grant with poker winnings from local clubs. Bill currently lives near Philadelphia.

Jerrold Ankenman grew up in Orange County, California. He received a BS in Business Management from Pepperdine University in 2004. Jerrod picked up poker as a hobby from reading rec.gambling.poker, and as a career from collaborating on poker research with Bill (and others). He is currently a professional player and lives in Avon, Connecticut with wife Michelle and stepchildren Abby and Michael.

Bill and Jerrod started corresponding in 2000 and collaborating in 2002. The collaboration culminated not only in many new ideas about poker and this book, but also in success at the poker table, namely Bill winning two World Series of Poker events in 2006, Jerrod taking second in one event, and both placing in the money in other events.

About the Publisher

ConJelCo specializes in books and software for the serious gambler. ConJelCo publishes books by Lee Jones, Russ Fox, Scott Harker, Roy Cooke, John Bond, Bill Chen, Jerrod Ankenman, Lou Krieger, Mark Tenner, Kalthleen Keller Watterson, Nick Christenson and Dan Paymar. ConJelCo also publishes software including StatKing for Windows (record keeping software).

Complete information on ConJelCo products is on the Web at <http://www.conjelco.com> where you'll also find lots of other useful information about poker and gambling.