

Probability, Game Theory and Poker
Course notes for MAT1374

Pieter Hofstra
Department of Mathematics and Statistics
University of Ottawa

CONTENTS

Preface	ix
I Game theory: a first look	1
I.1 Sequential Games	3
I.1.1 The Investment Game	3
I.1.2 Game Trees	4
I.1.3 Rollback Equilibria	7
I.1.4 Monopoly Game	11
I.1.5 Sucker Bets and other mistakes	12
I.1.6 NIM	13
I.1.7 Pirates	14
I.2 Payoffs, Utility and Real Life	15
I.2.1 Utility	15
I.2.2 Ultimatum game	16
I.2.3 Behavioral Game Theory	18
I.2.4 The Centipede	19
I.2.5 Why do we play?	19
I.2.6 Rationality	21
I.3 Levels of Thinking	22
I.3.1 Simultaneous games	22
I.3.2 Some central concepts in simultaneous games	25
I.3.3 Chicken	29

I.3.4	Coordination game	30
I.3.5	Battle of the sexes	32
I.3.6	Nash Equilibria	32
I.4	Game Theory in context	34
I.4.1	Origins of game theory	34
I.4.2	RAND and the nuclear arms race	35
I.4.3	Schelling's Work	36
I.4.4	John Nash	37
II	What are the chances?	39
II.1	On Games of Chance	41
II.1.1	Pre-history of Probability	41
II.1.2	Girolamo Cardano	42
II.1.3	Cardano's Law	44
II.1.4	Roulette	45
II.1.5	The Missing Ingredient	47
II.1.6	Some Examples	48
II.2	Hand Combinations and Ranges	50
II.2.1	Starting Hands	50
II.2.2	Hand Ranges	52
II.2.3	Narrowing down a range	53
II.3	What to Expect?	56
II.3.1	Pascal's Wager	56
II.3.2	Expected Value	57
II.3.3	Roulette Expectation	59
II.3.4	Pot Odds	61
II.3.5	On Bluffing	62
II.4	The Tools of Probability Theory	64
II.4.1	The World of Events	64
II.4.2	The Laws of Probability	65
II.4.3	Conditional Probability	69
II.4.4	Independence and the Gambler's Fallacy	71
II.4.5	Intermezzo: Heuristics and Biases	73

II.4.6	Problems Revisited	76
III	Mixing it Up	79
III.1	Games and Probability	80
III.1.1	Incomplete Information	80
III.1.2	Truel	81
III.2	Mixed Strategies	84
III.2.1	Matching Pennies	84
III.2.2	Finding the Right Mix	86
III.2.3	Nash Equilibria in Mixed Strategies	88
III.3	The AKQ Game	92
III.3.1	Solving the AKQ game	93
III.3.2	Possible Extensions	96
III.3.3	Exploitive vs. Equilibrium Play	97
III.4	Strategic Moves	98
III.4.1	Binding Oneself	98
III.4.2	Commitments, Threats, Promises	99
III.4.3	Credibility	101
III.4.4	Irrationality	104
III.5	The Bottom Line	105
III.5.1	Risk of Ruin	105
III.5.2	The Kelly Criterion	107
III.5.3	On Free Lunches	110
	Resources and suggested reading	113

LIST OF FIGURES

I.1	Game tree for the Investment Game	5
I.2	Example Game Tree	6
I.3	Strategies in a Game	7
I.4	Example of a Play	8
I.5	Game tree for the Investment Game after first pruning	8
I.6	Game tree for the Investment Game after second pruning	9
I.7	Game tree with multiple equilibria	10
I.8	Game tree for the Monopoly Game	11
I.9	Game tree for the monopoly game, pruned	12
I.10	Game tree for the Ultimatum Game	17
I.11	Payoff Matrix for the Prisoner's Dilemma	23
I.12	A 3x4 game	27
I.13	Result of IEDS	29
I.14	Payoff Matrix for Chicken	30
I.15	Payoff Matrix for the Pure Coordination Game	31
I.16	Payoff Matrix for the Battle of the Sexes	32
I.17	A game without NEs in pure strategies	34
II.1	Roulette Layout	45
II.2	Union, Intersection, Complement and Difference of Events.	66
II.3	Restricting the Sample Space	70
III.1	Matching Pennies	85

III.2 Matching Pennies, Version 2	86
III.3 Rock-Paper-Scissors	89
III.4 Rock-Paper-Scissors, Variation	90
III.5 Risk of Ruin, $p = 0.55$	107

PREFACE

While every mathematics curriculum includes at least one course in probability theory and while many universities offer a course in game theory, a course with subtitle “Poker101” is certainly not an everyday occurrence¹. The reasons for developing the course were the following: first, that it is quite common to use games of chance as illustrative examples in courses on elementary probability theory; second, that poker, both live and online, is becoming more and more popular; and third, that our university does not offer many courses in game theory, even though it is an excellent topic to show to a broad audience that mathematics can be a greatly rewarding, entertaining and useful subject.

These notes form the backbone of the course and have been written having in mind a general undergraduate audience with no particular background in mathematics beyond basic high school algebra. Some knowledge of the rules of poker (in particular hand rankings) are assumed at various points; in the suggested reading material some references are provided. On occasion, the reader will find remarks which are only useful to more experienced players; these remarks are tangential and can be skipped without fear of losing the main thread.

One of my main aims for the course is to give students the intellectual tools to think their way through a wide variety of strategical problems, and to show them how the basic concepts from probability theory and game theory go a long way to help sound decision making. Of course, my aspiration was also to convey the usefulness of the subject and to give non-mathematics students a chance to see an entertaining and applicable branch of mathematics in action in various corners of the real world.

Thus if there is one overarching theme to the material, it is the problem of making decisions in situations of uncertainty in the widest possible sense. This of course includes finding good strategies in games, but also applies to auctions, negotiations, political processes, investing and much more. Inevitably, many interesting subjects had to be left out, and the choice of material very much reflects the author’s personal preferences. Moreover, several topics appear in the actual lectures which are not treated in these notes, most notably psychological aspects of decision making and societal aspects of gambling. Again, the suggested reading material contains some excellent resources for students who wish to know more about these subject areas.

The material is organized in three chapters: the first introduces the basic concepts of game theory insofar as they do not require probability theory. Chapter two deals with elementary probability theory

¹The full title of the course is “Probability and Games of Chance: Poker101”.

and its applications. Finally, the third chapter combines the first two by studying mixed, or randomized strategies in games; the analysis of such strategies rely heavily on the concept of expected value and can hence only be understood after the treatment to probability.

As for exposition, I have aimed for an informal interplay between theory, historical background and real life applications. When it comes to mathematical results, no formal proofs are given, and the text concentrates on conveying the plausibility of the ideas through examples.

There are several people to whom I'm indebted. Within the Department of Mathematics and Statistics, I should like to thank David McDonald for his continuing support and encouragement, as well as several other people who have helped make this possible. My teaching assistant Dynimul Mao has been a great help with the preparation of the material for the lectures and other supporting materials. Finally, I'd like to thank all the students who have given feedback and helpful suggestions for improvement.

CHAPTER I

GAME THEORY: A FIRST LOOK

The name *game theory* is not entirely fortunate: it has connotations with frivolity and playfulness, and doesn't quite suggest that this is a serious and rigorous field of study. Officially, game theory is classified as applied mathematics, and has branched out into economics, social sciences, biology and more. A better name would perhaps be: *theory of strategy*, because at heart, game theory is about understanding strategic interactions between individuals. This may be in the context of a friendly parlor game, or it may involve two or more nuclear powers who are each trying to decide how much money to spend on weapons of mass destruction. It can be about negotiations or bargaining problems, or about politicians trying to regulate global warming or overfishing. It can be about genes competing for survival, or about a parent trying to change his or her child's behavior.

Just to whet your appetite a bit and to get you thinking about some strategic issues, here are a few questions game theory sheds light on.

- Why is it sometimes the case that it is better to know less? Here, I don't mean that not knowing something can be better for your peace of mind, as in: "I don't want to know what's this hamburger I'm about to eat is made of". What I have in mind here are situations where you are at a genuine strategic disadvantage because you have certain information.
- Why is it sometimes the case that it is better to have fewer options? Again, this should not be understood as saying simply that having fewer options makes for easier decisions. That would just be laziness. No, there are situations where the mere fact that you have a certain option makes your strategic position weak(er).
- Why would you ever want to be in a situation where you could get blackmailed? Could it be possible that at some point you would like to reveal to your competitor that you have a dirty little secret which could damage your reputation?

- What are good strategies for negotiating a price or a salary? Is a request for a small raise always more likely to be granted than for a large one?
- There are many different types of auctions. In an English auction, the price of the item goes up until only one bidder is left. In a Dutch auction, the price goes down until the first bidder raises his hand. In a silent Vickrey auction, each bidder submits a bid in a sealed envelope, and the winner pays the second highest price. How should we bid in these various types of auctions and avoid paying too much?
- How can we break through the “infinite levels of thinking” problem? By this, I mean the “I know that you know that I know...”-train of thought. What is a good way of reasoning in situations where we get trapped in such confusing patterns of thought?

Games come in many forms and guises, and there are many ways of categorizing them. For example, we may distinguish games with one, two or many players. Some games are purely strategic, such as chess; others, by contrast, have no strategic component and are pure luck, such as the lottery. In between, there are games which involve some luck but also skill, such as poker or backgammon. There are games of complete information, where all players have access to all the relevant facts about the game (again, chess would be an example); on the other hand, in some games the players may have varying degrees of uncertainty about various aspects of the game. Some are games of pure conflict, where one player’s gain is the other player’s loss; others are cooperative, in the sense that players need to work together in order to win. And of course, many of these features can be combined, so that we may consider, for example, purely strategic games of complete information, or cooperative games of incomplete information.

We won’t give a mathematical definition here¹, but whenever you are confronted with a game, it is good to identify the following four aspects:

- Players: who are they and how many are there?
- Moves: what can the players do and when can they do it?
- Information: what do the players know about the rules of the game and the other players? Do some players have information during the game which is not available to others?
- Payoffs: what do the various players win or lose in each of the possible outcomes of the game?

Our explorations begin with a class of games called *sequential games*. These include very complicated games such as chess, but we will focus on simple miniature games. The reason is that often even a very simple game can bring to light important ideas and concepts, which we need to understand before we can have any hope of understanding more complicated games.

Our first objective in this chapter is to introduce a few simple sequential games so that we get familiar with basic strategic ideas. We then learn how to represent such games by trees, and introduce some useful game-theoretic terminology. The important lesson we extract from our examples is that in order to determine the best strategy in a sequential game, you have to start your analysis at the end of the game, and reason backwards. We will also discover that even if all players play perfectly, the outcome of a game may be counterintuitive, or have certain undesirable features. Finally, we include

¹If you want a precise definition see for example the textbook by Myerson, but be warned that it is rather involved.

in the discussion of sequential games some suggestions on how to avoid strategic mistakes in practical situations.

Next, we address the issue of payoffs: in order to analyze a game, we need to express how the players value the different possible outcomes of the game. This is done using the concept of utility, which allows us to incorporate into the analysis of a game the fact that the players may care about other things than just winning the game or making as much money as possible.

Not all games are sequential: in many situations one has to make a decision without knowing in advance what the other person is going to do, and vice versa. Such situations are modelled using simultaneous games. There are lots of interesting miniature simultaneous games, each of them containing valuable insights about strategic principles. We study several classic examples, ranging from the Prisoner's Dilemma to the Battle of the Sexes.

Understanding these games requires some new concepts and ideas: the first of these are the notions of a dominating and of a dominated strategies. We learn how to recognize those and use them to draw conclusions about how games will play out. We then come to the most important concept of this chapter, namely that of a Nash Equilibrium. This is a solution concept for simultaneous games which makes precise what counts as optimal strategy in simultaneous games.

For many of the games we encounter we will observe that there is a discrepancy between the mathematical theory and observed practical play. This is interesting, because it indicates that human decision making in strategic situations often involves many considerations different from mathematical ones. The discipline concerned with understanding how people play in practice and why they play the way they do is called behavioral game theory, and we shall discuss some insights which emerge from the vast body of experimental results.

Finally, we will have a brief look at the origins of the subject, the political and military conflicts which were a driving force behind the development, and some of the main contributors.



I.1 SEQUENTIAL GAMES

Sequential games are games in which the players make their moves in a particular, predetermined order (as opposed to simultaneous games, in which the players move at the same time, without knowing what the other players do). Many typical parlor games, such as chess, go or checkers are sequential. We focus mainly on two-player games for now, and use our first example, the investment game, as an opportunity to introduce some basic ideas and terminology.

I.1.1 THE INVESTMENT GAME

Our first example of a sequential game is the *investment game*. This is a two-player game, which goes as follows.

- The first player chooses an amount to invest: \$0, \$1 or \$3, and puts the investment in an envelope.

If he invests \$0 (which simply means he doesn't invest at all) then the game ends. Otherwise, he passes on the envelope containing his investment to the second player.

- Next, the second player looks in the envelope, and has two options. Her first option is take the cash out of it, in which case the game ends and she gets to keep the money. Her second option is to match the investment, i.e. add \$1 if the envelope contained \$1, and add \$3 if it contained \$3. Then the envelope is passed on to the host.
- The host now pays interest on the investment; if the total investment is \$2 it grows to \$5, and if the total is \$6 it grows to \$10. In either case, the resulting amount is divided equally over the two players.

We assume that both players are simply interested in making as much money as possible, and that there are no considerations about fairness or fear of making a socially undesirable decision.

Our goal now is to understand how this game will play out. To this end, we ask ourselves how Player 2 should respond to each of Player 1's possible actions. Of course, if Player 1 doesn't invest at all, there is nothing to say. So suppose Player 1 decides to invest \$1. What is then the best response for Player 2? Well, taking the money earns her \$1, while matching the investment leads to half of the \$5 pot, giving a \$1,50 profit. Thus matching is the better choice. And what is the best response to Player 1 investing \$3? Well, taking the money gives her \$3, while matching gives her \$2 (after investing \$3 she gets \$5 back). Thus in this case, taking the money is better.

Now that we know what Player 2 will do in each of the scenarios, we can start thinking about what Player 1 should do. If he doesn't invest, he doesn't earn anything. If he invests \$1, then he knows that Player 2 will match it, and so he will earn \$1,50. And if he invests \$3, then he knows that Player 2 will take it, and so he loses his \$3. Of the three possibilities, the second is best: we have deduced that the best course of action is to invest \$1. In the next section we will explain how to make this type of reasoning more systematic.

I.1.2 GAME TREES

The method of reasoning we employed in order to analyse the investment game is sound, but we would like to make things a bit more precise. As a first step it is convenient to represent and visualise the game in the form of a *game tree*: such a tree gives a depiction of all possible sequences of moves in the game. In a game tree, the *nodes* (the coloured circles and diamonds) represent the players, the *edges* represent the moves, and at the *leaves* of the tree we record the *payoffs*. In this case, the payoffs represent the net monetary gain/loss by each of the players. The game tree corresponding to the investment game is depicted in Figure I.1. Our convention will be to display in red the information pertaining to Player 1, and in blue that to Player 2.

Of course, sequential games can have more than two players, in which case we have nodes suitably labelled for each of them. One player games exist as well (for example, sudoku, or solitaire), but since we are mainly interested in the interdependencies between the decision making processes of different individuals we shall not consider those games here.

In the literature on game theory you may find two more features of game trees. The first has to do with the role of *chance*. The investment game is a game of pure strategy, like chess: no element of luck

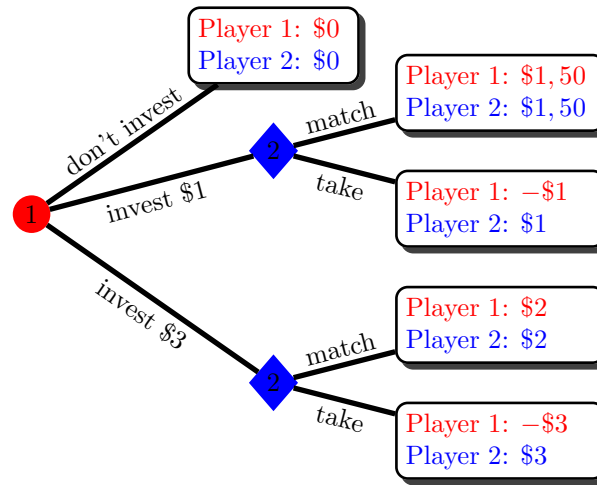


Figure I.1: Game tree for the Investment Game

is present. If we want to investigate sequential games in which chance plays a role (for example, where certain outcomes depend on the roll of a die) then it is customary to introduce an extra player called Nature, or Chance, and regard the possible chance events as moves by that player. Another feature of the investment game is that it is a game of *complete information*: at every stage of the game, both players know everything which is relevant to the game. In order to model games where one player has information not accessible to other players we need to have a way to indicate this in the tree. This is usually done by adding extra labels to the nodes, which tell us what information the player has at that stage.

Now that we can represent games as trees, we may introduce the central notion of a strategy in a game.

Strategies

A *strategy* for Player i consists of a specification of a choice of move at each node labelled by i .

Thus, a strategy for a player may be thought of as a complete set of instructions, which tell the player what to do in each possible situation he may find himself during the game. Note that this includes a choice of moves at nodes which may never be reached (for example because they lie on branches which are incompatible with earlier strategic choices).

To illustrate this, consider the game tree in Figure I.2: In this game, Player 1 starts by choosing one of three moves a, b, or c. If Player 1 chooses a, then Player 2 may respond with k or with m. If she plays

m, the game ends, but if she plays k, Player 1 gets to move, and may choose between u and v. When Player 1 plays b, Player 2 chooses between k and m; if she plays l the game ends, while if she plays k, Player 1 must choose between u,v and w. When Player 1 starts with c, the game ends when Player 2 plays k; otherwise she plays m, and Player 1 chooses between u and v.

Note that we have simply given the payoffs as numbers, one for each player; for the purposes of illustrating the notion of strategy these numbers don't matter.




Figure I.2: Example Game Tree

We now give examples of strategies in this game. In Figure I.3 we have indicated in red one of the strategies for Player 1 and in blue one of the available strategies for Player 2. Note that at *every* node which belongs to Player 1 (that is, every node which is labelled by 1), the strategy specifies a choice of move, and similarly for Player 2. Even though Player 1 chooses b and therefore actual play will never go through the branch (a,k), we must still indicate that he would choose v if he were to find himself at this node. Of course, there are many other possible strategies for either player, and the strategies indicated here are by no means the best ones.

Now suppose that we are given a strategy for Player 1, as well as a strategy for Player 2. Then this uniquely determines a path in the game tree starting at the root of the tree (the beginning of the game) to one of the leaves. In Figure I.4 we have indicated in orange the play arising from the two strategies given in Figure I.3.

Such a path is also called a *play* in the game. Thus a play is one possible way the game might play out, and there are as many possible plays in a game as there are leaves in the game tree. We may represent the play in Figure I.4 by (b,k,u).

Exercise 1. In the above example game, how many strategies does Player 1 have? How many does Player 2 have? How many plays are there?

Exercise 2. How does the definition of strategy and of play differ from everyday use of the word? What is a strategy for white in a chess game? What is a play?

I.1.3 ROLLBACK EQUILIBRIA

We have explained what we mean by a strategy in a sequential game. Of course, what we want to know next is how to distinguish the good from the bad strategies. Recall that when we analysed the investment game, we first thought about what Player 2 would do as a response to the various possible moves of Player 1. Thus, we started our analysis from the point of view of the last move in the game, not the first. This idea works in general: given a sequential game, we may look at the game tree and start at the leaves, eliminating unfavorable last moves one by one (by “pruning” the tree) and repeating this procedure until there are no inferior responses left. Whatever branches are left at the end of this pruning procedure are strategies that are part of a *rollback equilibrium*. Since backward induction results in a strategy for every player, this means that a rollback equilibrium consists of one strategy for every player, and that each of these strategies are optimal in the sense that none of the players could improve his or her payoffs by changing strategies.

Here is an illustration of this procedure using the tree of the investment game. Note that in essence, this takes us through the same reasoning as we used before to solve this game, but now we have a way of visualising it using the tree.

We start by examining the last moves, and see for each node what move would give the highest payoff for that player. In the middle branch, matching yields \$1.50 to player 2, while taking yields only \$1. Therefore we prune the “take”-branch at that node, because player 2 will never go down that branch.




Figure I.3: Strategies in a Game




Figure I.4: Example of a Play

Similarly, in the bottom branch, player 2 will never match, because that yields only \$2 as opposed to the \$3 she can earn by taking. Thus we prune that “match”-branch there. The result then is the tree in Figure I.5, where the dotted edges indicate that these have been cut off.

We now continue this process, by considering all the branches originating at the node for player 1. There are three, ending in payoffs of \$0, \$1.50 and $-\$3$, respectively, so we prune the first and the last

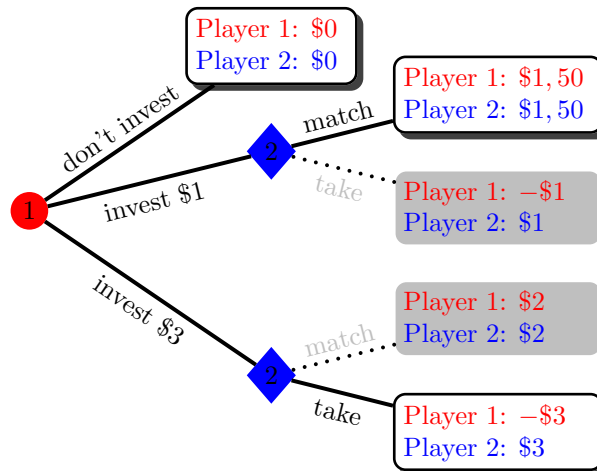


Figure I.5: Game tree for the Investment Game after first pruning

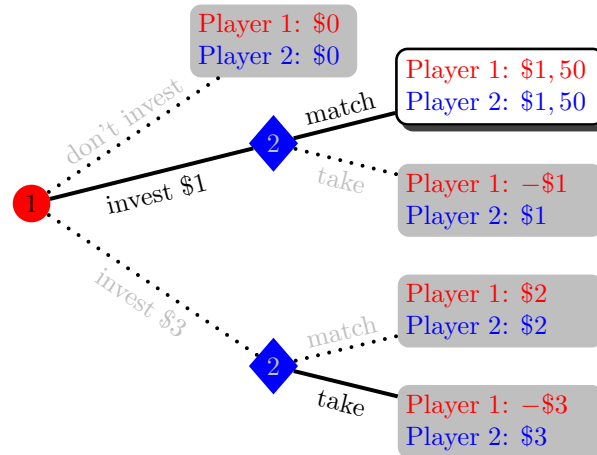


Figure I.6: Game tree for the Investment Game after second pruning

(Figure I.6).

Now what is left is a single strategy for each player; when we combine these, we see that when the players follow these strategies, Player 1 will invest \$1, and that Player 2 will match it.

The technique employed here is called *backward induction*, because it starts by considering the possible endings of the game and then reasoning back towards the beginning. This is our first important strategic principle, and is often summarized in the following slogan:

Look forward, reason backward!

Backward induction allows us to find equilibrium strategies. An equilibrium strategy is optimal, in the following sense: it is impossible for an individual player to obtain a better payoff by deviating from the equilibrium strategy. However, note that the optimal strategy in the investment game has the following features:

- Neither player gets the largest possible payoff. (The highest possible payoff for Player 1 is \$2, and for Player 2 it is \$3.)
- The collective payoff (the sum of both individual payoffs) isn't maximized either.

This from this example we see that an equilibrium may have some unsatisfactory characteristics, something which will turn out to be a recurrent theme in game theory.

In cases where the equilibrium in a game has counterintuitive properties, it is important to try and identify the causes. In the investment game, the payoff structure is to blame. If the second player was only allowed to take half the money from the pot, or if the rewards for investing were higher, then there

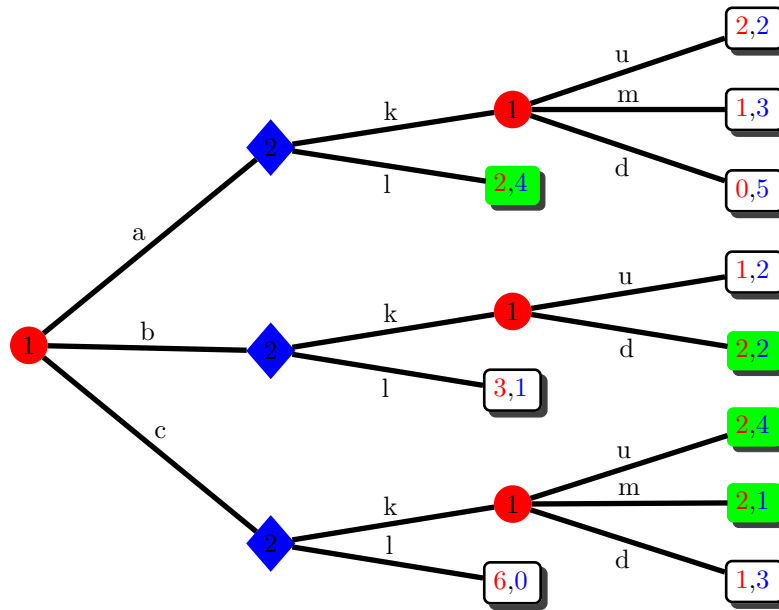


Figure I.7: Game tree with multiple equilibria

would be no incentive to take the money, and Player 2 would always invest, resulting in a maximal return for both.

It is tempting to think that a lack of communication between the players is also part of the problem. If the players could have a brief chat before the game, then they could agree to maximize their joint payoffs by both investing \$3, extracting the maximum possible from the host. However, as soon as Player 1 invests \$3, Player 2 has an incentive to change her mind and simply run with the money. Player 1 would be naive to think that Player 2 would accept \$2 instead of \$3.

Exercise 3. In the investment game, does it matter whether the players know each other or not? If you were player 1, would your strategy change if you knew your opponent was a computer? How would the game change if the payoffs were multiplied by a factor 100? Would you play differently against an opponent who you suspect doesn't know the rules of the game?

Exercise 4. The investment game illustrates how it is possible that the collective result is not always the best, even though each individual player acts in such a way as to maximize his or her personal gain. Are there real-life situations where this phenomenon occurs? What kind of measures can be taken in order to prevent this from happening?

We give one more practice example of finding rollback equilibria in a game tree. Consider the game tree in Figure I.7. To keep things in proportion, we indicate the payoffs simply by a pair of numbers; the first is the payoff for Player 1, and the second for Player 2.

In the a-branch the subbranches m and d will not be played by 1, so Player 2 chooses l. In the b-branch the subbranch u will not be played, so Player 2 chooses k. In the c-branch u and m are equally




Figure I.8: Game tree for the Monopoly Game

good for Player 1; Player 2 will choose k . Now at the beginning of the game Player 1 has several options leading to the same payoff (2), namely (a,l) , (b,k,d) , (c,k,u) and (c,k,m) . These four sequences are the rollback equilibria in this game, indicated in green.

I.1.4 MONOPOLY GAME

Our next miniature game is the *Monopoly game*. Imagine that you are about to open a new shoe store in a neighbourhood which currently has only one such store. Your start-up costs will be \$20,000, and the yearly demand for shoes is worth \$50,000 (meaning that customers are willing to buy \$50,000 worth of shoes per year). If you enter the market, one of two things will happen: either the other store will accept that their shoe monopoly has come to an end, and shares the market with you; in that case you both get half of the customers, and you both earn \$25,000 in sales. But the other store may also decide to fight you tooth and nail, in which case your revenue will only be \$10,000 and the other store will net \$30,000. (The difference between the \$50,000 and the \$10,000+\$30,000 arises from the fact that a price war means less revenue; moreover, your competitor will earn more than you, because he has loyal customers and knows the market better.)

The game tree is depicted in Figure I.8. The payoffs indicate the net profit or loss for each player in the first year. Note that in case of a price war, your net loss is \$10,000, because you've invested \$20,000 but earned only \$10,000 in revenue. In case the market is shared, you make \$25,000 in revenue, which translates to \$25,000-\$20,000 in profits.

Should you enter the market? In order to find out, we use backward induction and prune all sub-optimal branches. This leads to the tree in Figure I.9. We see that the remaining branch tells you not to enter. Ignoring this will cost you \$10,000.

Exercise 5. Now suppose that fighting will mean that you will earn nothing and that the other store will earn \$20,000. What will happen now? Draw the game tree for this variant of the game.

Exercise 6. As formulated, the game only takes into account what happens in the first year after you launch your company. Formulate a version where if you enter, you stay at least five years. (Revenue will not change for either party during those five years.)




Figure I.9: Game tree for the monopoly game, pruned

I.1.5 SUCKER BETS AND OTHER MISTAKES

Failure to look forward and reason backward may lead to various regrettable decisions, in particular the acceptance of “sucker bets”, by which we mean bets which don’t have any upside.

The following examples should get the point across:

- A classic mistake is to offer someone a bet that will only get accepted when you are certain to lose. For example, when I offer you a bet that you don’t know what the capital of Hungary is, you will gladly accept this bet when you know the answer, but when you don’t know it you will simply decline. Thus, this is a no-win proposition for me. (Of course, it is possible that you *think* that you know the answer, but that you have it wrong. This is my only hope to win.)
- Suppose you make a bid on an object of unknown value (to you, but not to the seller), for example, a second hand car. If the dealer is happy to sell at that price, it’s likely you’re paying too much. Similarly, when a company makes a takeover bid on another company without knowing the value of the company, then there is a genuine risk of the bid succeeding only when the company turns out to be worth less than the price offered.

This is closely related to the so-called *winner’s curse* in auctions. If your bid is accepted, that may well mean that you paid more than the other(s) think the item is worth.

- There is a classic example of a sucker bet in poker: making a large bet on the last street which will only get called by a better hand. For example, in No-Limit Holdem you hold $Q\heartsuit J\heartsuit$ on a $2\clubsuit 4\clubsuit J\spadesuit K\spadesuit A\heartsuit$ board against a single opponent, who checks. Here, you should usually² be happy to check behind. If you bet, then your opponent will likely call or raise with stronger hands (in which case you lose money), and will usually fold weaker hands (in which case you win the pot, just as when you would have checked down). What’s worse, by betting your opponent might even check-raise as a bluff, in which case you could be moved off the winning hand. In short, by betting here you will never win more than you can by checking, but you may lose more. Thus there is no upside to the bet.
- While not necessarily a sucker bet, a marriage proposal shares some of its characteristics. Indeed, before you make one, you are well-advised to think through the consequences of it being accepted!

²As always in poker, things are context- and player-dependent; there are always situations imaginable where general principles are overruled by other considerations.

This is particularly pertinent in social settings where marriage decisions are in part based on financial or societal considerations.

When both people involved are happy with a deal, for example when buyer and seller reach agreement over a price, or when two people walk down the aisle, it does not automatically mean that one of them is a sucker. However, it is important to realize that the mere fact that someone accepts a bet, a bid or an offer is valuable information in its own right!

I.1.6 NIM


The game NIM has many variations, but all revolve around the idea of players taking turns in removing tokens from one or more heaps; in the normal variant the person to take the last token wins, and in the *misère* variant the person to take the last token loses. In principle, one could draw game trees, but they would be very large, so we're not going to do that. Still, as in many games of pure strategy, we can benefit from the "look forward, reason backward" principle.

1. The first variation is relatively simple and is called 21flag. It was played in the tv show Survivor Thailand, Episode 6, as an immunity challenge; you can find it on youtube. In this variant, 21 flags are placed in a row, and players take turns removing 1,2 or 3 flags at a time. The player who takes the last flag wins. Can you formulate a winning strategy? Also, consider the *misère* variation where the player to take the last flag loses. What is the optimal strategy here?
2. Standard NIM is more complicated, but the game is solved, and there is an algorithm for optimal play. Here, one has several heaps of tokens, and each turn a player may remove as many tokens from one chosen heap as he/she wants. You can play this online at

<http://www.online-games-zone.com/pages/puzzle/nim-master.php>

and read more about the winning strategy on Wikipedia.

3. The third variation is called ZECK, also called *chomp*. This is played by taking an $n \times m$ grid of flags; each player chooses one of the flags and removes all the flags above and to the right of the chosen flag (including that flag itself). For example, in the 2×3 variant the following would be a legal sequence of moves (each time, the chosen flag is indicated):



The person who takes the last flag loses. This game is unsolved, in the sense that it is unknown what the winning strategy is for general values of n, m .

Exercise 7. Work out the winning strategy for the 2×3 version, the 3×3 version and the 3×4 version.

Exercise 8. Show that if we start with an $n \times n$ grid, then the first player has a winning strategy, no matter what value n has. Explicitly describe the winning strategy.

I.1.7 PIRATES

We end this section with a nice example of a game which illustrates how a seemingly reasonable but subtle set of rules and payoffs can lead to weird outcomes.

Five pirates find a treasure consisting of 100 gold pieces. They wish to divide this treasure according to pirate law and etiquette.

- First, the most senior pirate proposes a division. Pirate 1 is most senior, followed by 2, then 3, then 4 and then 5.
- Next, there is a vote on the proposal. If the proposal is approved by at least half of the pirates (the proposer's vote also counts) then the division is made accordingly. If the proposal is rejected, then the proposer is killed and the procedure is repeated.

The preferences of the pirates are as follows: first and foremost, they don't want to be killed. Next, they wish to get as much gold as possible. Finally, all other things being equal, they want to kill other pirates.

What will happen? The game tree is enormous, but we can still use a form of backward induction. Let's work on a simpler problem first, namely where there are only two pirates. Then we simply have a dictator game: the first pirate will propose a 100-0 division, and since the vote swings his way in case of ties, he gets what he proposes.

Now think of the case where there are three pirates, and where the oldest one proposes a division. As soon as one of the two younger pirates supports it, it will be accepted. The second-ranking pirate knows that if the proposal fails, then the oldest one will be killed, so that they will end up in the case with two pirates, where he will get everything. Thus, the second-ranking pirate hopes the proposal fails. The youngest pirate, however, can also see this coming, and therefore supports any proposal where he gets anything at all. Thus the senior pirate proposes a 99-0-1 division, and the junior pirate supports him. (Question: why can't he propose a 100-0-0 division?)

Next, consider the case of four pirates. Now the second pirate will vote against any proposal where he receives 99 or less, while the third one will vote in favor of any proposal giving him 1 or more. The senior pirate wants to survive, hence proposes a 99-0-1-0 division, which gets accepted.

Finally, the case with five pirates. Now a proposal needs two additional supporters, so it must be more attractive to at least two people than the 99-0-1-0 which remains if it is voted down. Thus the proposed division is 98-0-1-0-1, which is then supported by pirates 3 and 5.

Exercise 9. Check for yourself that any other division will either lead to the senior pirate's demise or result in a lower payoff.

Exercise 10. Can you predict what the outcome will be if there are 50 pirates dividing 100 coins?

Exercise 11. The solution to the pirates game is not immediately obvious and requires some careful thought. Suppose you were the senior pirate, and you were in doubt as to whether the other four pirates had a solid understanding of backward induction. Would you still make the 98-0-1-0-1 proposal? Why (not)?



I.2 PAYOFFS, UTILITY AND REAL LIFE

In the games we saw so far, the objective was to maximize the amount of money won, or to simply win the game. In many games, the stakes are more complicated. There are several possible reasons for this:

- First of all, money does not always capture all aspects of the value of certain goods: personal value may differ from monetary value. An item which is worthless to you may have sentimental value to me.
- Next, even in situations where winning money is the central objective, the prestige of winning may be an important motivating force for the players. Many of us have a tendency towards competitiveness and often place more value on beating our opponent(s) than on maximizing our monetary payoff.
- Certain commodities have value, but are hard to quantify in monetary terms. For example, we all value clean air and peace, but it is difficult to put an exact price on those things.
- In many real-life situations ethical considerations play a role (at least for some people). Also, in many games and sports there are (unwritten) conventions about fairness and sportsmanship, and most players find it important to adhere to those principles.
- Finally, it is important to recognize that often the game is not played in isolation, but can be part of a series of games, or a larger strategic setting. In those cases, one should consider not only the payoffs associated to the current game, but also those of the subsequent games.

I.2.1 UTILITY

From the above considerations, it is clear that we need to make sure that the payoffs we use in our representations of games allow us to include other considerations than just financial. In fact, what we want is to take all considerations into account which determine the desirability of the various outcomes to the individual players, because they all could influence a player's decision. The notion of *utility*, commonly used in economics and decision theory, serves that purpose.

Utility is the overall desirability of the various possible outcomes of a game to the various players. The unit of measurement for this quantity is called a *util*.

One may think of utils as tokens, or points which players try to win. By default, all considerations relevant to the players are taken into account. If a player truly values outcome A more than outcome B, this will be reflected in the payoffs. Usually we shall simply represent the payoffs using numbers, where a larger number means that the outcome is preferred.

The idea of introducing a unit of measurement which is more relevant than simple dollars won or lost goes back to the Swiss mathematician Nicolaus Bernoulli. He observed that the degree of satisfaction derived from winning a certain amount is not directly proportional with the amount won. Rather, it is a function of one's actual wealth and the relative gain or loss. Put in plain language, winning \$1,000 is more exciting for a poor student than for a millionaire.

The question arises how we can know what the actual utilities of the various players are. How can we determine utility a certain player associates to a certain outcome? It seems almost impossible to assign values to one's preferences, let alone those of others, and even more difficult to compare them. However, proponents of utility theory argue, while it is perhaps not feasible to directly measure utility, one can indirectly measure it through behavior. If we wish to determine certain utility values for Mr. F, we can design a clever series of questions which ask Mr. F to choose between certain lotteries or bets. The answers to these questions will then reveal the utility values, and we can increase the accuracy of this process by asking more and more such questions³.

Utility theory presupposes that our preferences satisfy certain technical conditions. These will not be discussed here (interested readers can find suggested material in the references), but we remark that experiments show that these assumptions are easily violated. Thus it should be kept in mind that utility is a theoretical concept, and that reality is more complicated. Several alternatives to utility theory have been proposed over the years, including the influential prospect theory by Tversky and Kahneman.

We will temporarily leave these issues aside now. For the purposes of game theory, we can simply assume that we are provided with the various utilities relevant to the game in question, and then start our analysis there. Thus the determination of these utilities is assumed to be taken care of already.

It is also important to keep in mind that for many games, we don't care about the precise values of the payoffs; all that matters is how they compare. For example, in the investment game it doesn't matter what the precise amounts are; we could give the first player \$0.50 for not investing, or we could give the second player a \$100 bonus when he takes the \$3 investment. This wouldn't disturb the relative order of the various preferences, and hence wouldn't influence the decisions. In the situation where we only care about the ordering of the preferences, we speak of *ordinal payoffs*. If the exact values matter, we speak of *cardinal payoffs*. It is usually simpler to work with ordinal payoffs, but cardinal payoffs can't always be avoided.

I.2.2 ULTIMATUM GAME

We now play a game which highlights some of the considerations about payoffs. A stack of 10 quarters is placed in the middle. Player 1 proposes a division of this stack between herself and Player 2. (Both players should get at least one quarter in the proposed division.) Next, Player 2 may accept or reject. If he accepts, then both players keep the money according to the proposed division. If he rejects it, neither receives anything.

First, let us have a look at the game tree in Figure I.10.

We can do backward induction. In each of the cases, we see that Player 2 gets more money when he accepts than when he rejects. Even in the case where Player 1 proposes a 9-1 split, Player 2 will accept

³It has been objected, however, that making the meaning of utility dependent on observable behavior introduces a certain circularity.

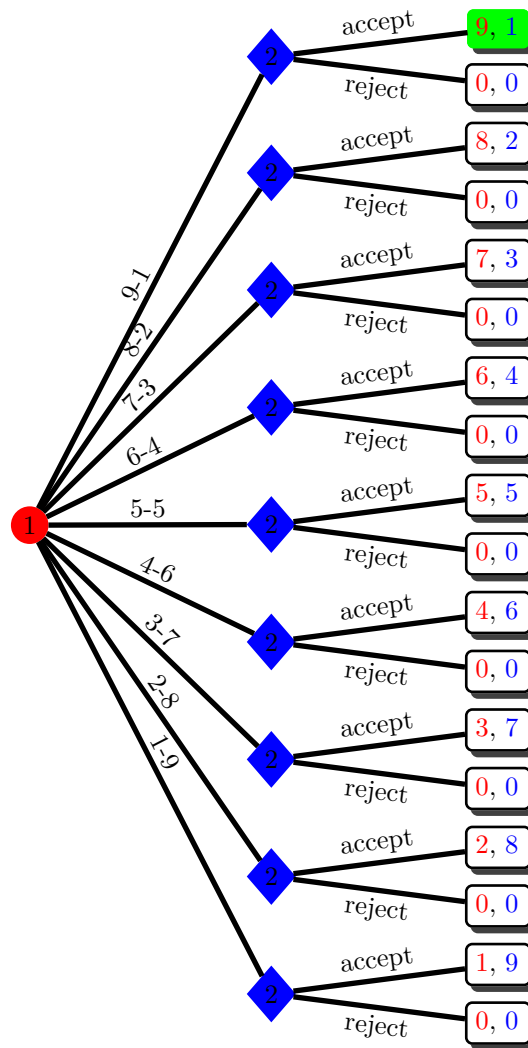


Figure I.10: Game tree for the Ultimatum Game

(perhaps grudgingly), because 1 quarter is better than no quarter at all. Thus Player 1 knows that any proposal she makes will get accepted, and hence selects the proposal which, when accepted, will earn her the most. This means that the proposed division is 9-1 and that it is accepted.

Exercise 12. Suppose that we allow the first player to offer a 10-0 division as well. Will she propose this division or will she stick to the 9-1 proposal?

Exercise 13. We can modify the ultimatum game in the following way: the first player proposes a division, and then the second player accepts or rejects. If she accepts, then the division is a done deal. But if she rejects, then the host of the game removes one of the quarters, and the game is played again (with the same roles). What is the optimal strategy for both players in this game?

Exercise 14. Another variation arises by allowing the first player to make a better offer after a refusal of the previous offer. The offer has to be more generous than the first offer. At any point, the first player can decide to quit and not make more offers. (Nobody gets paid anything then.) What is the optimal strategy here? How do you think people will play in practice here?

I.2.3 BEHAVIORAL GAME THEORY

The ultimatum game is interesting because it reveals a discrepancy between how we *should* play and how we actually play in practice. Experimental, or behavioral game theory tries to understand practical play. As such, it may be regarded as an *extension*, rather than an alternative to game theory in the classical sense. In experiments, subjects hardly ever play the optimal strategy: the first player tends to avoid 9-1 divisions in favor of more fair divisions, while the second player sometimes rejects offers. In fact, the typical offer is 6-4, and 5-5 offers are most frequent.

Since the game is simple and the optimal strategy is easily found, we cannot blame this on a lack of understanding on the side of the subjects in question. It is not surprising that it takes people a while to reason their way through the pirate game, but the ultimatum game is hardly an intellectual challenge. It can be made even simpler by letting the first player *dictate* a division and not giving the second player a veto at all (so that the second player simply has to be happy with whatever the first leaves him). This is called the *dictator game*, and even in this game we observe that the first player usually feels compelled to leave more than the bare minimum for the second. (Dictators tend to be a bit more greedy than the players in the ultimatum game though.)

Numerous experiments have been conducted, with the aim of understanding which factors contribute to this phenomenon. First of all, one might think that the players diverge from the optimal strategy because the stakes are insignificant, so that social- or fairness considerations get the upper hand in the decision. In order to test this, researchers have done experiments varying the stakes; they did not find significant differences in play between penny stakes and games where a hundred dollar is being divided. Even when highly meaningful amounts (in the order of magnitude of several months' pay) were on the line, people still showed an inclination towards fair decisions. One possible explanation is that when the payoffs are much larger, so is the fear of the offer being rejected.

However, experiments in different societies show that the decisions are influenced by social norms and values. Certain tribes in South America for example do not seem to regard a 7-3 division as selfish or unreasonable, while in other cultures sharing is much more engrained in the value system.

Other experiments involve varying the formulation of the rules or the setting of the experiment. In

one case, the two players first competed in another game, and were told that as a reward the winner would be the first player in the ultimatum game. After that, the ultimatum game was played as usual. It turns out, that the winners make more lopsided proposals than the average player. A possible explanation is that the players who have been told that their role in the game is a reward for their superior play in the preliminary game feel a sense of entitlement, and hence do not feel bad about taking a large share.

One has even investigated how testosterone levels affect play. It turns out that males with high testosterone are more likely to reject low offers than those with lower testosterone levels.

For an in-depth discussion of experiments with various games, the book *Behavioral Game Theory* by Colin Camerer is recommended.

I.2.4 THE CENTIPEDE

The Centipede is another game which nicely illustrates how we diverge from theoretically optimal play. The game owes its name to the fact that it usually has 100 rounds. In the first round of the game, the first player is presented with two amounts of money, say \$3 and \$1. She has two options: the first is to take the larger amount and giving the smaller one to her opponent, in which case the game ends. The second option is to ask for an increase in stakes. Then the host of the game increases the amounts to \$6 and \$2, and the same procedure is repeated, but now with the roles of the two players reversed. He may take the \$6 and give away the \$2 to the first player, or ask the stakes to be raised to \$9 and \$3 and turn it over to Player 1 again, and so on.

How will this play out? If we assume that the maximum number of rounds is 100, then there is a potential \$300 for the winner and \$100 for the other to be earned. However, the logic of backward induction ruins the dream of a happy ending again. For in round 99 the stakes are \$297 and \$99, and Player 2 (whose turn it is then) will definitely prefer the former over the latter. Thus the game will not make it to round 100. Similarly, in round 98 the stakes are \$294 and \$98, and Player 1 will take the \$294, knowing that if he turns it over to Player 2 he will only get \$99. Continuing this reasoning, we find that the game ends in round 1 with Player 1 taking the \$3.

How do people play in practice? Well, very few players stop in round 1. Most players try to cooperate with their opponent in an attempt to get a higher payoff, at least for a couple of rounds. An interesting finding is that chess players (who are familiar with the principle of backward induction) tend to end the game earlier than average, and that very strong players end the game at the first possible opportunity.

An interesting observation is that the Centipede game is a model for certain behavior among criminals. When criminals engage in certain deals, for example drug sales, there is always the risk that one party will try to double cross the other. Usually, to build trust, transactions of this sort start out small, but gradually increase over time. At some point, the temptation for one party to betray the other and take everything will become too big...

I.2.5 WHY DO WE PLAY?

We have seen that our decisions in games can be highly dependent on non-monetary factors. It is always a good idea to ask yourself what you really wish to achieve when playing a certain game. In poker, there is a category of recreational players who play poker for the same reason they play roulette,

craps, blackjack or bingo: they regard it as a form of entertainment, and they value the social component of the game. These are the people who are likely to have a monthly budget for playing the game, and who are willing to lose some money in return for an enjoyable night of play.

Then there is a category of people who will say that they want to make as much money as possible, and that this is their sole objective when playing. However, closer inspection reveals that there are several phenomena which contradict that.

- First of all, the game takes place in a social context. While we may sit down at a table with the intention of winning money, we also hope to have a good time. We care about what others think about us; some of us want to be liked, others feared, yet others want to be seen as smart or good at the game. Each of these attitudes will have some influence on our decisions.
- Also, many players tend to have certain emotional attitudes towards their opponents. If they perceive a particular opponent as less deserving to win, for example because he has been excessively lucky, or has been bullying the table, they will often start caring more about beating that particular player than about winning money. It is very common to see players get overly focused on “putting someone in his place”. Your decision to target a certain player should be based on the fact that you think you can exploit weaknesses in his or her game, not on whether you think he or she deserves to win.
- In the other direction, some players tend to feel compassion towards unlucky and weaker players, especially when they have been running them over. As a consequence, they will slow down and play less aggressively against them. If this is a conscious decision to keep your opponent at your table then this is fine; if it is merely because you feel bad for him, then it’s not.
- Another common observation, due to Mike Caro, is that people tend not to bluff or get out of line after getting lucky in a big pot⁴. Usually, it is our instinctive response not to push our luck. While the flow of the game (or table dynamics, as it is also called) should be taken into account when making decisions, the mere fact that you won the hand before in a lucky way should not in itself change your strategy, even temporarily.
- A costly weakness for many players is their ego. They want to prove that they are superior, and refuse to leave the table when they are outclassed or when the circumstances are otherwise unfavourable. The desire to “get even” is also detrimental.
- Finally, the biggest problem of all is self-control. Very few players can honestly say that they have never gone on tilt, and financial losses are almost invariably the result. Knowing when to call it a day before you get frustrated is one of the most important skills to have as a player.

More detailed insights into how your emotions affect your decision making at the tables can be found in Alan Schoonmaker’s *The Psychology of Poker*.

⁴Only strong players will turn this around on you if they think you’re a Mike Caro fan.

I.2.6 RATIONALITY

In decision theory and game theory, we usually stipulate that players try to make decisions which are in accordance with their preferences, and are not willingly giving up utility. In short, they always try to get the largest possible payoff.

A person is said to be *rational* when he/she acts in a way which maximizes his/her (expected) utility.

The word *expected* here refers to the notion of expectation in probability, which we will investigate later. However, for now it suffices to identify rationality with making decisions in such a way that it leads to the highest payoff.

This notion is quite different from the colloquial meaning of the word (according to the dictionary, rationality means the quality or state of being agreeable to reason or understanding). In particular, rationality in the game-theoretic sense implies that individuals are always capable of finding the best possible strategy through analysing the situation, no matter how complicated.

Normative game theory describes how we should play, and assumes perfect rationality of all players. This is a very strong, and usually unrealistic assumption. Games between perfectly rational players should be regarded as a mathematical abstraction, of which “real life”-games are an approximation, much in the same way as the mathematical concepts of a perfectly straight line and of a circle are abstractions.

Questions to think about:

1. Usually, when scientific theories fall short when it comes to making accurate predictions, this is regarded as a serious problem. How worrisome is it that game theory predicts outcomes which sometimes differ significantly from observed play (as in the ultimatum game)?
2. Does the definition of rationality given above meet our intuitive (i.e. pre-theoretic) ideas about what rational behavior is?
3. Is rationality an all-or-nothing property? Or is it possible to be somewhat, but not completely rational? How would you classify a chess player who is in the top 1% of the world but doesn't manage to win the world championship?
4. Suppose that I am playing a poker game, and that at the end of the hand I have two moves to choose from. (For example, I have a choice of calling or folding a marginal hand which can only beat a bluff.) I decide to flip a coin and leave the decision to chance. Is this behavior irrational? Why (not)?



I.3 LEVELS OF THINKING

We have seen that in determining the best strategy in a game one has to think about what the other players are going to do. In a sequential game, this means reasoning about the possible responses to your moves first. However, not all games are sequential, and even if they are there may be complicating factors (such as information asymmetries or uncertainty) which make a straightforward application of backward induction impossible. For example, one cannot use backward induction to determine the correct strategy in Texas Holdem, even though this is a sequential game.

In this section we consider games involving two players who make their decisions *simultaneously*, without knowing what the other player does⁵. In determining what strategy to use in such games we tend to get lost in the “levels of thinking”-problem: we try to decide what to do, but also should think of what the other player is going to do. The other player’s decision is based on what he thinks we will do, so we have to think about that as well, and so on. We need to break through the “I know that you know that I know...”-spiral, and this is what game theory helps doing.

Poker players are acutely aware of the levels of thinking: they speak of first, second or third level players, and of leveling wars. A player is said to be on the first level when she only thinks about her own hand. On the second level, she recognizes that the value of her hand relative to her opponent’s hand is more important than the absolute value, and hence she gives thought to the possible holdings of her opponent. On level three, she not only thinks about the opponent’s hand, but also thinks about what the opponent thinks she has. (Does my opponent’s raise mean that he thinks I have a weak hand?) One level higher, she thinks about what the opponent thinks she is thinking about him. (Does he think I’m raising because I think he has a weak hand?) And so on.

If you identify someone as an n -th level player, then the way to beat him is by moving to the $n + 1$ -th level. However, you don’t want to go higher than that: if you’re two levels above your opponent then you’re bound to misinterpret his actions. For example, if your opponent is a level 1 player, and he makes a large bet, then there is no point in trying to figure out whether he might think that you think that he can’t be bluffing. He’s not capable of that kind of thought process. His bet simply means that he likes his hand, and you should play accordingly.

I.3.1 SIMULTANEOUS GAMES

The best-known simultaneous game is the so-called *Prisoner’s Dilemma*. This game was invented in 1950 by Merrill Flood and Melvin Dresher, both employed by RAND (see below). The background story is due to Albert Tucker, a Canadian-born mathematician who served as Nash’s Ph.D. supervisor in Princeton. While this story adds color to the game, it also adds to the risk of bringing ethical or emotional aspects into the picture. One has to assume that the two criminals in the story are not hampered by such considerations, and that their sole objective is to minimize their own time in prison. There are nice alternative stories to illustrate the game, and you should try to come up with one yourself.

Two criminals, let’s call them criminal I and criminal II, have been arrested for a serious crime, but

⁵In principle, we can also have such games with more than 2 players, but most of the important features already come to light in 2 player games, and the exposition is much simpler in that case.

the police doesn't have enough evidence to convict them; they can only put them away for 3 years based on some petty crimes. The criminals are being held in separate cells and cannot communicate with each other. The police makes each of them the following offer:

"You are both getting the following choice: testify against the other or keep quiet. If you testify and the other doesn't, then you are free to go, while the other gets 20 years in prison. If you don't testify, but the other does, then it's the other way around, and you get 20 years and he goes free. If both of you testify you'll each get 10 years. Finally, if neither of you chooses to testify, then you'll both get three years."

It is common to refer to the option/strategy of ratting out the other person as *defecting* and to keeping quiet as *cooperating*. We first represent the game in a 2x2 matrix (Figure I.11) called the *payoff matrix* for the game. The numbers in the matrix represent prison time; in order to get the genuine payoffs we should take their negatives.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	3y / 3y	0y / 20y
	Defect	20y / 0y	10y / 10y

Figure I.11: Payoff Matrix for the Prisoner's Dilemma

What will happen? In deciding what to do, criminal I has to think about what criminal II might do, but he knows that criminal II is thinking about what he is going to do: the decisions are interdependent. The way out of this leveling problem is as follows: place yourself in the shoes of criminal I. We can reason as follows:

"I'm not sure whether criminal II is going to cooperate or defect. So let's think through each of these two scenarios. First, if he cooperates, then I'm better off defecting, because then I go free (if I would cooperate we'd both get 3 years). Second, if he defects, then I'm also better off defecting, because then I get 10 years as opposed to the 20 I'd get when cooperating. *Thus in either case it's better for me to defect.*"

Of course, the situation is perfectly symmetric, so criminal II goes through the exact same line of reasoning, and the upshot is that both end up defecting.

First, note that even though both players act in their own best interest, neither gets the best possible outcome. The joint interests of the players are not maximized either: together they spend 20 years in

jail, while they could have limited their joint time to 6 years. These are observations similar to the ones we made with regard to the investment game.

Another observation is that allowing the two criminals a couple of minutes of communication before making their decisions will not change the outcome: as soon as the two are back alone making their decision they fall back on the same reasoning as before.

The Prisoner's Dilemma is called a dilemma because the outcome is counterintuitive and because it seems so obvious that intelligent players will cooperate as to minimize their sentences. You might think that it is an artificial example, and that in real life such absurdities do not occur. But you would be very mistaken. Examples of prisoner's dilemmas abound, and some of the most pressing issues of today's world can be understood in these terms.

For example, consider the problem of overfishing. As is well-known, several species of fish such as bluefin tuna and atlantic halibut are seriously endangered because they are in such high demand. Think about this problem through the eyes of a fisherman. Clearly, it is not in your interest to see tuna become extinct. On the contrary, if quota would regulate fishing in such a way that the tuna population could grow back to the size where everybody can catch a moderate amount without endangering the population, that would be highly valuable. However, every day when you leave the dock, you face the choice of trying to catch a tuna or catching some sardines or other cheaper fish. Tuna will earn you many times more than sardines; moreover, one tuna less is not going to make a significant difference to the population. On top of that, you can figure that other fishermen go through the same line of thought and hence will probably go for tuna; it wouldn't be fair that they all earn a lot more than you while you try to do the right thing now would it? Soon there will be no tuna left, so you might as well grab one while you can. Thus all fishermen go out to catch tuna, and our last bite of sushi is on its way.

The same reasoning applies to all situations where solving a collective problem requires individual sacrifices: pollution, traffic jams, global warming, and many other social or environmental issues could be solved provided everyone would contribute. However, each individual knows that his or her actions have negligible impact on the overall state of affairs; the reward for defecting is often substantial. And hence the vast majority of people defects. Such situations are usually called *tragedies of the commons*, after the example of the pastures in Oxford where the public was allowed to let its cattle graze. Because of overgrazing, there soon was very little grass left on the commons, but even though the farmers realized that they were depleting it, for each of them it was better to try to get their cow to feed on the last bit of grass before it was all gone than to not feed their cow at all.

A closely related kind of problem is called the *free rider dilemma*. You are about to take the bus; a ticket costs \$3.25, which you don't feel like paying⁶. It is raining so a ticket control is highly unlikely. So even though you don't have a pass, you jump in via the back door and take a free ride. You know that your transgression is not making a significant difference to the bus company's financial situation; the \$3.25 means more to you than to them. Even though this is rational behavior on your behalf (provided you are right that the risk of getting caught is negligible), we do get into trouble when we all reason that way, for then the bus company is sure to go broke. Ideally, you would like to be the only person defecting, and have the rest of the population support your unethical behavior.

⁶I can't disagree with you, because it is way too much for the service you get.

I.3.2 SOME CENTRAL CONCEPTS IN SIMULTANEOUS GAMES

In our initial analysis of the Prisoner's Dilemma, we used the following argument: for each player, defecting is better than cooperating no matter what the other player does, and therefore each player (provided they care only about minimizing their own sentence) will defect. We will now take the key aspect of this idea and make it a bit more precise, so that it can be applied to other games as well. In order to do so, we first introduce some terminology concerning simultaneous games which makes discussion easier.

Definition I.3.1 (Strategy, strategy profile).

1. A (*pure*) *strategy*⁷ for a player consists of a specification of a valid move for each possible situation in which he can find himself.
2. A *strategy profile* is a collection of strategies, one for each player in the game.

Thus in the Prisoner's Dilemma, each player has two possible strategies, namely cooperate and defect. As a result, there are $2 \times 2 = 4$ strategy profiles, namely

- Both cooperate. This we denote (cooperate, cooperate), or (c,c).
- Player 1 cooperates, Player 2 defects. This is denoted (cooperate, defect), or (c,d).
- Player 1 defects, Player 2 cooperates. Denote this by (defect, cooperate), or (d,c).
- Both defect. We denote this profile by (defect, defect), or simply (d,d).

To specify a strategy profile thus means to specify one available strategy for each of the players. Given that information, the payoff matrix tells us what the payoffs for each of the players will be when they play these strategies.

Example I.3.2. Consider the 3×4 matrix in Figure I.12. Player 1 has three strategies, namely a,b and c. Player 2 has four, namely p,q,r and s. Thus there are 12 strategy profiles, namely:

$$(a,p),(a,q),(a,r),(a,s), (b,p),(b,q),(b,r),(b,s), (c,p),(c,q),(c,r),(c,s)$$

Each of these profiles corresponds to a cell in the matrix. For example, the strategy profile (b,s) corresponds to the second row, fourth column of the matrix, and hence we see that if the players play this profile, Player 1 gets a payoff of 1, and Player 2 gets a payoff of 2.

The question of course is: what are good strategies, and how can we find them? The following is a useful aid:

Definition I.3.3 (Dominated strategy).

⁷Later we will also encounter mixed strategies, which include an element of randomization.

1. A strategy P for a given player *weakly dominates* another strategy Q for that player when it yields at least as high a payoff as Q , no matter what strategies the opponents play.
2. A strategy P for a given player (*strictly*) *dominates* another strategy Q for that player when it yields a strictly higher payoff than Q no matter what strategies the opponents play. If a strategy P dominates every other strategy, then P is *dominant*.
3. A strategy P for a given player is (*strictly*) *dominated* when every other strategy strictly dominates P .

For example: in the Prisoner's Dilemma, defection is the dominant strategy for both players.

Here's another example:

		Player 2	
		A	B
Player 1	C	3 / 3	0 / 7
	D	10 / 5	4 / 2

In this game, Player 2 has a dominant strategy, namely A. Why? Because no matter what Player 1 chooses to do, Player 2 always gets a higher payoff by choosing A than by choosing B. Also, strategy B is dominated (if there are only two strategies for a player, and one is dominant, then the other must automatically be dominated). Player 1 does not have any dominant or dominated strategies.

From now on, whenever you encounter a simultaneous game, always look for dominated and dominant strategies. The reason that these strategies are so important is that a rational player will never prefer a dominated strategy to a dominating one: playing a (weakly) dominated strategy would mean that the player is not maximizing his payoffs. This fact is highly useful when analysing simultaneous games. We formulate this as a strategic principle:

Rational players never play a dominated strategy, and always play a dominant strategy when such is available.

This means that when we wish to analyse a game, we can disregard all dominated strategies, because they will never be played. As an example, consider the following game, in which Player 1 has three strategies and Player 2 has four:

The first player doesn't have any dominated strategies. However, for the second player, strategy q is dominated. Therefore we remove the corresponding column from the payoff matrix. The result is the following 3×3 matrix:

		Player 2			
		p	q	r	s
Player 1	a	0 1	-2 3	-1 -1	-1 0
	b	2 2	0 2	1 3	2 1
	c	1 3	-1 0	0 1	2 2

Figure I.12: A 3x4 game

		Player 2			
		p	q	r	s
Player 1	a	0 1	-2 3	-1 -1	-1 0
	b	2 2	0 2	1 3	2 1
	c	1 3	-1 0	0 1	2 2

However, there is no reason to stop here! We have removed one dominated strategy, and now we can look if there are more. In this case, strategy a is a dominated strategy for Player 1. Removing it gives the following:

		Player 2			
		p	q	r	s
Player 1	a	0 1	-2 3	-1 -1	-1 0
	b	2 2	0 2	1 3	2 1
	c	1 3	-1 0	0 1	2 2

Again we look in the new matrix for dominated strategies; this time r is dominated:

		Player 2			
		p	q	r	s
Player 1	a	0 1	-2 3	-1 -1	-1 0
	b	2 2	0 2	1 3	2 1
	c	1 3	-1 0	0 1	2 2

Next, remove dominated strategy b:

		Player 2			
		p	q	r	s
Player 1	a	0 1	-2 3	-1 -1	-1 0
	b	2 2	0 2	1 3	2 1
	c	1 3	-1 0	0 1	2 2

		Player 2			
		p	q	r	s
Player 1	a	0 1	-2 3	-1 -1	-1 0
	b	2 2	0 2	1 3	2 1
	c	1 3	-1 0	0 1	2 2

Figure I.13: Result of IEDS

Finally, we see that now p is dominated. When we remove p, what remains is the single entry (c,s). (We will see later that this is an equilibrium.)

Note that during this process, which we call *Iterated Elimination of Dominated Strategies* (IEDS), strategies which were initially not dominated became dominated after the deletion of some other strategies. For example, in the original matrix, strategy a is not dominated, but it became dominated after strategy q was deleted.

The IEDS procedure is an analogue for simultaneous games of backward induction for sequential games. Not all games are fully solvable by means of IEDS; indeed, some games don't have any dominated strategies; also, the procedure may stop before reaching a single entry in the game matrix. Still, one should always try to simplify games by applying the procedure as far as possible.

Iterated Elimination of Dominated Strategies (IEDS): keep removing dominated strategies until there are no such strategies left.

I.3.3 CHICKEN

In the 1955 movie *Rebel without a Cause* (starring James Dean) two cars are driving towards a cliff. The driver to jump out first is the chicken. In a cheaper but no less exciting variation two cars are speeding towards each other on a narrow road. The first to swerve is the chicken. Figure I.14 shows how we model this in a matrix.

Obviously, both players prefer not to die in a heads-on collision, which is why the upper-left corner is the lowest payoff for both. Both players swerving is also not what they hope for, but it's still better

		Player 2	
		Straight	Swerve
Player 1	Straight	0, 0	2, 10
	Swerve	10, 2	5, 5

Figure I.14: Payoff Matrix for Chicken

than a crash. Above all, each player prefers the other to swerve while driving straight himself. And they prefer being a living chicken to being a dead hero too.

Questions to consider:

1. Does any player have a dominant strategy here? Are there dominated strategies?
2. What kind of things could one do to improve one's chances in this game?
3. Can you think of other situations which can be modeled in this way?

We need a bit more theory to predict what perfectly rational drivers do in this situation. For now, notice that even though no dominant strategy exists, the outcomes where one player swerves and the other goes straight have a special feature: after learning what the other player chose, neither player has any regrets about his own decision!

I.3.4 COORDINATION GAME

A couple has planned to go out for dinner (Thai food) but they have forgotten to decide on a restaurant. They both like Khao Thai and Som Tum. Without being able to communicate they have to decide where to go. They are indifferent where they end up, as long as they're dining together. Figure I.15 shows the matrix for this game.

Again, there is no dominant strategy, but there are two outcomes where neither player has any regrets. A few points to notice:

1. Unlike in the previous game, the two players are not playing *against* each other, they are trying to work together.
2. In the Prisoner's Dilemma, or Chicken, communication does not help, but in the coordination game it certainly would.

		Player 2	
		KhaoThai	SomTum
Player 1	KhaoThai	2, 2	0, 0
	SomTum	0, 0	2, 2

Figure I.15: Payoff Matrix for the Pure Coordination Game

- Because of the symmetry of the game, both strategies are equally good, and there are no grounds for playing one instead of the other.

As presented, the coordination game is a bit silly, in that it seems that we cannot do better than flip a coin. Indeed, as we shall later prove, flipping a coin and choosing accordingly is a perfectly fine strategy in this game! Nevertheless, even (or especially!) without any knowledge of game theory most couples have a successrate higher than 50% when playing this game. So again it seems as if theory and practice are wildly diverging here.

2005 Nobel Prize in Economics winner Thomas Schelling recognized that other factors come into play in games such as the coordination game, and devised an experiment to test this. Pairs of mutual strangers (who did not have any information about each other except for that they were playing the same game) were dropped in different parts of New York City early in the morning, with the instruction to find their partner. They were remarkably successful; almost all participants went to one of the main landmarks (Empire State Building, Times Square, Grand Central Station), and decided that noon was the obvious meeting time⁸.

The experiment shows that in coordination games, people do not choose at random (that would lead to an astronomically small chance of success in a big city) but converge towards *focal points*, in this case landmarks. In honor of the inventor, we now refer to these as *Schelling points*.

Other examples:

- When two people lose each other, say in a busy shopping mall, they may try to get reunited at an information booth, or a lost-and-found.
- Price negotiations are likely to converge to a round number.
- Land divisions usually follow some obvious natural features, such as mountain ranges or rivers.

⁸In the tv show "Life, The Game" (ABC, partly designed by Barry Nalebuff), this experiment was repeated.

		Player 2	
		KhaoThai	SomTum
Player 1	KhaoThai	1 2	0 0
	SomTum	0 0	2 1

Figure I.16: Payoff Matrix for the Battle of the Sexes

I.3.5 BATTLE OF THE SEXES

It should be clear that coordination problems are ubiquitous; however, it is unlikely in practice that both players have exactly the same preferences. Usually, while the overall aim is still to achieve the main goal of successful coordination, the players are not in complete agreement on how this should be realized. The Battle of the Sexes game captures the essence of such situations. The game has the same background story as the coordination game, except for that now Player 1 slightly prefers Khao Thai, while Player 2 slightly prefers Som Tum (but they still both prefer to dine together rather than alone). See Figure I.16 for the corresponding payoff matrix. This is only a small variation, but it has important consequences. Now each player has an incentive to manipulate the situation to his/her advantage: the game is a combination of pure coordination and competition. Again, games like these lead to the consideration of *strategic moves* such as the ones Schelling was interested in. Particularly interesting is the question of how communication would influence the situation. While for the pure coordination game a brief phone call would solve the problem, this is no longer the case here because both players would try to convince the other to come to the restaurant of their choice, possibly with a stalemate as a result (especially if the preferences are strong).

Some more sophisticated tricks are needed here. For example, in the above situation one would like to send a message without being able to receive a reply (“I’m going to Som Tum right now - sorry, my phone is out of battery, bye!”).

Exercise 15. In what other ways could you change the Battle of the Sexes to your advantage or increase the chances of a favourable outcome?

There are many variations on the backstory of this game; a classic version is the *Stag Hunt* by Rousseau, in which two cavemen can join forces to catch a big prey, but risk ending up with nothing if they don’t coordinate this.

I.3.6 NASH EQUILIBRIA

We now turn to one of the central concepts in game theory, namely that of a Nash Equilibrium for simultaneous games.

A *Nash Equilibrium* (NE) is a strategy profile in which neither player can get a higher (expected) payoff by unilaterally changing his/her strategy.

In the game of Chicken, there are two Nash Equilibria: (swerve, straight) and (straight, swerve): in both cases, neither player would like to change his strategy if he knew what the other player was going to do. The other two strategy profiles, namely (swerve, swerve) and (straight, straight) are not Nash Equilibria: if you swerve and you see that your opponent also swerves you wish you would have gone straight, because that would have given you a higher payoff. Similarly, if you go straight and your opponent does the same, then you wish (if you live to tell) that you had swerved.

In the Coordination game, as well as in the Battle of the Sexes, there are also two Nash Equilibria (check this for yourself). It is also clear that the two strategy profiles where the two players end up at different restaurants are not Nash Equilibria, because in that situation both players would benefit from changing his/her strategy after learning about the other player's move.

Finally, we come back to the Prisoner's Dilemma. What are the Nash Equilibria in this game? First, both cases where one player defects and the other cooperates are not Nash Equilibria; clearly the cooperating person would like to change his answer after learning of the betrayal of the other! The case where both defect, however, is a Nash Equilibrium: in that situation neither player has any reason to change from defection to cooperation, for that would increase his sentence by 10 years. Finally, what about the case where both cooperate? Unfortunately, that is not a Nash Equilibrium: if one of the two criminals would learn that the other was cooperating, he would like to change his answer and defect, so that he would go free.

We make the following observations and comments about Nash Equilibria:

1. In a NE, each player plays his/her *best response* to the other player's best response. The idea of best response to another player's move is a useful one⁹, which we'll use later to determine the Nash Equilibria in games more systematically.
2. A Nash Equilibrium is self-reinforcing, in the sense that none of the players have any incentive to deviate from it.
3. Another side of the same coin is that neither player has any regrets about his/her choice of strategy. Of course, this is not a formal concept, but it is a good heuristic: if, after learning what your opponent did, you wish you'd done differently, then you're not in a NE! However, it is quite possible, as in the Prisoner's Dilemma, that one or both players regret the overall outcome of the game.
4. Be aware of the fact that games may have no NEs in pure strategies at all. For example, game in Figure I.17 is easily seen not to have any NEs. (Check this for yourself.) What type of game could this be? Can you think of other games where this is the case?

⁹However, the term can be a bit misleading, because it suggests a sequential move order; the way to think of it is as follows: I'm playing a best response to your strategy, when I'm maximizing my payoffs assuming that you've already made your move.

		Player 2	
		H	T
Player 1	H	1 -1	-1 1
	T	-1 1	1 -1

Figure I.17: A game without NEs in pure strategies



I.4 GAME THEORY IN CONTEXT

Now that we have had a taste of game theory and are starting to understand some of the basic principles of strategic thinking, we take a step back and have a look at some of the history of the subject. The reader is advised that what follows is a biased selection, and that the references offer a more detailed (and balanced) overview.

I.4.1 ORIGINS OF GAME THEORY

The founding father of game theory is usually considered to be John von Neumann (1903-1957), an exceptionally talented scientist¹⁰ who made profound contributions to almost all areas of mathematics, to physics and economics, who was one of the pioneers of computer technology, helped develop the atomic bomb and who, as a government consultant, played an important role in US politics.

Originally motivated by a desire to better understand the game of poker, he wrote the pioneering article *On the Theory of Parlor Games (Zur Theorie der Gesellschaftspiele, 1928)* in which he showed that for certain games there exist optimal strategies, a key result which goes by the name of the Minimax Theorem. It should be mentioned that the French mathematician Emile Borel wrote on *La Théorie du jeu* around 1921, and that there is considerable overlap with von Neumann's ideas. Borel also analysed a form of poker and considered the problem of bluffing. However, von Neumann's Minimax Theorem is considered the first significant result, and as such it gave the subject mathematical weight.

Von Neumann's ambition was much greater than just understanding when to bluff though, as he hoped to provide a solid mathematical basis for economics. In 1944 he joined forces with the Austrian economist Oskar Morgenstern, and together they wrote the classical text *Theory of Games and Economic Behavior*. This book, whose primary objective was to demonstrate that the theory of games would be the beginning of economics as an exact science, was a great success and brought the possible virtues of game theory to the attention of a wide audience. This was the beginning of a period (1945-1960) where

¹⁰Many have claimed that he had the best brain in the world.

economists, politicians and generals alike had high hopes that the discipline would help them better understand the strategic issues of the day.

The next important result was due to John Nash, who managed to overcome one of the serious limitations of Von Neumann's Minimax Theorem, namely that it applies only to zero-sum games. (Zero-sum games are games where one player's gain is another player's loss, so that the payoffs add up to zero.) Most games of practical interest are not zero-sum: for example, in a war, it is quite possible for both sides to lose heavily; similarly, both sides can win if a satisfactory diplomatic solution can be found. Nash introduced the solution concept which we now call Nash Equilibrium, and proved that every finite game admits at least one such equilibrium, provided players are allowed to employ mixed strategies. In 1994 he received the Nobel Prize in Economics for this work, joint with game theorists Reinhard Selten and John Harsanyi who refined and extended Nash's work.

The history of game theory is also closely related to that of international conflict and military strategy. The parallels between strategic games and warfare are old. In the classic book *The Art of War* by the Chinese military expert Sun Tzu (544-496 B.C.) various strategic principles are discussed, many of which are still applicable today¹¹. The concept of war as a strategic game became much more concrete with the invention of the so-called *Kriegspiel* (war game). This game, which is played on a miniature battlefield with toy soldiers, was hugely popular in Prussia: it was compulsory study material for military personnel, and was said to lie at the basis of some Prussian military victories. This prompted the US navy to begin their own study of the game, and eventually the game would be played in many of the institutes where research on game theory took place.

I.4.2 RAND AND THE NUCLEAR ARMS RACE

A great deal of work in game theory was done in the period immediately following World War II, at the RAND Corporation in Santa Monica (RAND stands for Research and Development). Because of the world-wide scale of warfare and the newly developed atomic bomb, military strategy had become so complex that military leaders found an increasing need for technological research as well as scientific studies regarding feasibility and possible outcomes of various strategies and tactics. Many scientists had contributed to the war effort, either by breaking enemy codes (Alan Turing is the most famous example), helping safely encrypt allied communication (Claude Shannon played an important role here) or by developing new weapons and strategic plans. Robert Oppenheimer and John Von Neumann, who studied mathematics together in Göttingen under the famous David Hilbert, worked on *Project Manhattan* in Los Alamos, which resulted in the atomic bomb. Von Neumann was also among those who helped determine the most successful bombing strategies, and helped select appropriate targets for the Little Boy and the Fat Man, the two A-bombs which destroyed Hiroshima and Nagasaki, respectively.

Towards the end of the war, however, most scientists went back to their universities to conduct their usual research, and many military leaders lamented the brain drain. In 1945, the Douglas Aircraft company sent Franklin Collbohm, a veteran pilot and engineer, to lobby for the creation of a military research centre. The air force decided to dedicate a substantial sum of money to the project; for the first three years of its existence, it was officially part of Douglas and answering to the air force, although its legal status was vague. After three years the organization separated from Douglas, and became a non-profit organization, still subsidized by the air force.

¹¹It is popular among poker players as well. Although a bit terse, it has some useful advice.

RAND corporation was more of a university than a corporation; its members had complete freedom and could do research on whichever topic they liked. Over the years a wide range of subjects was studied, but in the beginnings game theory was especially important, and around 1950 the vast majority of top game theorists either worked for RAND or were paid consultants. The list includes John von Neumann, John Nash, Kenneth Arrow, George Dantzig, Melvin Dresher, Anatol Rapoport, Lloyd Shapley, Martin Shubik, John Williams and R. Duncan Luce. By 1960, the organization employed about 500 full-time researchers and 300 consultants. At present, over 1600 people worldwide work for the organization. (For more information, consult the RAND website www.rand.org, or see the book *Prisoner's Dilemma*, by William Poundstone.)

As the arms race progressed (in 1949 the USSR successfully tested its first atomic bomb in Siberia; in 1950 the US had several hundreds of atomic bombs at its disposal, and the hydrogen bomb was under development) game theorists at RAND investigated increasingly bizarre and terrifying scenarios, often with cold-blooded results and recommendations to follow. For example, scientists at RAND asked how the government should allocate its resources after a hypothetical nuclear strike on US soil. They concluded that it would be most advantageous to the population to withhold support to the elderly and to the ill, because these would be a burden to the recovery of the country.

Meanwhile, the philosophy of World Government became increasingly popular, both among certain scientists and in military and government circles. The tenet was that nothing would be gained if both the US and the USSR would manage to build up an enormous arsenal, and hence that the best option was to force the USSR to surrender by threatening annihilation. Von Neumann was a fervent proponent of this idea, and was in fact in favor of a pre-emptive nuclear strike on the USSR in order for the US to obtain complete control over the situation.

This led to a perception of game theory as an overly theoretical, cynical and inhumane view on the world, and in the 1960s the initial euphoria over the discipline largely faded out.

I.4.3 SCHELLING'S WORK

One of the people who was disappointed in the applicability of game theory to the conflicts at hand was political economist Thomas Schelling. Schelling found that the developments in game theory at that time were too abstract, and that, even though game theory had been a success in several areas, it didn't take into account various aspects of real-life conflicts he was particularly interested in, such as strategies of deterrence.

In a conflict such as the Cold War, it is virtually impossible to come up with accurate estimates of various essential pieces of data: there are simply too many factors and unknowns. For example, it was clear that the USSR didn't want to be annihilated, but for most of their other preferences one could only guess. A more pragmatic issue was that both the US and the USSR had very limited information on the exact size of each other's nuclear arsenal, as well of other technical capabilities¹² Also, traditional game theory places a strong emphasis on rationality, and this is not a comfortable assumption to make in times of nuclear threats. Thus if the cold war was to be modelled as a game, it was one in which the US didn't know the payoffs nor the possible moves, and was playing against an opponent who was unlikely to comply with mathematical directives to begin with.

¹²Bizarrely, even in the US administration and military top very few people knew how many atomic bombs the US actually had, especially in the beginning.

One of Schelling's main contributions was to introduce a variety of new concepts to the abstract theory, with the aim of making it more applicable to real-life conflicts. His 1960 book *The Strategy of Conflict* systematically analyses several concepts such as *threats*, *promises*, *commitments*, *credibility*, *signaling* and *brinkmanship*. Together, these go under the name *strategic moves*: these are not moves within the game, but actions designed to alter the game to one's advantage. These have become recognized as important part of the theory of strategy. Together with Robert Aumann, Schelling received the 2005 Nobel Prize in Economics for his contributions¹³.

I.4.4 JOHN NASH

While John von Neumann put game theory on a solid mathematical footing and convinced the academic and the political world that it was a subject with many potential applications, it was John Forbes Nash (1928) who, in a 29-page Ph.D. thesis written at Princeton University, took the discipline to the next level by proving the existence of Nash Equilibria.

While von Neumann and Nash both were exceptional mathematicians, both connected to Princeton University and the RAND corporation, they were very different in many other regards. Where von Neumann grew up in Budapest, one of the intellectual centers of the world at the time, Nash childhood and highschool years were spent in Bluefield, West Virginia. Where von Neumann's family discussed poetry, art, psychology and financial affairs at the dinner table and had governesses to teach the children German and French, Nash's parents, while encouraging his intellectual development, were not overly sophisticated. Von Neumann had a brain faster than anyone else, and tended to lose interest in problems unless he saw the solution right away. Nash was a quiet, slow thinker who enjoyed spending night after night working on seemingly unsolvable problems and who was eager to prove to the world that he could do things others couldn't. And while von Neumann was a sociable, generally well-liked man who loved to throw and attend parties with fellow brilliant scholars, it was mainly because of the awe and respect that Nash' superior mental faculties commanded that others accepted him in their social circles.

During his early career at MIT, where he was a C.L.E. Moore instructor, Nash worked on some of the most difficult open problems in pure mathematics and achieved highly original and powerful results. Somewhat paradoxically, these results are of much greater mathematical depth than his work on non-cooperative games and bargaining¹⁴ for which he received the Nobel Prize (there is no Nobel Prize for mathematics, but to his dismay Nash did not receive its mathematical analogue, the Fields Medal, either). However, the applications of Nash' game theoretic results are so widespread and useful that they have dramatically changed our outlook on many disciplines.

Nash's spectacular achievements and dramatic personal life, including a decades-long battle with mental illness followed by a miraculous recovery, have been chronicled in the book *A Beautiful Mind*, by Sylvia Nasar. The movie, while worth watching, often veers quite far from historical fact and is no substitute for the book, which contains a wealth of fascinating information about Nash' life, the people around him, and the intellectual communities to which he belonged. Today, at age 83, Nash is a senior professor at Princeton and still makes public appearances.

¹³His Nobel Prize award lecture can be watched on <http://nobelprize.org/>.

¹⁴For those with an interest in mathematics: the proof of the existence of equilibria is a clever, but not exceptionally difficult application of Brouwer's fixed point theorem. Von Neumann managed to upset Nash by calling it a triviality.

CHAPTER II

WHAT ARE THE CHANCES?

Very few things are certain in life, and almost all of the decisions we make, whether in games or in other situations, involve uncertainty. Whether that uncertainty comes from a lack of information, from an element of randomness such as the throw of a die, or from processes which are too complicated to predict by means of exact calculation, we need to have tools for correct reasoning and strategic decision making that takes this uncertainty into account.

In this chapter we begin the study of chance, in the form of probability theory. Curiously enough, it was not until the 16th century that the first steps towards a systematic understanding of randomness were made, even though people have gambled since the dawn of civilization. The development of the subject was driven to a large extent by problems directly stemming from gambling, and we will trace through of some of the historical highlights.

In order to determine whether a bet is favorable, or whether one decision is better than another, we introduce the central concept of expected value. This notion is indispensable for the analysis of games of chance, but widely applicable outside the casino as well. Our main illustration is poker in its various forms. Correct application of expected value is the basis of mathematically sound play, and without a firm grasp on this matter it is virtually impossible to be a long term winner. Because it is often too time consuming to do an exact calculation (especially during actual play) we also present some heuristics which cover some situations which come up a lot, such as deciding whether to call a possible bluff or whether to make one.

As situations get more complicated, it becomes more and more difficult to determine the relevant probabilities with elementary means. Fortunately the laws of probability theory tell us how to break down complex problems into simpler ones. Using these laws correctly often boils down to organizing the problem in a systematic way, making use of the mathematical concepts of combinations and permutations. Poker players make heavy use of this when they try to deduce the likelihood of their opponent having certain hands.

We shall see, however, that probabilities often behave in a counterintuitive manner, and that this is the cause of many grave errors in decision making. The most famous of these is called the *Gambler's Fallacy*, but that is by no means the only mistake we are prone to make when reasoning about uncertainty. We will look at several examples where misjudgment of probabilities and incorrect reasoning about them have serious consequences. The biases and fallacies to which we are prone have been studied in great detail by cognitive psychology, and the main insights from that discipline will be presented here in order to increase our understanding and awareness of the weakest links in our chains of reasoning.

Finally, we will look at the problem of using new information. How should we adapt the probabilities we want to use in our decision making when new information comes to light? In games, every action by your opponents reveals some information. When someone obtains the results from a blood test, he or she gains new information, but that information may not be 100% reliable. And when a new evidence is introduced in a legal case, that may change the likelihood that the defendant is guilty. Bayes' Rule tells us how to adjust our probabilities in the face of new evidence. While mathematically simple enough, its applications can be tricky and have deceived even the most sophisticated thinkers.

To further motivate the explorations to come, here are a couple of classic trick questions involving uncertainty. As the chapter progresses, we shall find out how to answer them.

1. **Monty Hall Problem.** This is perhaps the most famous problem in probability, originating from the quiz show *Let's Make a Deal*, and named after the host Monty Hall. The winner of the quiz faces three doors; behind two of them are goats, and behind one of them is a brand new car. You get to pick a door. After you made your choice, Monty Hall opens one of the remaining doors, which turns out to have a goat behind it, and offers you the option of switching to the other closed door. Do you want to switch or not? (Pretend for the sake of discussion that you are more interested in the car than the goat.)
2. **Boy or Girl?** This question comes in three variants, in increasing order of trickiness. Suppose we consider a random family with two children. You meet the oldest of the two, who turns out to be a girl. What is the probability that the other is a girl as well?

Next, suppose that you meet one of the two children, but that you don't know whether it is the older or the younger one. Again, it turns out to be a girl. What is the probability that the other is a girl as well?

Finally, suppose that you meet one of the two (again you don't know whether it is the older or the younger), but now she introduces herself as Kimberly. Does this change your answer from the previous scenario?
3. **Birthday Problem.** You're in a room with 40 people. What are the odds that two of them share a birthday? How many people would you need for this chance to be 50%? How many would you need to have a 99% chance?
4. **Doping Test.** You are organizing a sports tournament, and you wish to make sure that none of the athletes are gaining an unfair advantage by using doping. You have available a blood test which is 95% reliable. This means that when you test 100 athletes who didn't use doping, on average 5 will test positive (this is called a false positive). You may assume that the test always gives the correct answer when an athlete used doping. After the tournament you test the winner,

and he turns out to test positive. On the basis of this evidence, you decide to disqualify him. What are the chances that you are making a mistake here?

5. **A Legal Case.** In O.J. Simpson's murder trial, the prosecution established that Simpson was physically abusive towards his wife Nicole. This, they argued, was serious evidence that he was capable of murdering her. The defense however, mentioned that according to the available statistics, of the 4 million women who get battered by their spouse, only 1,432 actually end up being murdered. The defense went on to conclude that the chances that an abuser will turn into a murderer are only about 1 in 2,500. If you had been on the jury, how would you have evaluated these statistical arguments?



II.1 ON GAMES OF CHANCE

While we are all familiar with the fact that many factors in life are unpredictable, it is not easy to make mathematically precise how to approach problems involving uncertainty. While various branches of mathematics are thousands of years old, the successful study of probability did not start until the 16th century. Its pioneer was a physician, philosopher, mathematician and gambler called Girolamo Cardano, whose tumultuous life and ideas we learn about in this section. His fundamental insight into probability was extended by Galileo, who formulated a more general and powerful principle. Using this principle, we can already calculate some of the probabilities that are at the basis of decision making in several games of chance.

II.1.1 PRE-HISTORY OF PROBABILITY

Gambling is not a modern phenomenon: many ancient civilizations have been shown to enjoy various games involving a component of luck; archeological findings include various forms of dice and board games at least five thousand years old. It is outside the scope of this course to give a comprehensive account here, but it is interesting and instructive to compare and contrast the ancient Greeks with the Romans with regard to their attitude towards, and understanding of, probability.

The ancient Greeks could pride themselves on great intellectual achievements: art, philosophy and mathematics thrived and the works produced two and a half thousand years ago are still being studied and admired today. They were also enthusiastic gamblers; one of their games of choice was played with a precursor of the die called *astragalus*, made of the heel bone of sheep and goats. One would expect that the scholars of those days with their mathematical and philosophical inclination would not hesitate to study the nature of chance and try to formulate the principles that govern it. After all, from a modern perspective probability theory is just another branch of mathematics; why study geometry but not probability?

However, no attempts were made to gain any understanding concerning probability. There are several possible reasons for this. First of all, the polytheistic religion in ancient Greece was not conducive to a study of chance. When every circumstance outside one's control is directly attributed to one of the

many deities governing the universe, it is all too easy to simply assume that the outcome of a game is also in the hands of the gods. Second, the philosophical and mathematical tradition in Greece was one in which proofs and absolute certainty were of central importance, and which frowned upon non-deductive reasoning. Finally, it has been argued that the lack of a proper numerical system was an impediment to the development of certain types of mathematics.

In Rome, gambling was no less popular, and explanation of events and natural phenomena by means of reference to one of the many gods and goddesses was no less common than in Greece. Unlike the Greeks, The Romans did not have any mathematical tradition to speak of. Still, in a way they came to better terms with probability than the Greeks. The main contributor was philosopher and statesman Marcus Tullius Cicero (106-43 B.C.) who coined the term *probabilis*, or likelihood. Cicero, a great admirer of Greek philosophy, seems to have had an intuitive grasp of the phenomenon that if a random experiment, such as rolling a die, is repeated a large number of times then one should expect non-random looking patterns to appear. Concretely, he understood that if you play for long enough, then you should not be surprised to see, say, a long sequence of sixes. This is a non-trivial insight, and most people find it counterintuitive. At least part of the claims that online poker is rigged stem from the fact that some people seem to fail understand that if you play millions of hands you will likely encounter bizarre series of bad beats.

II.1.2 GIROLAMO CARDANO

For a long time, the origin of probability theory as we know it today was thought to be a famous exchange of letters between Blaise Pascal and Pierre de Fermat. However, even though Pascal and Fermat made great strides and had fundamental insights, they were not the first to try to subject chance to rigorous study: almost a century before them, Cardano had at least one of the key ideas towards a systematic understanding of probability. There are several reasons why his work has been overlooked by later generations and why he has not been credited for his pioneering role. First of all, *The Book on Games of Chance* was published a century after it was written. Moreover, his writing on the subject was very unlike a modern mathematics text. Various erroneous and confusing attempts at solving the problem appear in the book, and only later a correct formulation appears.

Cardano's life was quite the opposite of the sheltered and uneventful existence most people imagine mathematicians to lead. Born in 1501 as the illegitimate and unwanted child of Fazio Cardano and Chiara Micheria, his life was filled with adversity from the very beginning. Briefly before he was born, his mother, a widow from a previous marriage, turned her back on Fazio and the city of Milan. Her attempt to end her pregnancy had failed, and after three days of labor an almost lifeless child saw the light of day. A bath of warm wine (!) revived the child, but serious health problems would haunt him throughout his life. He survived the plague when he was three months old, but it left him scarred. His mother repeatedly placed him in foster homes, until at age three she took him back; this was the first time he met his father, who was a reasonably well-respected mathematician and doctor of law and medicine, and who had just lost his other children to the plague.

Neglect, physical abuse and illness characterized the first 18 years of his life¹; although the family reunited when he was 8, he was mostly deprived of attention and warmth. When he was not suffering

¹Cardano's biographers are not in agreement on how serious this actually was; Cardano, in his autobiography, took various liberties with the truth in order to make for more appealing reading material.

from ailments he had to act as a page for his father, carrying his books. Beatings and flagellation were part of the daily regimen, and there seemed to be no pleasant childhood memories to speak of. That this left scars is not a surprise: later in life Cardano was extremely superstitious, developed a tendency towards automutilation, and became, by his own admission, seriously addicted to gambling.

When Girolamo reached adulthood, he grew increasingly frustrated with his life and frequently turned to playing cards and dice games, at which he soon became very proficient. He had the ambition to achieve things that would earn him respect and make him remembered by posterity, but he was in a bad position to do so. To make things worse, his father refused him to accept an inheritance which would have allowed him to attend university. He tried joining a religious order in order to escape his parental home, but that was also prevented.

Eventually he was allowed to study medicine in Pavia, and later in Padua. Because of his birth and gambling habits it took him three times before he was given the Doctor of Medicine title². He started a practice in a nearby town (Milan was being tormented by wars and the black plague at the time) and managed to subsist on a meagre income, supplemented by whatever he managed to earn at the gambling tables. After five years, he married, much to his own surprise - being extremely superstitious, he had believed for a long time that he was determined to remain single. The couple went back to Milan, but Cardano was prevented repeatedly from practicing medicine, partly because of his birth, partly because he had written a critical treatise on medical (mal)practice³.

Cardano had to wait until he was almost 40 years old before he finally got recognition for his mathematical and medical abilities. During his life he published over 100 books, and he left over 100 unpublished manuscripts, the most important being the *Liber de Ludo Aleae*, or *Book on Games of Chance*. In this book, he discusses various aspects of games and gambling, ranging from mathematical to psychological. There are chapters on the ethics of gambling (he states that games should always be played for money, otherwise the waste of time cannot be justified), on teaching and on cheating.

Unfortunately, Cardano's life ended as tragically as it had started. His wife died in 1546, leaving him with three children to care for. His daughter Chiara got pregnant from her brother Gianbattista, became infertile after the subsequent abortion. Gianbattista in turn married the wrong woman: her family had no money, and within a short period of time she managed to give birth to three children, none of which were her husband's. Driven into desperation, Gianbattista ordered a servant to poison her, but this was soon discovered, and he was arrested for murder. The stress and costs that resulted from the trial (after which Gianbattista was found guilty and executed) impoverished and weakened Cardano, and he was exiled from Milan. To make things worse, his younger son Aldo, who had a penchant for sadism, was coerced by one of Cardano's mathematical rivals⁴ into testifying against his father in exchange for a prestigious position as torturer for the inquisition. Cardano was sent to prison for several years. The last few years of his life he lived in Rome, but the events had taken too much of a toll on his mind, and he would not produce any significant work.

²In order to obtain this title one had to publicly defend theses, and a majority of scholars had to approve of the doctorate.

³His first published book, in 1536, was *On the Bad Practice of Medicine in Common Use*, based on the treatise *Differing of Doctor's Opinions* which was written during his studies. The book was a hit, but did not make him popular at the College of Physicians in Milan.

⁴Cardano had a long feud with Nicola Tartaglia about solving cubic equations.

II.1.3 CARDANO'S LAW

The main mathematical insight Cardano had was the following. Suppose that we consider a random experiment⁵; for concreteness' sake, think of throwing a die. Suppose furthermore that there are certain specific outcomes that we are interested in, because they give us a win, say. We shall call those outcomes *favorable*. Then the probability that the outcome is favorable can be found by counting the number of favorable outcomes and dividing it by the total number of all possible outcomes.

Let us reformulate this a bit. The collection of all possible outcomes of an experiment is called the *sample space*. Write \mathcal{A} for this collection. When our experiment consists of flipping a coin, then $\mathcal{A} = \{\text{heads}, \text{tails}\}$. When it consists of throwing a die, we have $\mathcal{A} = \{1, 2, 3, 4, 5, 6\}$. When we throw two dice, we get

$$\mathcal{A} = \{(1, 1), (1, 2), (1, 3), (1, 4), (1, 5), (1, 6), (2, 1), \dots, (2, 6), \dots, (6, 1), \dots, (6, 6)\},$$

a collection of 36 possible outcomes.

Next, write \mathcal{U} for the collection of favorable outcomes. In the case where throwing a five or a six is considered favorable, we would have

$$\mathcal{U} = \{5, 6\}.$$

Finally, let us write $P(X \in \mathcal{U})$ for the probability that the outcome of the experiment is favorable. The X represents the -unknown- outcome of the experiment⁶. We pronounce $P(X \in \mathcal{U})$ as “the probability that X is in \mathcal{U} ”, or “the probability that the outcome is in \mathcal{U} ”. In the above example, we would simply say “the probability that the outcome is 5 or 6”.

Cardano's Law: suppose all outcomes of a random experiment are equally likely. Then the probability of the outcome being favorable is:

$$P(X \in \mathcal{U}) = \frac{|\mathcal{U}|}{|\mathcal{A}|}.$$

Here, $|\mathcal{U}|$ stands for the number of elements of \mathcal{U} , i.e. the number of favorable outcomes, while $|\mathcal{A}|$ stands for the number of elements of \mathcal{A} .

Continuing the above example, we have $|\mathcal{A}| = 36$, and $|\mathcal{U}| = 2$, so the probability of throwing a five or a six is $2/36 = 1/18$.

A subset \mathcal{U} of favorable outcomes is commonly called an *event*. Thus $P(X \in \mathcal{U})$ may be reformulated as the probability that the event \mathcal{U} occurs. In our example we also write $P(X = 5 \text{ or } X = 6)$ to denote the probability of throwing a five or a six.

⁵Later, we will be more specific what we mean by a random experiment; for now the reader should think of processes with unpredictable outcomes, such as throwing a die, drawing a card from a deck, etcetera.

⁶In probability theory, X is referred to as a *random variable*, but we don't need the details of that.




Figure II.1: Roulette Layout

II.1.4 ROULETTE

The game of roulette does not involve a great deal of skill, but it gives us ample opportunity to see Cardano’s Law in action. The pockets of the roulette wheel are numbered 0 – 36. Thus there are 37 pockets in total⁷; after a spin of the wheel, the ball ends up in one of these pockets, resulting in one of the 37 possible outcomes in this game. The difficulty in roulette is predicting in which pocket the ball will end up; because the chance of getting this right is small, the casino allows for bets not just on single numbers, but also on various groups of numbers. For example, one may bet on even or on odd numbers, on red or black, on two adjacent numbers or on four adjacent numbers (see Figure II.1).

Suppose I bet \$50 on red. If I win I get paid 1 to 1 on my money, and so I stand to win \$50. What are my chances of winning? Well, there are 18 red numbers; these are the favorable outcomes. We already said that there are 37 possible outcomes in total. Thus

$$P(X = \text{red}) = \frac{18}{37} = 0.486486\dots$$

This probability is slightly below one half. Thus, in the long run I should expect to win about 486 in 1000 times. It’s precisely this small difference which makes the casinos rich and most visitors go home empty handed.

Exercise 16. Calculate the following probabilities:

- (i) The probability that one of the numbers 1 – 12 falls.
- (ii) The probability that 0 falls.
- (iii) The probability that either an even number or a black number falls.
- (iv) The probability that a black, even number falls.

⁷In American roulette, there is an additional pocket labelled 00.

Exercise 17. Occasionally one witnesses someone placing equal bets on red and on black at the same time. (One may make any combination of bets one likes, so this is not against any of the rules.) Is this a good strategy? Why (not)?

Even though the roulette table is usually not the place where great intellectual battles can be witnessed, exciting events do occur occasionally. One of the most famous roulette stories is that of Joseph Jagger, a man whose engineering experience let him to speculate that roulette wheels might well show a tiny but exploitable bias as a result of mechanical imperfections. In 1873, He took this idea to Monte Carlo, and instructed his assistants to collect data about the outcomes of various roulette wheels in one of the casinos. After analysing the data, he found that while most wheels turned out to be reasonably fair, one of them showed a bias towards certain numbers. Jagger started to bet heavily on those, and within a short period of time he managed to earn an amount of money which today would be several millions of dollars. Of course, the casino quickly grew suspicious, and took countermeasures. Eventually Jagger realised that he no longer had the edge, and took his remaining winnings which allowed him a comfortable early retirement.

Less of a financial success but scientifically more interesting were the endeavours of the *Eudaemon collective*, a group of young students, scientists and computer hackers led by Norman Packard and James Doyne Farmer. Their insight was that while there is no mathematically sound roulette strategy, it is possible to improve one's odds by exploiting on the fact that the roulette wheel has to obey the laws of physics. In principle, if one knows the speed of the rotor (the revolving part of the table) and the speed and position of the ball with respect to the rotor, then, given enough information about the rate of deceleration of the ball and rotor, one can predict where the ball will end up. This is harder than it seems, because many variables come into play: tilt of the wheel, atmospheric conditions, the material of the ball, etcetera. Still, one can estimate these parameters by repeatedly measuring the velocities of the rotor and the ball. Given sufficiently accurate estimates, a computer can solve the mathematical equations needed to predict the outcome of subsequent spins.

The Eudaemons decided to work in pairs; one player would first collect data about the wheel and enter it into a portable computer; then once the parameters of the wheel were estimated, the first player would broadcast predictions to the second player, who would place bets accordingly. Since the computer would predict quite reliably in which sector of the wheel the ball was going to fall, the expected return on investment was close to 44%, a lot better than the -2.7% the common man has to accept.

The process of working out the mathematics needed for the predictions, writing the software, building the portable computers, endless testing, debugging and fine tuning cost years. While some successes were booked, a systematic exploitation of their invention ran into practical difficulties. The book *The Eudaemonic Pie* by Thomas Bass tells the full story of their adventures.

Independently, Edward Thorp and Claude Shannon had experimented with a similar scheme. Thorp had been a successful card counter in blackjack, and was the author of the classic book *Beat the Dealer*. (Later, he would run an extremely successful hedge fund called Princeton-Newport Investments, and would write the follow-up book *Beat the Market*.) However, he was no longer welcome in many casinos, and he felt that the risk of getting caught and roughed up by one of the pit bosses was no longer worth the reward. Shannon, the inventor of the branch of science called *information theory* and professor at MIT, had a passion for inventing and constructing electronic gadgets. Together with Thorp, he built a computer system similar to that of the Eudaemons. However, they also ran into the practical issues which ultimately threw off the Eudaemons; it turned out to be too difficult to get the system to work

smoothly in the noisy casino environment.

II.1.5 THE MISSING INGREDIENT

Cardano's Law is correct and useful, but it only applies when all outcomes of the experiment are equally likely to occur. The probability of a coinflip landing heads is $1/2$ because of the two possible outcomes {heads, tails} precisely one is favorable. However, this presupposes that the two outcomes are equally likely, i.e. that the coin is fair! As is clear, the probability of a biased coin landing heads is different from $1/2$. (This almost is a tautology.)

It took several decades before Galileo managed to answer the more subtle question of how to determine probabilities when the outcomes are not equally likely. Galileo was not a gambler, but he was led to consider probability theory because the Grand Duke of Tuscany had observed during his playing that if you throw three dice, then the total sum 10 is more likely to occur than 9. He commissioned a study by Galileo to investigate this phenomenon. The key observation made by Galileo is the following:

The probability of an event occurring is dependent on the number of ways in which the event can occur.

To determine probabilities if not all outcomes are equally likely, one has to count the number of ways in which certain events can occur. Let us do this for the Grand Duke's problem. First, we list all the possible ways of throwing a total of 10:

(1, 3, 6), (1, 4, 5), (1, 5, 4), (1, 6, 3),
 (2, 2, 6), (2, 3, 5), (2, 4, 4), (2, 5, 3), (2, 6, 2),
 (3, 1, 6), (3, 2, 5), (3, 3, 4), (3, 4, 3), (3, 5, 2), (3, 6, 1),
 (4, 1, 5), (4, 2, 4), (4, 3, 3), (4, 4, 2), (4, 5, 1),
 (5, 1, 4), (5, 2, 3), (5, 3, 2), (5, 4, 1),
 (6, 1, 3), (6, 2, 2), (6, 3, 1)

There are 27 of them. Similarly, we may check that there are 25 ways of throwing a total of 9. (Check for yourself!) Since there are $6 \times 6 \times 6 = 216$ possible outcomes in total, we have

$$P(\text{sum is } 10) = \frac{27}{216} = 0.125, \quad P(\text{sum is } 9) = \frac{25}{216} = 0.1157\dots$$

In this example, we were able to list and count the number of ways in which a favorable outcome can occur. In general, this can be tricky, and later on we shall develop some more systematic approaches to such counting problems.

Let us look at another example: suppose I offer you an even money bet that a throw with two dice gives a total between 5 and 8. Thus, if the total ends up between 5 and 8 you pay me one dollar, if it

is less than 5 or more than 8 I pay you one dollar. Should you accept? To find out, determine first the set of favorable outcomes:

$$\mathcal{U} = \{(1, 1), (1, 2), (1, 3), (2, 1), (2, 2), (3, 1), (3, 6), (4, 5), (4, 6), (5, 4), (5, 5), (5, 6), (6, 3), (6, 4), (6, 5), (6, 6)\}.$$

Counting these gives $|\mathcal{U}| = 16$. There are still 36 possible outcomes, so the probability comes to

$$P(X \in \mathcal{U}) = \frac{16}{36}.$$

This is slightly less than $1/2$, so you should expect to lose this bet more than half of the time, and hence you shouldn't accept it.

It is easy to go wrong in probability calculations. Below is another calculation of the probability that the sum of the dice is less than 5 or greater than 8 which at first sight looks plausible. See if you can spot the flaw in the reasoning.

—Wrong calculation starts here—

We are interested in the total sum of the dice. The collection of all possible sums of two dice is $\{2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$. The favorable outcomes are $\{2, 3, 4, 9, 10, 11, 12\}$. Thus of all 11 possible outcomes 7 are favorable. Hence by Cardano's Law the probability is $7/11$.

—Wrong calculation ends here—

Where did we make a mistake? Every line is true, except for the last one, which is a non-sequitur. We overlooked the fact that some sums are more likely to occur than others and hence that Cardano's Law does not apply. For example, the sum 3 can occur in two different ways, namely $(1, 2)$ and $(2, 1)$. The sum 4 can arise as $(1, 3)$, $(2, 2)$ or $(3, 1)$, and so on. Thus the condition for Cardano's Law is not satisfied: not all outcomes are equally likely.

II.1.6 SOME EXAMPLES

Being able to calculate the probability of events is indispensable if you wish to be successful at any game of chance. One of the most common situations in poker is the matchup between a made hand and a drawing hand. The terminology "made hand" and "drawing hand" is not very precise; however, for the present purposes we simply consider situations in which one player needs help from the deck in order to make a hand, while the other player currently is ahead.

Here's an example from Texas Holdem: suppose you hold $7♥8♥$, and the board on the turn shows $2♥K♣5♥T♦$. What is the probability that you make your flush on the river? First, determine the set of all possible outcomes (river cards): there are 52 cards in a deck, but 6 of these are already gone, so there are $52 - 6 = 46$ unseen cards left in the deck. Next, determine which of these are favorable, i.e. make your flush: these are $3♥, 4♥, 6♥, 9♥, T♥, J♥, Q♥, K♥, A♥$. Thus there are 9 outcomes which are favorable.

(In poker, the favorable outcomes are usually called *outs*.) Hence the chance that you make your flush on the river is $9/46$ (a bit more than 20%).

There are some points worth noting in the above example:

- There may be other people in the hand, each holding two cards. Does this change our calculation? The answer is no; as long as we haven't seen these cards there are 46 unseen cards, regardless whether some of these have been dealt to others.
- However, if one of our opponents would flash his cards to us, revealing $4♥4♦$, then we would have new information. Now there are only 44 unseen cards left, and of those 8 are hearts. This would change the probability of hitting our flush to $8/44$, a lower probability.
- None of the above considerations says anything about what we should do in the hand! This depends on a lot of other factors, which we shall discuss in detail later on.

The next example comes from Pot Limit Omaha. Omaha is similar to holdem, but in Omaha each player receives four hole cards, and has to use *exactly two of his hole cards* together with three from the board to form a poker hand. (It takes a bit of practice to read hands correctly in Omaha.) Omaha is most commonly played with a *pot limit* betting structure, meaning that one's bet size cannot exceed the size of the pot.

Suppose player 1 holds $A♥J♠J♦T♥$, and player 2 holds $K♠K♦Q♥9♦$ (both quite reasonable PLO starting hands). On the turn, the board is $K♥8♠4♥7♠$, so that player 2 currently is in the lead with a set of kings. What is the probability that player 1 makes a winning hand on the river?

First, determine the total number of possible outcomes: we know 12 cards, so that leaves $52 - 12 = 40$ unseen cards. Next, figure out which cards make a winning hand for player 1. This is a bit tricky, because we have to take into account the possibility of making a flush but losing to a full house or quads. For example, the $8♥$ would give player 1 a flush, but player 2 a full house; for this reason such a card is often called a "tainted" out. The "untainted", or "clean", outs to a flush are $2♥, 3♥, 5♥, 6♥, 9♥$ and $J♥$. (Note that board-pairing hearts are not counted, nor is the $K♥$.) Player 1 can also win by making a straight: this happens if the $9♠$ or $9♣$ falls (we have already taken the $9♥$ into account). Thus, there are 8 winning river cards, and hence the probability that player 1 draws out is $8/40 = 1/5$. This kind of matchup of a set versus a flush- and/or straight draw is very common in Omaha.

Exercise 18. The following Texas Holdem hand occurred during the Durr Million Dollar Challenge (Episode 1). Tom Dwan was holding $J♣5♣$, while his opponent Sammy George held $K♦K♠$. The board on the turn read $3♣2♣9♠5♥$. What are the chances that Dwan draws out on the river?

Note that there is a difference between a "third person"-calculation where we, as observers, know the cards of both players, and the calculation you do as a player, where you only know your own cards and the board. In the latter case you have less information. In particular, it can be difficult to decide which outs are tainted and which are not. Later we'll come back to this problem and see how we can handle these kinds of issues using hand ranges.



II.2 HAND COMBINATIONS AND RANGES

In poker, you are constantly wondering about what your opponent may have. In this section we explain some methods for reasoning about your opponent's possible holdings. First we look at the likelihood of various possible types of starting hands. Then we consider hand ranges, and learn how to break these down into various possible types of hands. Finally, we see how new information can help us narrow down our opponent's hand range.

II.2.1 STARTING HANDS

In Texas Hold'em, each player receives two *hole cards* at the beginning of each hand; these cards are visible only to that player and not to the rest of the table. The two cards which you receive at the beginning are referred to as your *starting hand*; mathematically, they are just the result of drawing two random cards from a shuffled deck. The order of the two cards is irrelevant for the game. Thus it doesn't matter whether you first receive, say, an ace of clubs and next a ten of spades or the other way around. Put differently, $A\clubsuit T\spadesuit$ and $T\spadesuit A\clubsuit$ are considered the *same starting hand*. By contrast, $A\clubsuit T\spadesuit$ and $A\spadesuit T\clubsuit$ are different starting hands; however, we will sometimes say that they are the same *type*: both are ace-ten hands.

Problem 1. How many different starting hands are there? There are 52 cards in the deck. If you draw one card, you thus have 52 different possible outcomes. After you've drawn one card, there are 51 cards left, and for your second card you then have 51 possible outcomes. That gives $52 \cdot 51$ different ways of drawing two cards. However, this way of counting means that we take the order of the two cards into account, which we shouldn't (at least not for the purposes of Hold'em). So we're counting every hand double: hence the answer is:

In hold'em, there are $\frac{52 \cdot 51}{2} = 1326$ different starting hands.

Of course, some starting hands are better than others⁸: a pair of aces is the best starting hand, while 72o (seven-deuce offsuit) is the worst.

Exercise 19. Why is 72o a worse hand than, say, 53o? Why is 72o worse than 72s?

Problem 2. What is the probability of getting dealt AA in hold'em? To answer this question (and many others) we need to consider the concept of *hand combinations*. By a *combination*, we mean a particular way a certain type of hand can occur. In the case of AA, there are 6 combinations:

$$A\heartsuit A\diamondsuit, A\heartsuit A\spadesuit, A\heartsuit A\clubsuit, A\diamondsuit A\spadesuit, A\diamondsuit A\clubsuit, A\spadesuit A\clubsuit.$$

(Make sure you understand why this is so!) All of these are equally likely. That means that the probability of getting dealt AA is

$$P(\text{AA}) = \frac{6}{1326} \approx 0.45\%.$$

⁸Technically, this is because they have a higher probability of making a winning hand.

Problem 3. What is the probability of getting dealt 72o in hold'em? Again we need to count combinations. We list all 72o hands:

$$\begin{aligned} &7♥2♦, 7♥2♠, 7♥2♣, 7♦2♥, 7♦2♠, 7♦2♣, \\ &7♠2♥, 7♠2♦, 7♠2♣, 7♣2♥, 7♣2♦, 7♣2♠. \end{aligned}$$

Note that we're not counting the four remaining combinations 7♥2♥, 7♦2♦, 7♠2♠, 7♣2♣, because these are suited. We have 12 combinations, hence:

$$P(72o) = \frac{12}{1326} \approx 0.90\%.$$

A *suited connector* is a hand in which the cards are of consecutive rank and of the same suit. For example, 4♣5♣ is a suited connector, as is J♥T♥.

Problem 4. What is the probability of getting dealt a suited connector in hold'em? Let's first count all possible suited connectors involving one suit only. Here are all the possible combinations in the suit of hearts:

$$\begin{aligned} &2♥3♥, 3♥4♥, 4♥5♥, 5♥6♥, 6♥7♥, 7♥8♥, \\ &8♥9♥, 9♥T♥, T♥J♥, J♥Q♥, Q♥K♥, K♥A♥. \end{aligned}$$

This gives 12 combinations. Of course, each of the three other suits also has 12 of such combinations, giving a total of $4 \times 12 = 48$ suited connector hands. Thus

$$P(\text{suited connector}) = \frac{48}{1326} \approx 3.6\%.$$

A *Broadway hand* is a hand consisting of cards ten or higher. One usually doesn't count pairs as Broadway hands.

Problem 5. What is the probability of getting dealt a Broadway hand in hold'em? Again, we need to do some counting. This time, the number of combinations is getting too large to list everything, and we need to be systematic. We split the problem into a bunch of smaller ones. There are the following different types of Broadway hands:

$$AK, AQ, AJ, AT, KQ, KJ, KT, QJ, QT, JT.$$

Each of these can occur in 16 possible ways (four suits for the first card, and four for the second). Adding these up gives

$$P(\text{Broadway}) = \frac{160}{1326} = 12.1\%.$$

Exercise 20. How many combinations of Broadway hands are suited? Suppose we know our opponent has a Broadway hand. What is then the probability that his hand is suited?

Exercise 21. Find the probability of getting dealt two cards lower than 5 (including pairs lower than 5).

Exercise 22. A *suited A* is a hand of the form Axs, where x can be any non-ace of the same suit as the ace. Find the probability of getting dealt a suited A before the flop in hold'em.

II.2.2 HAND RANGES

In practice, it is usually impossible to deduce which exact hand your opponent has. Instead, we try to put our opponent on a *range* of hands: this means that we consider a number of different types of hands we believe our opponent might hold in a particular situation. The hand range we assign to our opponent is a function of our past experience with him/her, his/her actions during the hand and possibly other information as well (for example statistical information about his/her play⁹).

A typical hand range for a conservative opponent (let's call him Buster) who raises before the flop in hold'em might be something like this: any pair 66 or bigger, any Broadway hand, any suited connector 78s or higher.¹⁰

Problem 6. How frequently can Buster expect to get a hand he can raise with? This is a counting problem again. We can combine results from earlier problems, but we have to be careful not to count certain hands twice.

Type	Hands	Combinations
Pairs	66, 77, 88, 99, TT, JJ, QQ, KK, AA	$9 \times 6 = 54$
Broadway	AK, AQ, AJ, AT, KQ, KJ, KT, QJ, QT, JT	$10 \times 16 = 160$
Suited Connectors	78s, 89s, 9Ts	$3 \times 4 = 12$
Total		226

Note that we're not counting the suited connectors JTs and higher, as they've already been included in the Broadway hands. Now Buster can expect a raising hand about 226 in 1326 times, which is roughly 1 in 7 hands.

Problem 7. You're watching Buster play (you're not playing yourself, so you haven't seen any of the cards), and he raises. What is the probability that he has a pair of aces?

First, the important point here is that it is NOT what we calculated in Problem 1. The reason is that we have more information about Buster's hand: it is no longer a random starting hand, but it is one of the hands in his preflop raising range. There are 226 hand combinations in this range, and we need to use that as the number of possible hands he can have. Of these, 6 are AA hands. This gives a probability of $\frac{6}{226} \approx 2.7\%$.

Note that the narrower someone's range is, the higher the probability that (s)he has a big hand. For example, if Buster was even more conservative and didn't raise any unsuited Broadway other than AK, his range would consist of only 118 combinations (Check!), making the probability that he was raising with AA about 5.1%.

Exercise 23. Mr. F plays any pair and any two suited cards. How many hand combinations does this amount to? If you see Mr. F play a hand (without knowing any of the other cards), what is the probability that she has a pair?

Exercise 24. In his book "Play Poker Like The Pros", Hellmuth advocates playing the following "top ten" hands:

⁹Many players make use of statistical software to gain insight into their opponents' playing tendencies.

¹⁰In reality, someone who plays these hands will probably also play some suited one-gappers such as T8s, but we ignore this for now.

AA, KK, QQ, AK, JJ, TT, 99, 88, AQ, 77.

(This is Limit hold'em advice for beginners.) What is the probability of getting dealt one of these hands?

Exercise 25. Sometimes hold'em players have a side bet which pays a certain amount to anyone winning a pot with 72 (suited or unsuited). Suppose you're in such a game and one of your opponents makes a big raise which can only mean AA, KK or 72. What is the chance he has 72?

II.2.3 NARROWING DOWN A RANGE

As soon as we obtain new information, we can try to use it to narrow down our opponent's range of possible holdings. This already starts when we look at our own cards: we know for sure our opponent can't have these!

Problem 8. Buster raises again. You look at your cards and see a pair of Jacks. How does this affect Buster's range?

Intuitively, the fact that you have two jacks makes it less likely Buster has a hand containing one or two jacks. The table now becomes:

Type	Hands	Combinations
Pairs other than JJ	66, 77, 88, 99, TT, QQ, KK, AA	$8 \times 6 = 48$
Pairs JJ	JJ	1
Broadway without J	AK, AQ, AT, KQ, KT, QT	$6 \times 16 = 96$
Broadway with J	AJ, KJ, QJ, JT	$4 \times 8 = 32$
Suited Connectors	78s, 89s, 9Ts	$3 \times 4 = 12$
Total		189

(Note that we get, for example, 8 combinations of AJ, because there are four aces left in the deck and two jacks, giving $4 \times 2 = 8$ possible combinations.)

Problem 9. (Previous problem continued.) How likely is it that Buster has a higher pair than you?

Of his 189 possible holdings, QQ, KK and AA are higher pairs. There are $3 \times 6 = 18$ such hands, giving a probability of $\frac{18}{189} \approx 9.5\%$.

Community cards also allow us to narrow down our opponent's range. Conversely, understanding our opponent's range allows us to determine the likelihood of him/her connecting with the flop in various ways.

Problem 10. Buster raises, you call with JJ, and the flop comes $2\clubsuit 6\heartsuit K\spadesuit$. What is Buster's range now? And how likely is it that you currently have the best hand?

First, narrow down Buster's range by removing hands containing one of the three flop cards.

Type	Hands	Combinations
Pairs other than JJ, KK, 66	77, 88, 99, TT, QQ, AA	$6 \times 6 = 36$
Pairs JJ	JJ	1
Pairs KK	KK	3
Pairs 66	66	3
Broadway without J,K	AQ, AT, QT	$4 \times 16 = 48$
Broadway with J, without K	AJ, QJ, JT	$3 \times 8 = 24$
Broadway without J, with K	AK, KQ, KT	$3 \times 12 = 36$
Broadway with J,K	JK	6
Suited Connectors	78s, 89s, 9Ts	$3 \times 4 = 12$
Total		169

Now we need to count how many of these hands beat our JJ.

Type	Hands	Combinations
Higher pairs	QQ, AA	$2 \times 6 = 12$
Sets	66, KK	$2 \times 3 = 6$
Top pair	AK, KQ, KT, JK	$3 \times 12 + 6 = 18$
Total		36

This means that the probability that we are currently behind is $\frac{36}{169} \approx 21\%$.

Finally, we can narrow down our opponent's range by considering his/her actions. This is the trickier part, and you need to know your opponent's style and tendencies before you can make good use of this.

Problem 11. Buster raises, you call with JJ, and the flop comes $2\clubsuit 6\heartsuit K\spadesuit$. Buster shoves all-in. What is his range now?

Unfortunately we can't say for sure without knowing Buster's style of play in spots like these. But we can think about the different types of hands he can have and with which of them an aggressive play like this would make sense¹¹. Generally, a large bet can mean several things, such as:

- "I have a big hand and I'm hoping you have a hand you can call with."
- "I have nothing and I'm hoping you will fold."
- "I have a draw and am not terribly worried about getting called, but if you fold I'm fine with that too."

In this particular hand, there are no possible draws, so Buster either has a strong hand or has missed and is bluffing. Which of the two is more likely? First, let's see which hands we can eliminate from Buster's range given this play. Even if Buster would get this excited with top pair or better, he would probably not push all-in with a made hand weaker than top pair. This eliminates hands such as 99. His range becomes:

¹¹Of course, we're ignoring many vital bits of information, such as stack sizes, and so on. The analysis in this example is only meant to illustrate one aspect of the kind of reasoning which can be used in spots like these.

Type	Hands	Combinations
Top pair or better:	AA, KK, 66, AK, KQ, KJ, KT	$6 + 6 + 12 + 12 + 6 + 12 = 54$
Broadway bluffs	AQ, AJ, AT, QJ, QT, JT	$16 + 8 + 16 + 8 + 16 + 8 = 72$
Suited Connectors	78s, 89s, 9Ts	$3 \times 4 = 12$
Total		138

Out of these, 54 are made hands, and the rest are bluffs. That means that the chance that Buster has a real hand is $\frac{54}{138} \approx 39.1\%$.

If you would know (from past experience, say) that Buster would play a hand like KT more passively here then you could use this information to narrow down his range even further. Indeed, it is unlikely that you are going to call an all-in with a hand weaker than top pair here; so when Buster pushes all-in with a weak top pair, he cannot get expected to get called by weaker hand, and hence there would be little point in playing that way. Hence, provided we deem Buster capable of that analysis (and if you haven't given him reason to believe you'd stack off with mid pair here) then we can eliminate most top pair hands from his range.

When a player has only very strong hands and bluffs in his range, we say that his range is *polarized*. This is not necessarily a bad thing; as long as the proportion of bluffs versus strong hands is appropriate, we cannot immediately exploit such a situation. And in some cases there are good reasons for reserving certain plays for your weakest and strongest hands only, and to take a different line with medium-strong hands. (Tournament play would be an example: in the early stages of a tournament when you have an edge over your opponents you don't always want to take big risks with hands such as 99, even though you might be a slight mathematical favorite. In this situation, surviving is more important than pushing every small edge you have, as you would do in a cash game.) However, polarized ranges are often easier to play against and reason about than non-polarized ones, if only because for the latter there are more possibilities to take into account.

Exercise 26. Suppose your opponent only plays the “top ten” hands AA-77, AK, AQ. Suppose you're holding AJ, and the flop comes J52 rainbow.

- How many hand combinations are in your opponent's range?
- How many of those are ahead of your hand?
- What is the chance that you have the best hand?

Exercise 27. Suppose your opponent plays any suited ace and any pair. You're holding $7\clubsuit 7\spadesuit$, and the flop is $3\clubsuit 8\heartsuit K\heartsuit$.

- How many hand combinations are in your opponent's range?
- How many of those combinations give your opponent a hand which is currently ahead of yours?
- How many give him a flush draw? (With or without a pair.)
- Suppose your opponent pushes all-in on this flop, and that you think he would do this with a flush draw, with top pair or better, or if he had nothing. (In other words, he wouldn't push all-in with a hand like second pair.) What is the probability that you currently have the best hand?



II.3 WHAT TO EXPECT?

Knowing the probability of an event occurring is one thing, but using that information to your advantage is another. In this section we learn how to determine the most profitable course of action by means of the concept of Expected Value. This allows us to decide whether we should accept or reject a bet, whether to make a bluff or not, and whether it is a good idea to play in a certain lottery. What is common to all those situations is that each of them has several possible outcomes, and that to each of these outcomes we associate a certain payoff. For example, winning the lottery might earn us three million, while if we don't win we wasted the price of the ticket. The concept of Expected Value allows us to combine the information about the payoffs of the various possible outcomes with the information about how likely each of these outcomes is. After that, our decision is simple: choose the option which gives us the highest Expected Value.

Below, we will first have a brief look at the origins of this concept. After that, we give a precise definition, and we follow with numerous examples of how to apply this in real life.

II.3.1 PASCAL'S WAGER

According to some, the first documented use of Expected Value (EV from now on) to make a decision occurred in a theological work by Blaise Pascal (1623 – 1662), called *Pensées*. This work, which was never fully completed, was quite the opposite of the Gambler's handbook left to us by Cardano; instead, the *Pensées* are a justification of Christianity¹². While he was younger, Pascal had been one of the central figures in French mathematics and physics, but after a mystical experience he turned to philosophy and theology. Throughout adulthood, he was vexed by health problems¹³, and after his religious conversion he led an ascetic life. One would not expect a discussion of profitable betting in a book dealing with profound philosophical issues such as life and death, faith and reason and the nature of the infinite, especially not when that book was written by a devout Christian such as Pascal!

While Pascal does not address gambling, he does introduce us to a key aspect of decision making. Pascal was convinced that we cannot deduce the existence or non-existence of God by logical, or rational means. We are fundamentally uncertain about the issue, and thus the question of whether we should believe in the afterlife is one worth investigating. In our terms, Pascal studied a game in which we have only two moves: believe in, or reject the afterlife. According to Pascal, our decision should depend on the following considerations: first, how likely or unlikely we deem the existence of God (and hence of the afterlife) to be; how much there is to be gained by correctly believing; how much there is to be gained by correctly rejecting; how much is lost when we incorrectly reject God; and finally, how much is lost when we incorrectly believe.

Pascal goes on to provide various estimates for these quantities. If you correctly believe in the afterlife, then you gain eternal life (an infinite gain according to Pascal); if you incorrectly reject the afterlife, then you are doomed, which is a very bad thing. The price of incorrectly believing is relatively

¹²Pascal was an adherent of the religious movement called Jansenism, which was considered heretic by the Catholic church.

¹³An autopsy showed severe abdominal problems and brain lesions, among other things.

small, and so is the gain of correctly rejecting. In modern terms, Pascal assigns utilities to each of the possible outcomes:

correctly believing:	$+\infty$
incorrectly believing:	-10
correctly not believing:	$+10$
incorrectly not believing:	$-\infty$

With these estimates in place, the key point of his argument is the following: even if we only think that there is only a small probability that God exists; because the potential gain from correctly believing is infinitely large, it outweighs all other considerations.

II.3.2 EXPECTED VALUE

Pascal's reasoning introduced two ideas, and in order to better understand what's going on we should separate the two. The first has to do with infinity. In Pascal's opinion, making it to the eternal afterlife was worth infinitely more than any possible earthly event. This idea is interesting, but does not make the road to understanding the second idea, namely that of expectation, any smoother. We therefore stick to finite utilities for now, and consider an easier example.

Example II.3.1. Suppose that you consider hosting a garden party coming Saturday. You are worried about the weather; if it rains that day, most of the fun will be gone. How should you decide?

If we follow Pascal's reasoning, then we should provide estimates for the probability that it will rain on Saturday, and for the values (payoffs) associated to the various scenarios. Here they are:

- Chance of rain on Saturday: 40%
- Sunny party: 40
- Rain on your party: 5
- Sunny, no party: 15
- Rainy, no party: 25

Of course these payoffs are personal and may be different for everyone.

Now for each of the four possible scenarios we multiply the probability with their utility:

- Sunny party: value = $0.6 \cdot 40 = 24$
- Rain on your party: value = $0.4 \cdot 5 = 2$
- Sunny, no party: value = $0.6 \cdot 15 = 9$

- Rainy, no party: value = $0.4 \cdot 25 = 10$

We now compare the total value for each possible decision: if we host the party, then we gain $24 + 2 = 26$. If we decide to cancel, we gain $9 + 10 = 19$. The expected value of hosting the party is greater than that of canceling, and hence you should go on with the organizations. One way to think of this is as follows. If you would make this decision over and over again, then your average payoff would tend towards 26 when you would always host the party, while your average payoff would tend towards 19 when you would always cancel.

We now formulate the principle behind this. Given a random experiment with n possible outcomes s_1, \dots, s_n to which we assign probabilities p_1, \dots, p_n , and values v_1, \dots, v_n , the *expected value* is

$$EV = p_1 v_1 + p_2 v_2 + \dots + p_n v_n.$$

Thus:

Expected Value is the weighted average of the values of the different possible outcomes, weighted by their probabilities.

It is important to keep in mind that the different possible outcomes should be *jointly exhaustive* and *mutually exclusive*. In plain English, this means that no two of them can happen at the same time, and that at least one of them should happen.

The term “value” in “Expected Value” is supposed to be interpreted in a neutral way, and can have many meanings depending on the particular situation we’re in. In the case of Pascal’s Wager and the above example, we computed the expected utilities of various outcomes. This is common in game-theoretic and economic situations. We also speak of expected payoffs in these cases. In other cases, the value we’re interested in could be simply a number.

For example, suppose we roll one die and that the value we’re interested in is simply the number of eyes we roll. What is the expected value? Well, there are six possible outcomes 1,2,3,4,5,6. The probability of each is $1/6$, and the value we associate to each of the outcomes is simply the number of the outcome. Thus:

$$EV = \frac{1}{6} \cdot 1 + \frac{1}{6} \cdot 2 + \frac{1}{6} \cdot 3 + \frac{1}{6} \cdot 4 + \frac{1}{6} \cdot 5 + \frac{1}{6} \cdot 6 = 3.5$$

What this number tells us is that if we roll the die many times and record the outcomes, then if we compute the average outcome over all those trials it will be close to 3.5.

If we roll two dice and ask about the sum, then we can do a similar computation. In this case, we need to consider all possible totals and their probabilities. The possible totals are 2,3,4,5,6,7,8,9,10,11,12, and their probabilities are computed as before. Then we get:

$$\begin{aligned} EV &= 1/36 \cdot 2 + 2/36 \cdot 3 + 3/36 \cdot 4 + 4/36 \cdot 5 + 5/36 \cdot 6 + \\ &\quad 6/36 \cdot 7 + 5/36 \cdot 8 + 4/36 \cdot 9 + 3/36 \cdot 10 + 2/36 \cdot 11 + 1/36 \cdot 12 \\ &= 7. \end{aligned}$$

Note that this is just twice the expected total for throwing one die. Often, it is possible to compute the EV of a complex event in terms of simpler events. In this case, that is possible because the outcomes of the two individual dice are *independent*. We shall come back to this later.

Exercise 28. What do you think is the expected value of the total sum when we throw 3 dice? What about 10 dice?

Exercise 29. Suppose I offer you a bet on the outcome of the throw of two dice. If both dice land on the same number, I pay you \$10. If they don't, you pay me \$2. Should you accept my bet?

Exercise 30. In a certain lottery there is only a single prize, and it is \$10 million dollar. Every ticket is a 7-digit number. How much would you be prepared to pay for a ticket?

Exercise 31. At the horse track, you believe to have found a horse that is about 30% certain to win the race. You contemplate betting \$100 on the horse. If you win, you will get paid \$400. However, the bookie takes a 10% cut of your winnings. Should you make the bet?

II.3.3 ROULETTE EXPECTATION

Earlier we claimed that roulette is not a game of skill, but of pure luck. With the concept of EV we can substantiate that claim. Let us calculate the expectation of some typical and not-so-typical bets.

1. Suppose I bet \$10 on red. What is the expected value of this bet? Well, as we determined earlier, our chances of winning are $18/37$, and our chances of losing are $19/37$. In case of a win, we gain \$10, and in case of a loss we lose \$10. Thus

$$EV = \frac{18}{37} \cdot \$10 + \frac{19}{37} \cdot (-\$10) = \frac{1}{37}(-\$10) \approx -\$0.27.$$

This means that, on average, we expect to lose 27 cent each time we place this bet. Of course, we cannot actually lose this exact amount, just as it is impossible to have the average number of 1.7 children. What this number means is that if we were to play the game many times in succession, then our average loss per game would be roughly 27 cents. If we were to bet a different amount, say \$100, then we would expect to lose \$2.70 on each play. In general, betting \$Y gives an expected loss of $\frac{1}{37} \cdot \$Y$.

2. Could we do better by placing our \$10 elsewhere on the table, say on a single number? If we place \$10 on number 8, for example, then we have a $1/37$ chance of winning \$350, and a $36/37$ chance of losing our \$10. Therefore

$$EV = \frac{1}{37} \cdot \$350 + \frac{36}{37} \cdot (-\$10) = \frac{1}{37}(-\$10) \approx -\$0.27,$$

just as before.

3. We could also bet on a group of numbers, say 20-21-22-23. In that case we gain \$80 when we win, and \$10 when we lose. Thus

$$EV = \frac{4}{37} \cdot \$80 + \frac{33}{37} \cdot (-\$10) = \frac{1}{37}(-\$10) \approx -\$0.27,$$

4. What about a combination bet, say \$3 on red, \$3 on even, \$2 on 0 and \$2 on 20-21-23-24? This can be computed in different ways. The first is to calculate the EV of the combination bet by figuring out what each possible outcome will earn us. The difficulty here is that some numbers win twice, for example if 12 hits then we earn \$3 from our bet on red and \$3 from our bet on even. We get the following:

Outcome	Probability	Value
0	1/37	$35 \cdot \$2 - \$8 = \$62$
1, 3, 5, 7, 9, 25, 27	7/37	$1 \cdot \$3 - \$7 = -\$4$
2, 4, 6, 8, 10, 22, 26	7/37	$1 \cdot \$3 - \$7 = -\$4$
12, 14, 16, 18, 28, 30, 32, 34, 36	9/37	$2 \cdot \$3 - \$4 = \$2$
20, 24	2/37	$1 \cdot \$3 + 8 \cdot \$2 - \$5 = \14
21, 23	2/37	$1 \cdot \$3 + 8 \cdot \$2 - \$5 = \14
11, 13, 15, 17, 19, 29, 31, 33, 35	9/37	$-\$10$

Thus our EV is:

$$EV = \frac{1}{37} \cdot \$62 + \frac{7}{37} \cdot (-\$4) + \frac{7}{37} \cdot (-\$4) + \frac{9}{37} \cdot \$2 + \frac{2}{37} \cdot \$14 + \frac{2}{37} \cdot \$14 + \frac{9}{37} \cdot (-\$10) \approx -\$0.27.$$

Alternatively, we can compute the EV of each of the bets separately and then combine them. Red wins 18/37 of the time, in which case we get \$3. Same for even numbers. A \$2 bet on a single number has a 36/37 chance of losing \$2, and a 1/37 chance of winning $35 \cdot \$2 = \70 . Finally, the bet on 20-21-23-24 has a 4/37 chance of winning $8 \cdot \$2$ and a 33/37 chance of losing \$2. We easily compute as before that these bets have the following EV:

$$\begin{aligned} EV(\$3 \text{ on red}) &= \frac{1}{37} \cdot (-\$3) \approx -\$0.081 \\ EV(\$3 \text{ on even}) &= \frac{1}{37} \cdot (-\$3) \approx -\$0.081 \\ EV(\$2 \text{ on } 0) &= \frac{1}{37} \cdot (-\$2) \approx -\$0.070 \\ EV(\$2 \text{ on } 20-23) &= \frac{1}{37} \cdot (-\$2) \approx -\$0.070 \end{aligned}$$

If we add up these numbers we get

$$EV(\text{combination bet}) = -\$0.081 - \$0.081 - \$0.070 - \$0.070 = -\$0.27,$$

which is the same as before.

Clearly, the second computation is preferable, because we reduce the problem to smaller and easier problems. But why is it valid? Imagine that instead of you making this four-way combination bet, there are four different people placing each single bet. They each have the expectation as computed above. Now if we want to know their combined EV, we should simply add their individual EVs together. This is called the *additivity property* of EV, and we observed earlier that we could use it to compute the expected total of two dice.

It's a bit disconcerting, but no matter how you place your bets on the roulette table, you're always expecting to lose 2.7% of your total bet amount. In this case, we say that in roulette, the *house advantage*, or the *house edge* is 2.7%. It's up to you to decide how fast you want to lose and how big you want the financial swings to be. Once and for all: there is no mathematical strategy to beat the game.

II.3.4 POT ODDS

In an earlier lesson, we showed how to compute certain probabilities in poker, for example the odds of completing a drawing hand on the river. However, we remarked that this is but one step in the decision making process, and that other factors bear upon the question of whether to call or fold with such a hand.

Using the concept of expected value, we can reduce many such decisions to simple calculations. As we have seen before, when we face a decision among various alternatives, we need to compute the expected value of each of the alternatives and take the one with the greatest EV. We now apply this to the poker hand we studied before: you hold $7♥8♥$, and the board on the turn shows $2♥K♣5♥T♦$. Suppose that we are confident that our opponent has a good hand, such as AK. He bets his last \$150. Should we call in the hope of drawing out? The answer is: we don't know, because we don't know what the payoffs are yet. It all depends on the size of the pot. To illustrate, we will do the calculation for various pot sizes. We already calculated that the probability of making your hand on the river is $9/46$, or slightly below 20%. This is the same in all three scenarios. Another thing which remains the same is that the EV of folding is always \$0.

- ❶ Suppose that the pot size is \$250 before your opponent bets his last \$150. That means you will have to call \$150 for a \$400 pot. We can calculate the EV of calling as follows.

$$EV = \frac{9}{46} \cdot \$400 + \frac{37}{46} \cdot (-\$150) = \$78.26 - \$120.65 = -\$42.39.$$

This is a negative number, which means that we expect to lose, on average, \$42.39 on the call. Since folding has an EV of 0, it is better to fold the hand.

- ❷ Suppose next that the pot was \$750 before your opponents bet. Then the \$150 bet will make the pot \$900. Now the EV of calling is

$$EV = \frac{9}{46} \cdot \$900 + \frac{37}{46} \cdot (-\$150) = \$176.09 - \$120.65 = \$55.44.$$

This number is positive; we expect to make, on average, \$55.44 on the call, so calling is better than folding.

- ❸ Finally, suppose that the pot is \$466.67 before the bet. Then the total pot comes to \$616.67, giving

$$EV = \frac{9}{46} \cdot \$616.67 + \frac{37}{46} \cdot (-\$150) = \$0.$$

This is the *break-even point*, where calling has an expected value of 0 and hence is equally good as folding.

Again, several caveats are in order. First, we have assumed reasonably accurate knowledge of our opponent's hole cards. If you don't know your opponent's cards, then there is a possibility that your flush could lose to a full house. There is also the possibility that your flush loses to a higher flush. This is a small chance, but it makes the expectation of calling slightly lower. On the other hand, if there is a slight possibility that your opponent doesn't have much and that a 7 or 8 might win you the pot, then that improves the expectation of a call.

Second, we made the simplifying assumption that there was no money left to bet on the river. If there were money left, then this would be to your advantage: if you miss your draw you simply give up on the river, losing the same (\$150) but if you make your flush, then you might be able to extract a bit more from your opponent. The prospect of getting paid off when you make your hand is usually referred to as *implied odds*.

Third, we should stress that it is completely irrelevant (at least from an EV perspective) how much money you've put in the pot up to the point where we started our analysis. All that matters is what the best play is now that we're at this point in the hand. Many people have a tendency to get attached to their investments and to throw good money after bad. This is a form of entrapment which you should avoid.

II.3.5 ON BLUFFING

Without bluffing, poker would be a boring game. Von Neumann, understanding that bluffing is a necessary component of a winning poker strategy, tried to mathematically analyse this and was thus led to game theory. However, in various cases a simply EV calculation can tell us what to do.

Let us first consider the case where we are facing a possible bluff and have to decide whether to call or fold. To make things concrete, suppose that we are holding $A\heartsuit 7\spadesuit$, and that we are looking at a board of $J\spadesuit 7\heartsuit 2\heartsuit 5\spadesuit 6\heartsuit$. The pot is \$40, and our opponent comes out betting \$30, which puts him all-in. Should we call hoping to win the \$70 pot?

Again, this can't be answered in a vacuum. We need to have some information on his possible holdings. Let us make the assumption that we know that this particular opponent is usually very cautious with medium-strong holdings, and thus that the bet signifies either a strong hand (top pair or better) or a complete bluff. Let us say call the probability that our opponent has a strong hand p . Then

$$EV(\text{call}) = p \cdot \$70 + (1 - p) \cdot (-\$30) = p \cdot \$100 - \$30.$$

We can now determine for what value of p the call is break-even. In that case, the expected value of the call would be \$0, i.e.

$$p \cdot \$100 - \$30 = 0 \Rightarrow p \cdot \$100 = \$30 \Rightarrow p = \frac{30}{100} = 30\%.$$

What do we conclude from this? Well, if we estimate the probability that our opponent has a real hand to be greater than 30%, we should fold. If we think that he is bluffing at least 30% of the time here, we should call.

The same sort of reasoning can be employed when we are deciding whether to make a bluff or not. Let's take an example from Pot Limit Omaha. We are holding $5\spadesuit 6\spadesuit 7\diamondsuit 8\diamondsuit$, a good starting hand. The

flop came $4\clubsuit 5\clubsuit K\heartsuit$, giving us a straight draw¹⁴. The turn was the $T\spadesuit$. On the river, the $2\clubsuit$ falls. We missed our big draw, and our only chance of winning the pot is by bluffing. The pot is \$150, and we have \$125 left to bet, covering our single opponent. From how the hand played, we are positive that our opponent has us beat. Should we try to bluff?

We have to ask ourselves how often a bluff is going to be successful. Just as in the previous example, there is a break-even point. Above that point, a bluff is mandatory, and below it, we should give up on the hand. To determine the break-even point, we call the probability of our opponent folding p , and calculate the expected value of the bluff:

$$EV(\text{bluff}) = p \cdot \$150 + (1 - p)(-\$125) = p \cdot \$275 - \$125.$$

This we set equal to 0, and solve for p :

$$p \cdot \$275 - \$125 = 0 \Rightarrow p \cdot \$275 = \$125 \Rightarrow p = \frac{125}{275} \approx 45\%.$$

Thus if we estimate our opponent to fold more than 45% of the time, we should bluff, otherwise we should fold.

One of the nice features of Omaha is that there are often various possible draws, and that hand values can change drastically with every new card. Even if our opponent had a set of kings on the flop, by the river he must seriously consider being behind to a straight or a flush. Bluffing on cards which complete a draw is common practice, and can be profitable especially against players suffering from the “monsters under the bed”-syndrome, i.e. who always suspect the worst.

Since situations like these frequently come up and since there is no time for detailed calculations at the table, some heuristics are useful. The following numbers are good to know (verify them for yourself).

- In order to profitably call a pot-sized bet on the river, you need to win at least 33% of the time.
- In order to profitably make a pot-sized bluff on the river, your opponent needs to fold at least 50% of the time.
- In order to profitably call a half-pot bet on the river, you need to win at least 25% of the time.
- In order to profitably make a half-pot bluff on the river, your opponent needs to fold at least 33% of the time.

What has been left open is the question on how to estimate the likelihood of your opponent bluffing or of your opponent folding to your bluff, respectively. This is where the game becomes interesting! What kind of player is your opponent? How does he view you? Does he know you know that this is a good bluffing spot? Does he think that that is making you less likely to bluff? As you can see, we’re quickly ending up in a leveling war, prompting the needs of game theory. We will revisit this topic when we know more about mixed strategies. For now, it should be stressed that in spite of the simplifications made here, the key lesson that our propensity to call a bluff is determined by an EV calculation is very valuable. It also helps keeping your head up if you make some calls only to see that the pot is going the other way: after all, you are only supposed to win 33% of the time!

¹⁴Recall that in Omaha you use exactly two of your cards and three of the community cards, so that we still need a 3,6,7 or 8.



II.4 THE TOOLS OF PROBABILITY THEORY

In the previous section we looked at some elementary applications of probability and the central notion of Expected Value. In this section we develop some basic tools for simplifying calculations, and for tackling questions which are too hard to approach naively. One of the most important results is Bayes' Theorem, which tells us how likelihoods change in the light of new information.

II.4.1 THE WORLD OF EVENTS

In various problems we encountered, we wanted to determine the probability of a certain compound event. For example, we computed the probability that a combination bet on the roulette table would win. To this end, we used a brute force method, by simply counting all the favorable outcomes. We now seek to explain easier and more intelligent ways of finding such probabilities. The first step is to take a closer look at the structure of the set of outcomes and the various events contained in it, and to investigate how relations between events translate into relations between their probabilities.

In line with common practice in probability theory, we shall refer to the set of all possible outcomes of a random experiment as the *sample space*, and denote it by \mathcal{A} . The possible outcomes themselves, i.e. the elements of the sample space are also referred to as *sample points*. Finally, we refer to a subset \mathcal{U} of \mathcal{A} as an *event*.

To illustrate this use of terminology, suppose that our random experiment consists of three successive flips of a fair coin. Then the sample space is

$$\mathcal{A} = \{\text{HHH, HHT, HTH, HTT, THH, THT, TTH, TTT}\}.$$

In this case, there are $2^3 = 8$ sample points, each corresponding to a particular sequence of results of the flips. The following are typical events which we might want to consider:

$$\begin{aligned} \mathcal{U}_1 &= \{\text{HHT, HTH, THH}\} = \text{All outcomes with exactly 2 heads} \\ \mathcal{U}_2 &= \{\text{HHH, TTT}\} = \text{All outcomes which have no alternations} \\ \mathcal{U}_3 &= \{\text{HHH, HHT, HTH, HTT}\} = \text{All outcomes starting with heads} \\ \mathcal{U}_4 &= \{\text{HHT, HTT, THH, TTH}\} = \text{All outcomes with one alternation} \\ \mathcal{U}_5 &= \{\text{HHT, HTT}\} = \text{All outcomes with one alternation starting with heads} \end{aligned}$$

On the one extreme, we have the event $\mathcal{U} = \mathcal{A}$. This is the event containing all possible outcomes. On the other extreme, there is the empty event $\mathcal{U} = \emptyset$, containing no outcomes. In total, there are 2^n events, where n is the number of sample points.

There are a couple of highly useful ways of constructing new events from old ones. For example, if \mathcal{U} and \mathcal{V} are two events, we may form their *union* $\mathcal{U} \cup \mathcal{V}$. This is the event containing those sample points

which are either in \mathcal{U} , or in \mathcal{V} , or in both. In the example of the three coinflips, we would have

$$\mathcal{U}_2 \cup \mathcal{U}_3 = \{\text{HHH}, \text{TTT}, \text{HHT}, \text{HTH}, \text{HTT}\}.$$

If we want a description of the union of two events in plain English, we use the word “or”. The event $\mathcal{U}_2 \cup \mathcal{U}_3$ may be described as the collection of those outcomes which have no alternations or which start with heads.

Similarly, we can form the *intersection* $\mathcal{U} \cap \mathcal{V}$. This is the event containing those sample points which are both in \mathcal{U} and in \mathcal{V} . In the coinflip example, we have

$$\mathcal{U}_2 \cap \mathcal{U}_3 = \{\text{HHH}\}.$$

We describe the intersection of two events using the word “and”: the event $\mathcal{U}_2 \cap \mathcal{U}_3$ may be described as the collection of those outcomes which have no alternations and which start with heads. Of course, a simpler description is to say that it is the collection consisting of the single outcome HHH.

There are two more constructions of interest. The first is called *complementation*. Given an event \mathcal{U} , we write \mathcal{U}^c for the event which contains precisely those sample points which are not in \mathcal{U} . For example, we have

$$\mathcal{U}_3^c = \{\text{THH}, \text{THT}, \text{TTH}, \text{TTT}\}.$$

The corresponding English word is “not”: \mathcal{U}_3^c is the event consisting of those sequences which do not start with heads.

Finally, we may, from \mathcal{U} and \mathcal{V} , form a new event $\mathcal{U} - \mathcal{V}$ called the *difference* of \mathcal{U} and \mathcal{V} , consisting of those sample points which are in \mathcal{U} but not in \mathcal{V} . For example,


$$\mathcal{U}_3 - \mathcal{U}_2 = \{\text{HHT}, \text{HTH}, \text{HTT}\},$$

the event consisting of those sequences which start with heads but have at least one alternation. Note that by definition, we have $\mathcal{U} - \mathcal{V} = \mathcal{U} \cap \mathcal{V}^c$.

Events and the operations thereon can be conveniently visualised using *Venn Diagrams*. The outer rectangle represents the sample space; events are simply regions inside the rectangle. Figure II.2 summarizes the operations on events, and illustrates their effect in a Venn Diagram. In all diagrams, the doubly shaded region is the one being defined.

II.4.2 THE LAWS OF PROBABILITY

The laws of probability explain how relations between events translate into relations between the corresponding probabilities. For simplicity, we write $P(\mathcal{U})$ for $P(X \in \mathcal{U})$.




Union

Definition: $U \cup V = \{a \in \mathcal{A} | a \in U \text{ or } a \in V\}$

English: The outcomes in U **or** in V (or both)

Diagram: Region covered by U and V together




Intersection

Definition: $U \cap V = \{a \in \mathcal{A} | a \in U \text{ and } a \in V\}$

English: The outcomes both in U **and** in V

Diagram: Overlap between U and V




Complement

Definition: $U^c = \{a \in \mathcal{A} | a \notin U\}$

English: The outcomes which are **not** in U

Diagram: Everything outside U



Difference

Definition: $U - V = \{a \in \mathcal{A} | a \in U \text{ but } a \notin V\}$

English: The outcomes in U **but not** in V

Diagram: U minus the overlap between U and V

Figure II.2: Union, Intersection, Complement and Difference of Events.

Laws of Probability

- ① $P(\emptyset) = 0$
- ② $P(\mathcal{A}) = 1$
- ③ $0 \leq P(\mathcal{U}) \leq 1$ for all events \mathcal{U}
- ④ $P(\mathcal{U}^c) = 1 - P(\mathcal{U})$
- ⑤ $P(\mathcal{U} \cup \mathcal{V}) = P(\mathcal{U}) + P(\mathcal{V}) - P(\mathcal{U} \cap \mathcal{V})$
- ⑥ $P(\mathcal{U} \cup \mathcal{V}) = P(\mathcal{U}) + P(\mathcal{V})$ when $\mathcal{U} \cap \mathcal{V} = \emptyset$
- ⑦ If $\mathcal{U} \subseteq \mathcal{V}$ then $P(\mathcal{U}) \leq P(\mathcal{V})$
- ⑧ $P(\mathcal{U} \cap \mathcal{V}) + P(\mathcal{U} \cap \mathcal{V}^c) = P(\mathcal{U})$
- ⑨ If $\mathcal{U}_1 \cup \dots \cup \mathcal{U}_k = \mathcal{A}$, and if $\mathcal{U}_i \cap \mathcal{U}_j = \emptyset$ for all $i \neq j$, then

$$P(\mathcal{U} \cap \mathcal{U}_1) + \dots + P(\mathcal{U} \cap \mathcal{U}_1) = P(\mathcal{U})$$

Let us see why these must hold. First, it is impossible that no outcome occurs. Thus the probability of the empty set of outcomes is 0. Since we are certain that one of the outcomes occurs, the probability of \mathcal{A} itself is 1. The third item says that the probability of any event is a number between 0 and 1. This is true because we defined the probability to be the ratio between the favorable outcomes and the possible outcomes, and this ratio must be between 0 and 1.

Item 4 says that the probability that an event doesn't happen is 1 minus the probability that it does happen. For example, the probability that we don't throw a 3 with a die is

$$P(X \neq 3) = P(X \in \{1, 2, 4, 5, 6\}) = \frac{5}{6} = 1 - \frac{1}{6} = 1 - P(X = 3).$$

Item 5 explains how to compute the probability of the union of two events. We cannot simply add the individual probabilities, because then the outcomes in the overlap are counted twice. Thus we need to subtract those from the total. For example, if we place a bet on red and on even in roulette, we have

$$P(\text{red or even}) = P(\text{red}) + P(\text{even}) - P(\text{red and even}) = \frac{18}{37} + \frac{18}{37} - \frac{9}{37} = \frac{27}{37}.$$

(Here we used the roulette layout to find that there are 9 numbers which are both red and even.)

Item 6 is an important consequence of item 5: if the two events are *mutually exclusive*, meaning that they cannot happen at the same time, then the probability that one of them happens is simply the sum of the individual probabilities. For example, the chance that a randomly drawn card from a deck is either a club or a red ace is

$$P(\text{club or red A}) = P(\text{club}) + P(\text{red A}) = \frac{13}{52} + \frac{2}{52} = \frac{15}{52}.$$

Item 7 says that if one event contains at least as many favorable outcomes as another, then it is at least as likely.

Item 8 is a useful computational tool. It tells us that we may compute the probability of an event by splitting the problem up in two cases. The two cases are given by considering another event \mathcal{V} . From the Venn diagrams one easily sees that \mathcal{U} splits up nicely into $\mathcal{U} \cap \mathcal{V}$ and the difference $\mathcal{U} - \mathcal{V}$.

Finally, Item 9 is a generalization of 8; it says that if one of the events $\mathcal{U}_1, \dots, \mathcal{U}_k$ must occur and when no two of them can happen at the same time, then we can reason by cases, where each event \mathcal{U}_i is one of the possible cases.

Example II.4.1. An unbiased coin is flipped ten times. The number of *alternations* of the resulting sequence is the number of times the outcome changes from H to T or vice versa. For example, HHT-THTHTTT has 5 alternations, while HHHTTTTHHH has 2. By a *run*, we mean a segment of the sequence without alternations. For example, HHHTTTHHTTT has four runs. (Clearly, the number of runs is the number of alternations plus one.)

What are the probabilities of the following events?

- (i) The sequence has no alternations.
- (ii) The sequence has at least one alternation.
- (iii) The sequence has exactly one alternation.
- (iv) The sequence has at most one alternation.
- (v) The sequence has a run of exactly 6 Hs.
- (vi) The sequence has a run of at least 6 Hs.
- (vii) The sequence has a run of length 6 or more.
- (viii) The longest run in the sequence is at most 5.

Note first that the total number of possible sequences is $2^{10} = 1024$.

- (i) If the sequence has no alternations, then it consists either of only Hs or of only Ts. Thus there are two favorable outcomes, and hence $P(\text{no alternations}) = \frac{2}{1024} = \frac{1}{512}$.
- (ii) This event is the complement of the previous one, so $P(\text{at least one alternation}) = 1 - P(\text{no alternations}) = 1 - \frac{1}{512} = \frac{511}{512}$.
- (iii) We can count the favorable outcomes by hand. The alternation can occur in 9 spots in the sequence (between 1 and 2, between 2 and 3, and so on). Thus there are 18 sequences with one alternation (we need to count twice because we can either start with H or with T). Thus $P(\text{exactly one alternation}) = \frac{18}{1024}$.
- (iv) Here we can use that “at most one alternation” is the union of “exactly one alternation” and “no alternations”. These two events are mutually exclusive, so

$$P(\text{at most one alt.}) = P(\text{exactly one alt.}) + P(\text{no alt.}) = \frac{2}{1024} + \frac{18}{1024} = \frac{20}{1024}.$$

- (v) We split the problem up in different cases, depending on where the run occurs in the sequence. The lower case x, y, z indicate that the value in the sequence can be either H or T.

$$\begin{aligned}
 P(\text{run of exactly 6 Hs}) &= P(\text{HHHHHHT}xyz) \\
 &+ P(\text{THHHHHHT}xy) \\
 &+ P(x\text{THHHHHHT}y) \\
 &+ P(xy\text{THHHHHHT}) \\
 &+ P(xyz\text{THHHHHH})
 \end{aligned}$$

Now how do we find $P(\text{HHHHHHT}xyz)$, say? Well, we count the number of such sequences. The first 7 values are fixed; the remaining three can be chosen arbitrarily. That means we have $2^3 = 8$ such sequences, and hence the probability of getting one of these is $\frac{8}{1024}$. Similarly, we get $P(x\text{THHHHHHT}y) = \frac{4}{1024}$, and so on. This gives

$$P(\text{run of exactly 6 Hs}) = \frac{8}{1024} + \frac{4}{1024} + \frac{4}{1024} + \frac{4}{1024} + \frac{8}{1024} = \frac{28}{1024} \approx 0.027$$

- (vi) Just as in the previous example, we may count the number of sequences with a run of exactly 7 heads: there are $4 + 2 + 2 + 4 = 12$ of those. There are $2 + 1 + 2 = 5$ sequences with a run of exactly 8 heads, 2 sequences with a run of exactly 9 heads and 1 sequence with a run of 10 heads. Therefore:

$$P(\text{run of at least 6 H}) = \frac{28 + 12 + 5 + 2 + 1}{1024} = \frac{48}{1024} \approx 0.0469$$

- (vii) We split this up in two cases:

$$P(\text{run of at least 6}) = P(\text{run of } \geq 6 \text{ H}) + P(\text{run of } \geq 6 \text{ T}) = \frac{48}{1024} + \frac{48}{1024} \approx 0.0938$$

- (viii) To have no run longer than five is the same as not having a run of at least six. Thus


$$\begin{aligned}
 P(\text{no run longer than 5}) &= 1 - P(\text{run of at least 6}) \\
 &= 1 - P(\text{run of } \geq 6 \text{ H}) - P(\text{run of } \geq 6 \text{ T}) \\
 &= 1 - 0.0496 - 0.0496 = 0.9062
 \end{aligned}$$

II.4.3 CONDITIONAL PROBABILITY

Often, one has estimates or exact knowledge of certain probabilities, and then one acquires a new piece of information which influences those probabilities. How should one revise the estimates in the light of this new information? This is the subject matter of conditional probability.

Example II.4.2. Suppose the weather forecast said last night that there would be a 25% chance of rain showers today. However, when we look out of the window we see dark storm clouds congregating above us. We should now revise our probability estimate that we will get rain.

Example II.4.3. Last week, you believed that the liberal party was likely to win the next elections. But in yesterday's poll, they only got 35% of the votes. You now need to adjust your expectations.



Conditional Probability

- Definition:** $P(\mathcal{U}|\mathcal{V}) = \frac{|\mathcal{U} \cap \mathcal{V}|}{|\mathcal{V}|} = \frac{P(\mathcal{U} \text{ and } \mathcal{V})}{P(\mathcal{V})}$
English: The probability of \mathcal{U} **given** \mathcal{V}
Diagram: \mathcal{V} becomes the new sample space.

Figure II.3: Restricting the Sample Space

Example II.4.4. Before our opponent reraised us, the chances of him having aces were very small, and he could have a wide range of hands. However, after the reraise we must narrow down his range and in particular revise our estimate that he has AA.

First, we introduce notation for these kind of probabilities, which are called *conditional probabilities*: we write

$$P(\mathcal{U}|\mathcal{V}) = \text{the probability of } \mathcal{U} \text{ given } \mathcal{V}.$$

How should we compute such a probability? The idea is that if we assume that \mathcal{V} has occurred, then we have effectively narrowed down the sample space to the outcomes in \mathcal{V} . We then focus on what remains of \mathcal{U} once we restrict to \mathcal{V} : this is $\mathcal{U} \cap \mathcal{V}$. This is almost what we need. However, we want that the probability of \mathcal{V} given \mathcal{V} equals 1; therefore we need to rescale everything down to the proportions of \mathcal{V} . This is achieved by dividing by $P(\mathcal{V})$.

$$P(\mathcal{U}|\mathcal{V}) = \frac{P(\mathcal{U} \cap \mathcal{V})}{P(\mathcal{V})}.$$

This can also be understood by going back to Cardano's Law: the probability that \mathcal{U} occurs assuming that \mathcal{V} occurs is the number of outcomes in \mathcal{U} that are also in \mathcal{V} divided by the number of outcomes in \mathcal{V} . See figure II.3.

$$\text{Note that } P(\mathcal{U}|\mathcal{V}) = \frac{|\mathcal{U} \cap \mathcal{V}|}{|\mathcal{V}|} = \frac{|\mathcal{U} \cap \mathcal{V}|/|\mathcal{A}|}{|\mathcal{V}|/|\mathcal{A}|} = \frac{P(\mathcal{U} \cap \mathcal{V})}{P(\mathcal{V})}.$$

Some examples will clarify this concept. Suppose we throw a die, and let \mathcal{U} be the outcomes less than or equal to 3, and let \mathcal{V} be the odd outcomes. Then $P(\mathcal{U}|\mathcal{V})$ denotes the probability of getting a number less than or equal to 3 when it is given that the outcome is odd. Then $\mathcal{U} \cap \mathcal{V} = \{1, 3\}$. We compute

$$P(X \leq 3 | X \text{ odd}) = \frac{P(X \leq 3 \text{ and } X \text{ odd})}{P(X \text{ odd})} = \frac{2/6}{3/6} = \frac{2}{3}.$$

Note that this is different from $P(\mathcal{U}) = \frac{3}{6}$.

Next, suppose we place a roulette bet on red. The chance of winning, as we've seen before, is $\frac{18}{37}$. However, suppose we can't see the outcome, but we see that our friend, who bet on the 2nd 12, is happy to have won. What are the chances we won? Well, there are 5 red numbers in the 2nd 12 (see layout). So

$$P(\text{red} | 2\text{nd } 12) = \frac{P(\text{red and } 2\text{nd } 12)}{P(2\text{nd } 12)} = \frac{5/37}{12/37} = \frac{5}{12} \approx 0.417.$$

Thus learning that the ball landed in the 2nd 12 actually diminished our chances of winning on red.

Conditional probability is of central importance to the correct interpretation of medical test results. Suppose your doctor runs a test on you, and tells you that you have a rare but very serious disease. How rare? Well, let's assume that only one in 5,000 people is affected by it. You wonder whether the test could give the wrong outcome, but the doctor assures you that the test is 99% reliable, meaning that only one in 100 people without the disease actually test positive (this is called a false positive). You go home depressed, thinking that it is 99% certain that you are doomed. But are you? Let's think again. Before you got the bad news, your chances of having the disease were $1/5,000 = 0.0002$. (This is called the *prior probability*, the chances you assigned before the new information came to light.) Now suppose we consider a population with 1,000,000 people. About 200 of those carry the disease. Imagine we test the entire population. The 200 people who have the disease will test positive. Of the remaining healthy 999,800, one percent will test positive. Thus in total, we have $200 + 9,998 = 10,198$ positive test results. Upon learning that you tested positive, you're now in this group. However, within this group only 200 out of 10,198 really have the disease. Therefore the chances of having the disease are only $200/10,198 = 1.98\%$.

The important lesson is that one cannot conclude anything from the reliability of the test alone. The prior probability needs to be taken into account. If the disease were much more common, say if the prior probability was 1%, then the likelihood that you have the disease are almost 50%.

We calculated the previous problem using an example population. The calculation is correct; however, we would like to know the principle behind it. It is called *Bayes' Theorem*, or *Bayes' Rule*. In its simplest form, it reads:

Bayes' Theorem

$$P(\mathcal{U}|\mathcal{V}) = \frac{P(\mathcal{U})P(\mathcal{V}|\mathcal{U})}{P(\mathcal{U})P(\mathcal{V}|\mathcal{U}) + P(\mathcal{U}^c)P(\mathcal{V}|\mathcal{U}^c)}$$

This formula helps when it is not easy to determine the probability of $\mathcal{U} \cap \mathcal{V}$ directly.

Let us apply this to the previous problem and see that it validates our earlier informal reasoning.

$$\begin{aligned} P(\text{disease}|\text{positive}) &= \frac{P(\text{disease})P(\text{positive}|\text{disease})}{P(\text{disease})P(\text{positive}|\text{disease}) + P(\text{no disease})P(\text{positive}|\text{no disease})} \\ &= \frac{1/5,000 \cdot 1}{1/5,000 \cdot 1 + 4,999/5,000 \cdot 1/100} = \frac{100}{5,099} = 1.98\%. \end{aligned}$$

II.4.4 INDEPENDENCE AND THE GAMBLER'S FALLACY

Intuitively, two events are *independent* if the occurrence of one does not have any influence on the probability of the other, and vice versa. For example, if we flip a fair coin twice, then the outcome of

the first has no influence on the outcome of the second, so the two tosses are independent. However, if we draw two cards from a deck, then the two outcomes are not independent, because the removal of the first changes the probabilities for the outcomes of the second.

It is easy to make precise what independence is: two events \mathcal{U} and \mathcal{V} are independent when $P(\mathcal{U}|\mathcal{V}) = P(\mathcal{U})$. Thus, knowing that \mathcal{V} occurred does not change the probability of \mathcal{U} occurring. Of course, this notion is symmetric in \mathcal{U} and \mathcal{V} , so it can be equivalently expressed as $P(\mathcal{V}|\mathcal{U}) = P(\mathcal{V})$. There is however one slight problem: one of the probabilities involved might be zero, and we can't divide by zero. Therefore we reformulate, using the fact that $P(\mathcal{U} \cap \mathcal{V}) = P(\mathcal{U})P(\mathcal{V}|\mathcal{U})$. Then we are led to the following definition:

Two events \mathcal{U}, \mathcal{V} are (*statistically*) *independent* when

$$P(\mathcal{U} \cap \mathcal{V}) = P(\mathcal{U})P(\mathcal{V}).$$

As an example, consider an urn containing 25 balls, 10 white and 15 black. We draw three balls one by one, placing back the ball after each draw. What are the chances that we get BWB? Because we replace the balls after each draw, the individual draws are independent: the outcome of the first doesn't influence the outcome of the second, and so on. Thus the probability is

$$P(\text{BWB}) = P(\text{B}) \cdot P(\text{W}) \cdot P(\text{B}) = \frac{15}{25} \cdot \frac{10}{25} \cdot \frac{15}{25} = \frac{18}{125} = 0.144$$

If, however, we do not replace balls then we get

$$\begin{aligned} P(\text{BWB}) &= P(\text{B}) \cdot P(\text{W}|\text{first is B}) \cdot P(\text{B}|\text{first two are BW}) \\ &= \frac{15}{25} \cdot \frac{10}{24} \cdot \frac{14}{23} \approx 0.152 \end{aligned}$$

One of the more serious errors often made is treating independent events as if they were dependent. This is called the *Gambler's Fallacy*: the perpetrator of this reasoning mistake incorrectly believes that the fact that the probability of an event happening are influenced by another event, whereas in reality these events don't influence each other. Casino visitors tend to be prone to this fallacious reasoning (hence the name of the fallacy): for example, people often believe that just because the previous five spins of the roulette wheel fell on black, it is now more likely that red will fall on the next spin. This, of course, is incorrect because the wheel has no memory: each spin is independent of all the ones before it. Casinos stimulate this incorrect reasoning by displaying the results of the last ten spins or so above each roulette table. Even worse is the situation for baccarat, where players are given sheets on which they can record the outcomes of all the previous rounds of play so that they can work out a "system". The reason why casinos do this, is because it is known that players who believe that they might be able to spot patterns and come up with a winning system are much more engaged in the gambling activity and hence much less likely to stop playing.

Let's look at the Gambler's Fallacy a bit closer, because there are several interesting aspects to it. First of all, there are actually two versions of the fallacy, which are sometimes referred to as a *Type 1*

and a *Type 2* Gambler's Fallacy. The first one is the more common one, and is the misconception that because a certain outcome of an experiment has come up more than expected, it is now more likely that another outcome will appear. Thus, if we do ten successive coinflips and the first eighth land on Tails, a Type 1 Gambler's Fallacy would be to believe that Heads is more likely than Tails on the ninth and tenth toss. A Type 2 fallacy also ignores independence of events, but errs in the opposite direction. It mistakenly concludes from the fact that a certain outcome has come up more than expected, that this pattern will continue in the future and that that same outcome is more likely than what probability theory says it is. In the example of the ten coinflips, a Type 2 fallacy would be to predict that Tails are more likely to appear on the ninth and tenth toss.

The actual prediction of outcomes is not the only situation where gamblers fall prey to both types of fallacies. For example, you will often hear gamblers say that they are in the middle of a "winning streak" or a "losing streak". For some, this just means that they have won more than their fair share over their past couple of gambling sessions (or lost more, in case of a losing streak) and this may just be an accurate factual statement. But more often than not, they tend to draw the unwarranted conclusion that their state of being (un)lucky is likely to be continued in the near future: if they are on a winning streak now, it's time to keep playing before their luck runs out! And if they are going through a losing streak, it's probably best to ride it out playing small stakes until their luck returns. Both these types of reasoning are an instance of the Type 2 fallacy. Some gamblers also draw opposite conclusions: when they are on a losing streak, they will start betting bigger because they are now due for a win. That is the Type 1 fallacy at work.

II.4.5 INTERMEZZO: HEURISTICS AND BIASES

Why is the Gambler's Fallacy so widespread and pervasive? And why is it that even when we learn about independence of events, there still is this nagging feeling that after seeing a long run of spins landing on red, we feel that black is somehow due? An explanation for this, and many other fallacies as well, is given by the so-called *Heuristics and Biases* framework, due to cognitive psychologists Amos Tversky and Daniel Kahneman.

Our hunting and gathering ancestors could not afford the luxury of taking a lot of time for detailed estimations and calculations of probabilities. When a large animal is dashing towards you you have to decide instantly whether to fight or to flee. From a survival perspective, deciding quickly is often more important than deciding very accurately.

When we need to judge situations and make decisions without the time for detailed analysis we tend to resort to heuristics. *Heuristics* are simple rules of thumb which, though not fully accurate, allow for a reasonable approximation. For example, when we wish to estimate the probability that crossing the street is safe, we don't proceed by calculating velocity vectors for all vehicles approaching us, and then deducing whether our own speed will be sufficient to avoid collision. Rather, we use our experience and intuition to tell us whether the probability is below a certain threshold. We are virtually always successful in this strategy. Similarly, in order to decide whether to fight or flee, our ancestor might simply use the fact that the last three times he saw a similar animal it always turned out to be false alarm.

Even when we could in theory try to find an optimal solution to a problem, the price of the time and energy this would cost is much too high compared to giving up a bit of utility because our quick

heuristic solution is not 100% optimal. Herbert Simon summarized this idea in the following slogan: “Humans are satisficers, not optimizers”. We rather have a satisfactory solution reasonably quickly than spend great effort finding an optimal one.

While the usefulness of such heuristics is undeniable, there are downsides as well. In many situations, heuristics lead us astray, giving rise to *biases* in our judgment. These are deeply rooted in our way of thinking, and can often only be overcome by explicit awareness and training. The heuristics and biases program explains various judgment and reasoning mistakes in terms of the heuristics we employ and why they are not appropriate for the problem at hand. It turns out, that in situations where abstract probability theory is involved, heuristics very often lead to the wrong answer. Below we briefly discuss three important heuristics: anchoring, availability and representativeness, and we see how they give rise to misjudgment of probabilities.

Anchoring

Suppose I ask you whether you think the probability of a military coup in Colombia within the coming decade is larger than or smaller than 25%, and to give your best estimate of that probability. Even though you may not follow Colombian politics very closely, try to make an educated guess. Now suppose I had asked instead whether you think the probability of a military coup in Colombia within the coming decade is larger than or smaller than 2%, and to give your best estimate of the probability. Do you think your estimate would have been the same?

One would think that the first part of the question should play no role in the answer to the second part. Theoretically, the answer to the first part should be logically deduced from the second. However, in practice the average estimate found among subjects answering the version with the higher percentage (25%) will be much higher than for the version with the lower percentage (2%). The explanation is that the mentioned percentage will influence the estimate, because it provides a frame of reference. This is called the *anchoring effect*. Anchoring can be useful, because it is hard to make accurate judgments without a frame of reference. But we often forget that the frame of reference we use can be inaccurate.

Sales people use this to their advantage by putting a high sticker price on items, so that the discount looks more impressive. The customer, instead of focusing on whether the offered price is reasonable, feels that he’s getting a great deal. Similarly, in negotiations one often observes one or both parties making an unreasonable demand. As a consequence, the other party will perceive a substantial concession as big progress, rather than concentrating on the merits of the current offer.

One might object that guessing probabilities concerning Colombian politics is too much of a crapshoot because most of us lack relevant information about the situation. So take instead a subject which you ought to have at least some knowledge: yourself. Do you think you have a larger or smaller chance than average (within a group of comparable subjects, say your fellow classmates in this course) of becoming a drug addict? Of getting a good job shortly after graduation? Of having a long-lasting and satisfying relationship? While people vary greatly in their optimism or pessimism about their own future, empirical studies show that most people estimate their chances of positive things happening to them as bigger than average, while negative events are judged to be less likely than average. This is called the *valence effect*: we tend to overestimate probabilities of positive events, while underestimating those of negative events.

A cruder manifestation of this phenomenon was observed in an experiment in which subjects were

presented with a sequence of cards, containing either a picture of a smiling face or a picture of a sad face. In one version, the proportion of happy to angry faces was 70-30. Subjects were quite accurate at guessing this proportion. In another version, however, the proportion was reversed; this caused subjects to overestimate the proportion of happy faces.

Availability

Another possible source of distortion of probability estimates is the fact that we tend to place unwarranted faith in information which is readily available to us. Judging the likelihood of an event according to the ease with which occurrences are brought to mind is called the *availability heuristic*.

As an example of a bias arising from this heuristic, consider the following question: which of the following two is a more frequent cause of death in North America, stomach cancer or car accidents? Most people think that car accidents cause more deaths, but they are wrong. Car accidents are reported in the media very frequently, making it much easier to recall examples, while stomach cancer gets less publicity but kills many more people.

Experiments have shown that people judge events to be more likely when they are instructed to imagine them in detail. Also, when events are difficult to imagine or to visualise, probability estimates are low. Events described in great and vivid detail are also deemed more likely than nondescript events; for example, a study showed that legal evidence which is presented with vivid detail is given more weight than when its presentation is more pallid. Finally, there is the *recency effect*: events which are still fresh in our memory are deemed more likely than those which happened further in the past.

Representativeness

The third type of heuristics we discuss is called the *representativeness heuristic*. The representativeness heuristic is a rule of thumb which tells us to judge the probability of an event by how representative that event is of the general population or process. In plain language: if something looks more typical, then it is more likely.

In a famous experiment which nicely illustrates the representativeness heuristics at work, Tversky and Kahneman presented subjects with a description of a woman named Linda:

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations.

Subjects were then asked to check which of the following statements about Linda was most likely to be true:

1. Linda is a bank teller.
2. Linda is a bank teller and is active in the feminist movement.

Most people think option 2) is more likely. Why? Because given what we know about Linda, this option seems to fit best. It is *more representative* of the type of person Linda is. But the laws of

probability are unforgiving: the probability of A and B happening are always smaller than or equal to the probability of A happening. There are more bank tellers than there are bank tellers active in the feminist movement, so the chances of Linda belonging to this smaller group are smaller.

The representativeness heuristic also explains the Type 2 Gambler's fallacy. Indeed, this fallacy may be reformulated as "the mistaken belief that a small sample is representative of the overall population".

In order to explain the Type 1 Gambler's Fallacy, Tversky and Kahneman introduce the concept of "local representativeness". This refers to the idea that a small sample from a large series of experiments should resemble the population. In the case of coinflips, for example, we know that if we perform a large number of flips of a fair coin, the ratio between Heads and Tails will be approximately 1/2. (In probability theory, this is called the Law of Large Numbers.) However, a small sample, of eight flips say, need not resemble this, and it can easily have a very different proportion of Heads and Tails. Local representativeness is the mistaken belief that the small sample should be similar to the large one. This is precisely the Type 1 Gambler's fallacy: the incorrect belief that chance is "self-correcting". Thus a Type 1 fallacy starts with knowledge of, or assumptions about, the general population or process at hand, and draws incorrect conclusions about how small samples should behave. Note that this is precisely opposite to what happens in a Type 2 fallacy, where from a small sample an unwarranted conclusion about the population as a whole is drawn.

A lot of research has been done in the area of heuristics and biases. A good start is Scott Plous' book *The Psychology of Judgment and Decision Making*. For a collection of representative research papers, the reader may wish to consult *Judgment under Uncertainty: Heuristics and Biases*, by Kahneman, Slovic and Tversky, as well as the more recent collection *Heuristics and Biases*, by Gilovich, Griffin and Kahneman.

II.4.6 PROBLEMS REVISITED

We have more than enough theory under our belt to answer the questions posed at the beginning of this chapter.

Monty Hall. Yes, you should switch. One way of seeing this is as follows. Initially, you have a 1/3 chance of guessing the right door. In case you guessed right, switching will cost you the prize. In the remaining 2/3 of cases, switching is guaranteed to win. Thus if you always switch, you win 2/3 of the time, and you only lose the in those cases when you guessed right to begin with.

It should be pointed out that there is an assumption at work about the role of the host of the quiz. It is assumed that (a) the host knows behind which door the car is, (b) always opens one of the doors, regardless of whether you picked right or not, and (c) always opens a door with a goat. Part of the enormous discussion this problem generated can be traced back to an ambiguous description of the role of the host. A clearer way might be to say that the host follows the following recipe: if you guess right, then he randomly opens another door. But if you guess wrong, he'll tell you honestly which door is correct.

If you're still not convinced, try this: imagine the same problem, but now with 1000 doors. You pick one, and the host opens 998 doors with goats behind them. Again, if you guessed wrong initially (and now there is a 999/1,000 chance that that was the case), the host tells you where to find the car.

Presumably the hordes of mathematicians and other PhDs who argued that there is no point in switching had the following conditional probability computation in mind:

$$P(\text{door 1}|\text{not door 2}) = \frac{P(\text{door 1 and not door 2})}{P(\text{not door 2})} = \frac{1/3}{2/3} = \frac{1}{2}.$$

This would be correct when the host would always open a door at random. However, as we explained before, the host gives away crucial information.

Boy or Girl? This is a problem which we can solve by carefully looking at the sample space. For a couple with two children, there are four possibilities: boy-boy, boy-girl, girl-boy and girl-girl. Here, we write the firstborn first. Each of the four are assumed to be equally likely.

For the first part of the question, we are given that the oldest is a boy. That prunes the sample space to {boy-boy,boy-girl}. These are still equally likely, so there is a 1/2 chance that the second child is a boy as well. A formal verification using conditional probabilities is easy.

For the second part, we only learn that one of the two is a girl. This reduces the sample space to {boy-girl,girl-boy,girl-girl}, each of which is equally likely. Thus there is a 1/3 chance that the other is a girl as well.

Finally, we learn that one of the two is a girl called Kimberly. Does the name really matter? Strangely enough, it does. Let us write girlK for a girl called Kimberly, and girlNK for a girl with any name other than Kimberly. Initially, the sample space then consisted of the following:

boy-boy	boy-girlNK	boy-girlK
girlNK-boy	girlNK-girlNK	girlNK-girlK
girlK-boy	girlK-girlNK	girlK-girlK

After removing those outcomes without any Kimberly's in them what remains is

$$\{\text{boy-girlK,girlK-boy,girlNK-girlK,girlK-girlNK,girlK-girlK}\}.$$

We may assume that no sane parent calls both their daughters Kimberly, so the last option may be assigned probability 0. The four remaining options are equally likely again, which means that there the chance of the other child being a girl are 1/2. (If you insist that it is possible to have both daughters named Kimberly then you should assign a small probability to that and factor that into the computation.)

Birthday Problem. The easiest solution here is not to compute the probability that two people have the same birthday, but the probability that no two people have the same birthday. If there's only one person in the room, then that probability is 1. If there are two, then there is a $\frac{364}{365}$ chance that the second person doesn't have the same birthday as the first. Enter a third person; there is now a $\frac{364}{365} \frac{363}{365}$ chance that his/her birthday is different from the other two. And so on: with n persons in the room, the chances of two sharing a birthday are

$$1 - \frac{364}{365} \frac{363}{365} \dots \frac{365 - n + 1}{365}.$$

We can now plug in any value for n we like. For $n = 23$, the probability is about 50%. For $n = 53$, we can already be 99% certain.

Doping Test. This is a typical application of conditional probability, similar to the one discussed in the section thereon. One thing should be clear: no answer is possible unless we are given the prior probability, i.e. the percentage of doping-using athletes. Suppose you expect 10% of athletes to be using doping. Then we can invoke Bayes' Rule (write g for guilty, ng for not guilty, p for positive):

$$P(g|p) = \frac{P(g)P(p|g)}{P(g)P(p|g) + P(ng)P(p|ng)} = \frac{1/10 \cdot 1}{1/10 \cdot 1 + 9/10 \cdot 1/20} = 18.18\%.$$

A Legal Case. The defense was correct that the chances of an abuser ending up a murderer are very small (roughly 1 in 2500). However, this is not the probability we're interested in. This applies to a random couple where the man is abusive. Simpson and Smith were no longer a random couple, since it had come to light that Smith had been murdered. That information changes a lot. The relevant statistic is the percentage of murdered women with an abusive spouse who were actually killed by their spouse: this is around 90%.

CHAPTER III

MIXING IT UP

Having looked at games and at probability, it is now time to put the two together. For one thing, this allows us to study games in which chance or uncertainty plays a role. But equally importantly, the use of randomization is the key to understanding games in which pure strategies are not good enough. Everyone who has ever played Rock-Paper-Scissors knows that if you always play the same throw, then your opponent will sooner or later pick up on this and start exploiting your predictability. Thus, you need to be unpredictable, and the way to do that is by randomizing. (Expert RPS players will probably object that producing truly random throws is impossible, and that it is better to have series of random-looking throws memorized, but my point about having to be unpredictable stands.)

The mixing of strategies is one of the most important ideas in game theory and is essential for poker and many other games. Nash' famous equilibrium theorem says that provided players are allowed to mix their strategies then every finite player game has at least one Nash Equilibrium. Concretely, that means that in such games there is an optimal strategy for each player; when a player plays an optimal strategy, it is impossible for other players to exploit him/her.

In the first section of this chapter we first consider a number of games where mixed strategies are easily found and understood. We then turn to the question of how we should find such strategies. As it turns out, this is easy for 2x2 games and boils down to solving two linear equations in two unknowns. In larger games finding the Nash Equilibria can be more complicated computationally, but fortunately there are software packages that will do the hard work for you. We also consider an instructive miniature poker game called the AKQ game where mixing strategies provides us with a mathematical explanation of why, when and how often we should bluff.

Next, we turn to another aspect of games and strategy which has been mentioned before but not studied in detail: the possibility of altering the game to your advantage using strategic moves. We will first introduce several strategic moves such as threats, commitments and signalling; then we try to understand why they work and what risk there is in using them; and we see how they apply to real life situations as well as specific poker-related issues.

Finally, we look at some other aspects of the mathematics of gambling, in particular bankroll management and -growth, variance and risk of ruin. Since even the best player loses from time to time, and even the most favorable bet can go the other way, we wish to maximize our financial growth, while minimizing our risk of going broke through bad fortune. The Kelly criterion helps in balancing these two desiderata.



III.1 GAMES AND PROBABILITY

We take a first look at how uncertainty and randomness interact with game theory. There are many possible sources of uncertainty in games, and we begin by giving a brief overview. We then consider a nice game in which randomness plays a key role: a three-person duel.

III.1.1 INCOMPLETE INFORMATION

In Chapter I, we studied sequential and simultaneous games under the key assumption of *complete information*: this meant that all players had all the information relevant to the game, such as the structure of the game, the available moves, the various payoffs, and so on. In fact, we were assuming something stronger: the players didn't just know all this, they also knew that everyone else knew these things; moreover, they knew that everyone knew that everyone knew, and so on.

In many real-life situations, this assumption is highly unrealistic. Most of the time we are dealing with various degrees of uncertainty about some of the crucial aspects of the game. Sometimes we don't even realize what game we're playing! Let us therefore look in a bit more detail at the various types of uncertainty we might face, and how this influences our strategic position.

1. Some games have a built-in component of randomness. In many board games, one uses dice as a random experiment, as to introduce an element of luck into the game. Without that element of randomness most of these games would not be worth playing, or at least be much less exciting.
2. Many games also revolve around one player knowing something the others don't. In poker, only you know your hole cards; in risk, only you know your mission; and in a game such as Werewolves and Villagers, the entire point is that players don't know to which category the others belong. Note that in poker, both these aspects work together, as the hole cards are the result of a random experiment, of which only you learn the outcome.
3. Another potential source of uncertainty is about your opponents' payoffs. That is, even though you know the rules of the game, you may not be aware of your opponents' true motivations and his/her preferences. Is your poker opponent trying to win as much money as possible? Or does he care more about appearing to be a smart and seasoned pro? Part of the reason why negotiations can be complicated is that it is not always clear what the opposing party is trying to achieve, and whether their resistance to your offer is caused by the fact that they feel it is disadvantageous to

them or by the fact that accepting may be a loss of face, even though the offer seems reasonable to them.

4. Sometimes you're not aware of all the options you have. This may be something as simple as not knowing you're allowed to castle in chess, or something as serious as not knowing you have certain legal rights. It used to be difficult to negotiate the price of a new car, because only the dealer knew how much they actually paid for the car and thus how much their profit margin was. Nowadays, you can (and should!) find out exactly how much the car is worth to the dealer through various available internet services. Of course, if you're not aware of this, the salesperson is going to be all too happy not to reveal this possible move.
5. In the most extreme case you're in a game without realizing it. This happens any time you fail to realize that the situation calls for strategic thought.

Most of the time it's better to have more information. When you know more, you can make a more informed decision, while if you know less, you're more dependend on guesses and hunches. But there are notable exceptions. Sometimes, if you realize what you don't know, you can anticipate strategic problems and simply refuse to play, or play a conservative strategy which won't get you in trouble. And there are situations where having certain information can actually hurt you. Anytime you know something which is potentially damaging to others there is a risk that they will take actions against you in order to neutralize the threat of you revealing that information.

III.1.2 TRUEL

We now introduce our first game which involves randomness. It is a game theory classic called *Truel*, which is supposed to mean three-person duel. In the classic "The Good, the Bad and the Ugly" there is a famous scene where the three protagonists stare each other down for several minutes before all reaching for their guns. In our version, we will make things a bit more structured by having the players take turns shooting at each other.

The three players in this game (we'll call them The Good, The Bad and The Ugly, although the scenario will diverge from the movie plot) have different shooting abilities, which are summarized in the following table:

Player	Accuracy
The Good	1/3
The Bad	5/6
The Ugly	2/3

This means, for example, that when The Ugly tries to shoot an opponent, he will succeed 2/3 of the time. To make things a bit more fair, we let the The Good, who is the worst shooter, go first. Then it's The Ugly's turn (if he's still alive of course), and The Bad, being the most accurate shot, goes last.

Here are the complete rules:

- Players take turns, with the less accurate shooter going first.

- When it is a player's turn to act, he may aim at any target.
- Each player gets only one bullet.

Of course, we need to know the preferences. We will assume that each player does whatever he can to maximize his chances of survival. In addition, he will try to kill others, but everything else being equal is indifferent between whom he kills. We assume that if a player has the choice of killing two others and his choice does not influence his chances of survival, he will simply flip a fair coin to decide whom to shoot at.

We place ourselves in the shoes of the first shooter, The Good. At first, it seems like a good plan to try to kill The Bad. After all, it appears sensible to attempt to take out the most dangerous opponent first. There is some truth to that: if you kill The Ugly, then it's between The Bad and you, and he's such a good shooter that you're very likely to die. In order to see how good exactly the plan of aiming at The Bad is, we'll do a calculation of our chances of survival.

If you aim at The Bad, there are two cases to consider: the case where you hit (which happens with $1/3$ probability) and the case where you miss (which happens with $2/3$ probability).

Case 1: You kill The Bad, probability $1/3$ Then it's The Ugly's turn. Since you're his only remaining opponent, he'll take a shot at you. This means that with probability $1/3$ you survive.

Case 2: You miss, probability $2/3$ Then again it's the Ugly's turn, but this time he can choose between aiming at you and aiming at The Bad. Now you used your only bullet, so you are no longer a threat, whereas The Bad still has a bullet and knows how to use it. Therefore The Ugly, in order to maximize his own survival chances, will aim at The Bad. Now we again have two cases:

Case 2a: The Ugly kills The Bad, probability $2/3$ then the game ends, and you're still alive.

Case 2b: The Ugly misses, probability $1/3$ The Bad gets to move. Since both his opponents are unarmed, he'll flip a coin to decide whom to take a shot at. Then with $1/2$ probability he shoots at you, in which case you survive with $1/6$ probability. Of course, when he shoots at The Ugly, you stay alive.

We know from the previous chapter that we may calculate the probability of survival by adding up the probabilities we get in the individual cases (since they are mutually exclusive and jointly exhaustive). Within cases, we multiply probabilities. That means that the chance you survive is

$$P = \frac{1}{3} \cdot \frac{1}{3} + \frac{2}{3} \cdot \left(\frac{2}{3} \cdot 1 + \frac{1}{3} \cdot \left(\frac{1}{2} \cdot \frac{1}{6} + \frac{1}{2} \cdot 1 \right) \right) \approx 0.69$$

Can you do better than that? We already claimed that aiming at The Ugly isn't going to help.

Exercise 32. Find that the probability that you survive when you aim your shot at The Ugly.

However, there is a third option available: deliberately miss your shot! At first, this may seem silly, since you're giving up what little chance you have to take out an opponent. But let's calculate the probability of survival again by looking at what will happen. After you miss, it is The Ugly's turn, and he will aim at The Bad, just as before. And just as before, he either kills The Bad, in which case you survive, or The Bad flips a coin and shoots at one of you. Thus deliberately missing results in always ending up in Case 2. Then the probability is

$$P = \frac{2}{3} \cdot 1 + \left(\frac{1}{3} \cdot \left(\frac{1}{2} \cdot \frac{1}{6} + \frac{1}{2} \cdot 1\right)\right) \approx 0.86.$$

Thus missing deliberately raises your survival chances quite substantially!

This is counterintuitive, but there is some logic behind it: when facing two opponents who are both substantially stronger than you, it may be better to let them battle it out between them and do whatever it takes to avoid a direct confrontation with either of them.

Exercise 33. Analyse, by looking carefully at his chances of survival, why it is not optimal for The Ugly to copy your trick and deliberately miss as well after you've done so.

The shooting accuracies have a big impact on how this game should be played. To illustrate this, we change things a bit so that The Ugly is now the worst shooter:

Player	Accuracy
The Good	1/3
The Bad	5/6
The Ugly	1/4

Now The Ugly may go first. What should he do? You might think that since he's now the weakest, he should reason as in the first example, and miss on purpose. But things have changed, and we should do the math again. Let's calculate first what happens when The Ugly aims at The Bad. Again, we put you in the shoes of The Good.

Case 1: The Ugly kills The Bad, probability 1/4 That's great for you, because you survive. You also get a shot at The Ugly, and his chances of surviving that are 2/3.

Case 2: The Ugly misses, probability 3/4 Then it's your turn, but this time you can choose between aiming at The Bad or at The Ugly. The latter is unarmed, so you shoot at The Bad.

Case 2a: You kill The Bad, probability 1/3 Then the game ends, and both you and The Ugly are still alive.

Case 2b: You miss, probability 2/3 The Bad gets to move. Since both his opponents are unarmed, he'll flip a coin to decide whom to take a shot at. Then with 1/2 probability he shoots at you, in which case The Ugly survives and you survive with 1/6 probability. When he shoots at The Ugly, you stay alive and The Ugly survives with probability 1/6.

Again we must put these probabilities together, summing over cases and multiplying within cases. The probability that The Ugly survives is thus

$$P(\text{The Ugly survives}) = \frac{1}{4} \cdot \frac{2}{3} + \frac{3}{4} \left(\frac{1}{3} \cdot 1 + \frac{2}{3} \cdot \left(\frac{1}{2} \cdot \frac{1}{6} + \frac{1}{2} \cdot 1 \right) \right) \approx 0.72.$$

Now let's compare this with a deliberate miss strategy. That effectively brings us to Case 2: you get to take a shot at The Bad, and if The Bad survives that, he flips a coin and shoots at one of you. Then

$$P(\text{The Ugly survives}) = \frac{1}{3} \cdot 1 + \frac{2}{3} \cdot \left(\frac{1}{2} \cdot \frac{1}{6} + \frac{1}{2} \cdot 1 \right) \approx 0.71,$$

which is slightly lower than when The Ugly tries to hit The Bad.

Why is the situation different this time? Well, in the previous scenario there was one weak shooter (you) and two strong ones. Now, there are two weak shooters and one strong one. The two weak shooters are better off teaming up by both trying to take out the strong opponent.

Exercise 34. Explain why it is not optimal for The Ugly to shoot at you. Also explain why you shouldn't deliberately miss when it's your turn.

There are many variations on this game. Aside from varying the accuracies of the shooters (which you should experiment with to make sure you can do the calculations!) and the number of players, one can allow multiple rounds (of ammo and shooting), or one can change the preferences of the players. For example, one can stipulate that, everything else being equal, The Bad will try to kill The Good when given the chance.

Aside from being a nice brain teaser, there are real-life applications. For example, the analysis of the truel helps understand why it is often the case that an apparently less qualified candidate wins in a competition against several stronger candidates. And if you know the game show *The Weakest Link*, you now understand the strategic nuances of the voting process at the end of each round, where players are allowed to vote which of them will be sent home.



III.2 MIXED STRATEGIES

In poker, one often hears the phrase “mixing it up”. This refers to the idea of diverging from one's standard way of playing a certain type of situation, in order not to become too predictable. In many other situations mixing strategies is required as well. This section explains the logic behind the idea, how to determine when and how to mix strategies, and discusses a number of examples from game theory and poker.

III.2.1 MATCHING PENNIES

Let's look at a simple 2×2 game which illustrates the necessity of mixing strategies. This game is called *Matching Pennies*. It is a silly game: both players choose Heads or Tails. If they choose the same, Player 1 wins. If they choose differently, Player 2 wins. See Figure III.1 for the payoff matrix.

		Player 2	
		H	T
Player 1	H	-1 1	1 -1
	T	1 -1	-1 1

Figure III.1: Matching Pennies

How would you play this game? It is easily seen that there is no dominant strategy. It is also clear that no pure strategy can be part of a Nash Equilibrium: suppose Player 1 always plays Heads. Then Player 2 has a best response to that strategy, namely by always playing Tails. Player 2 is perfectly content, but Player 1 regrets his choice and would like to change.

If there are no dominant strategies or Nash Equilibria in pure strategies, then what is the best way to play, and how can you avoid being exploited by an astute opponent who is on the lookout for patterns in your play? The answer is: one needs to use a *mixed strategy*.

A *mixed strategy* is a probabilistic combination of the available pure strategies.

Concretely, that means that with certain probability one plays strategy a, with another probability one plays strategy b, and so on. In order to specify a mixed strategy one simply indicates the *relative frequencies* with which the pure strategies will be played. For example, in the matching pennies game, one of the possible mixed strategies would be:

$$35\% \text{ H}, \quad 65\% \text{ T}$$

Of course, the percentages have to add up to 100%.

How do we determine the best possible mix? For most problems, a bit of calculation is needed. But for the matching pennies problem we can use easier considerations. Suppose Player 1 is more likely to play heads than Tails. If you were Player 2, how would you respond? The answer is simple: by always playing Tails. In that case, the majority of times you would win. There is no reason for Player 2 to have a mixed strategy in this situation, for Player 1 can be exploited using a pure strategy. This shows that playing Heads more than 50% of the time is not a good idea for Player 1. By considerations of symmetry, it follows that it is never wise in this game to play a pure strategy with greater than 50% probability. The only possible conclusion is that both players have to play a mixed strategy where they mix both pure strategies in equal proportion. Note that if one player plays this 50-50 mix, the other player cannot exploit that. In fact, this mix makes the other player *indifferent between his pure strategies*.

		Player 2	
		H	T
Player 1	H	1, -1	-1, 1
	T	-1, 1	10, -10

Figure III.2: Matching Pennies, Version 2

The above reasoning shows that the best way to play the matching penny game is to flip your penny and let chance decide the outcome (assuming your penny is fair of course). There is something counterintuitive about this: we tend to think that our best decisions are those which are based on thorough reasoning leading to choosing the best option; leaving the outcome to chance seems to be tantamount to giving up on rationality. However, it is often unavoidable to introduce an element of randomness into your game plan in order to avoid exploitation.

The idea that an unexploitable strategy is one for which your opponent cannot increase his or her payoffs by switching to a certain pure strategy is a very important one, so we highlight it for emphasis:

A strategy is *unexploitable* when it makes each opponent indifferent between his or her pure strategies.

III.2.2 FINDING THE RIGHT MIX

In the matching pennies game, symmetry led us to the intuitively clear result that a 50-50 mix between both strategies is optimal. However, if the game is not as symmetric we need to do some math to find the right proportions of strategies. Suppose that in the pennies game we change the payoffs: everything stays the same, except for when both players choose Tails, then Player 1 wins 10 pennies and Player 2 loses 10 (see Figure III.2).

It is clear that Player 1 is now more inclined to play Tails, and that Player 2 wants to play Tails less. But Player 1 can't start playing Tails all the time, because then Player 2 would exploit by always playing Heads. He needs to play Heads occasionally to prevent this from happening. How often should he play Heads?

Recall from the original version that the optimal mix is the one which makes your opponent indifferent towards his/her various pure strategies. In other words, if you play your optimal mix then your opponent expects the same payoff regardless of what he/she plays. This principle is the key to finding the optimal mix, which we now determine for Player 1 as follows.

We wish to find the probability p with which to play Heads. It should give Player 2 the same expected

payoff for both of his/her strategies. When Player 2 plays Heads, his expectation is:

$$EV_2(H) = p \cdot (-1) + (1 - p) \cdot (1) = 1 - 2p.$$

When Player 2 plays Tails, his expectation is:

$$EV_2(T) = p \cdot (1) + (1 - p) \cdot (-10) = 11p - 10.$$

To say that Player 2 is indifferent between those options is to say that

$$EV_2(H) = EV_2(T),$$

i.e. that

$$1 - 2p = 11p - 10.$$

We can solve this equation for p :

$$\begin{aligned} 1 - 2p = 11p - 10 &\Leftrightarrow -13p = -11 \\ &\Leftrightarrow p = \frac{-11}{-13} = \frac{11}{13} = 0.84615\dots \end{aligned}$$

Thus Player 1 should play Heads about 85% of times, and Tails about 15% of times.

What about Player 2? Denote the optimal mixing proportion for Player 2 by q . Now compute

$$EV_1(H) = 1 \cdot q - 1(1 - q) = 2q - 1, \quad EV_1(T) = -1q + 10(1 - q) = 10 - 11q.$$

Setting these equal and solving for q gives

$$2q - 1 = 10 - 11q \Leftrightarrow 13q = 11 \Leftrightarrow q = \frac{11}{13}.$$

Thus Player 2 should play Heads 85% of times and Tails 15% of times as well. It is no coincidence that these numbers are the same as for Player 1. The reason is that this game is a zero-sum game. In general though, the optimal mixing proportions are not the same for both players.

We can now also compute what both players should expect to win or lose in this game when both play their optimal mix. Player 1 expects to win

$$EV_1(H) = 1 \cdot \frac{11}{13} - 1 \cdot \frac{2}{13} = \frac{9}{13}$$

per game. We only have to compute the expected value of Player 1 choosing Heads, because he is indifferent between Heads and Tails. Since the game is zero-sum, Player 2 expects to lose $\frac{9}{13}$ per round. This number ($\frac{9}{13}$ for Player 1, and $-\frac{9}{13}$ for Player 2) is usually referred to as the *value* of the game.. When a game has a positive value for Player 1 but a negative value for Player 2, that means that the game is stacked in favour of the first player.

One of the counterintuitive aspects of the mixed strategy is that Player 2 plays Tails at all. We learn two lessons from this: first, that you sometimes have to do something seemingly unfavorable in order to protect yourself from being exploited. Second, that it is often impossible to judge a move in isolation; one has to judge the overall strategy.

III.2.3 NASH EQUILIBRIA IN MIXED STRATEGIES

We have seen that there are games which have no Nash Equilibria in pure strategies, such as the matching pennies game. However, Nash proved that if we allow mixed strategies, then every finite game always has at least one Nash Equilibrium.

In order to understand this, let us recall that a strategy profile (i.e. a choice of strategy for each player) is a Nash Equilibrium when no player can improve his/her payoff by unilaterally changing his/her strategy. This definition also applies to mixed strategies: we only have to replace payoff by *expected payoff*. Concretely, this means that no player can benefit by changing the proportions in his/her mix of pure strategies.

Let us try to understand this in the context of a pure coordination game, such as the one from Figure I.15. We already saw that this game has two Nash Equilibria in pure strategies. However, there is a third equilibrium in mixed strategies: if both players mix their pure strategies with equal probability, then neither player can improve by changing unilaterally. However, as soon as one of the players would have a different mix, the other player will want to change his strategy from the 50-50 mix to a 100-0 or a 0-100 mix.

Note however, that this equilibrium where both players mix both pure strategies with equal probability is of a different nature than the two Nash equilibria in pure strategies: the latter two are stable equilibria, in the sense that if one of the players would diverge a little from it, then the other player couldn't take advantage by unilaterally changing his strategy. However, in the mixed strategy equilibrium the tiniest divergence from the 50-50 mix by one of the players induces the other to switch to a pure strategy to take maximal advantage.

For another example, consider again the Chicken game:

		Player 2	
		Straight	Swerve
Player 1	Straight	0 / 0	2 / 10
	Swerve	10 / 2	5 / 5

We already saw that there were two Nash Equilibria in pure strategies, namely (Straight, Swerve) and (Swerve, Straight). However, if you intend to play this game you may wish to consider a randomized strategy as well. Thus, with probability p we're going straight, and with probability $(1 - p)$ we swerve. How should we choose p ? Well, as always we want to make our opponent indifferent between his pure strategies, in order to be unexploitable. In this particular case that means that his EV of going straight is the same as that of swerving. In a formula:

$$EV(\text{P2 straight}) = p \cdot 0 + (1 - p) \cdot 10 = 10 - 10p,$$

while

$$EV(\text{P2 swerve}) = p \cdot 2 + (1 - p) \cdot 5 = 7 - 3p.$$

Setting these equal allows us to solve for p :

$$\begin{aligned} 10 - 10p &= 7 - 3p \Rightarrow -3 = -7p \\ &\Rightarrow p = \frac{3}{7} \end{aligned}$$

Exercise 35. Find the Nash Equilibria in mixed strategies in the Coordination game and in the Battle of the sexes game.

Exercise 36. Try to use the above method for finding Nash Equilibria in mixed strategies in the Prisoner’s Dilemma. What happens? Can you explain what goes wrong?

We give one more illustration, this time in a slightly larger game, namely Rock-Paper-Scissors. For ordinary RPS, we may represent the payoff matrix as in Figure III.3.

		Player 2		
		r	p	s
Player 1	r	0	-1	1
	p	-1	0	-1
	s	1	-1	0

Figure III.3: Rock-Paper-Scissors

It should be intuitively clear that there is only one Nash Equilibrium in this game: both players should play each of their pure strategies with a 1/3 probability. As soon as one of them diverges by playing, say, Rock more frequently, the other player can exploit that by switching to the pure strategy Paper. (See below for some more explanation of what it means to exploit an opponent’s strategy.)

But let’s see if we can find the Nash Equilibrium in a systematic way and demonstrate that there really is only one NE. We thus wish to find the pure strategy frequencies for Player 1 which make Player 2 indifferent between her options. Let’s call the frequency with which Player 1 throws Rock p and the frequency with which he throws Paper q . Then the frequency with which he throws Scissors is $1 - p - q$. Now given these -as of yet unknown- frequencies, Player 2 can calculate the expectation of each of her pure strategies as follows:

$$\begin{aligned}
 \text{EV}_2(\text{Rock}) &= p \cdot (0) + q \cdot (-1) + (1 - p - q)(1) = 1 - p - 2q \\
 \text{EV}_2(\text{Paper}) &= p \cdot (1) + q \cdot (0) + (1 - p - q)(-1) = -1 + 2p + q \\
 \text{EV}_2(\text{Scissors}) &= p \cdot (-1) + q \cdot (1) + (1 - p - q)(0) = -p + q
 \end{aligned}$$

To say that Player 2 is indifferent between each of her pure strategies is to say that these three quantities are equal. Thus, the frequencies p and q must satisfy

$$1 - p - 2q = -1 + 2p + q = -p + q.$$

To solve for p and q , we take the last equation and subtract q from each side, giving

$$-1 + 2p = -p.$$

Thus we find $p = 1/3$, as expected. Now we can substitute this value of p into the other equation $1 - p - 2q = -1 + 2p + q$ to get $1 - 1/3 - 2q = -1 + 2/3 + q$ and solve for q , giving $q = 1/3$. Since the payoff matrix is symmetric, the very same reasoning works to find that the optimal frequencies for Player 2 are $1/3$ of each. This solves the game.

Suppose next that we create a variation on the game, by giving Player 2 one bonus point when she beats Player 1 with Rock, and that we give Player 1 one bonus point in case of a tie. This gives the following game matrix:

		Player 2		
		R	P	S
Player 1	R	0 1	1 -1	-1 1
	P	-1 1	0 1	1 -1
	S	2 -1	-1 1	0 1

Figure III.4: Rock-Paper-Scissors, Variation

What are now the optimal strategies? And which player would you rather be in this variant? We proceed as in the standard version, by denoting by p and q the frequency of Player 1 throwing Rock and Paper, respectively. Then we calculate the EV for Player 2 of each of her pure strategies:

$$\begin{aligned}
EV_2(\text{Rock}) &= p \cdot (0) + q \cdot (-1) + (1 - p - q)(2) = 2 - 2p - 3q \\
EV_2(\text{Paper}) &= p \cdot (1) + q \cdot (0) + (1 - p - q)(-1) = -1 + 2p + q \\
EV_2(\text{Scissors}) &= p \cdot (-1) + q \cdot (1) + (1 - p - q)(0) = -p + q
\end{aligned}$$

Put these equal to get

$$2 - 2p - 3q = -1 + 2p + q = -p + q.$$

From the second equation we find again that $p = 1/3$. Plugging that into the first equation and solving for q gives $q = 5/12$.

We're not done: this game isn't symmetric, so we need to calculate things from the perspective of Player 2 as well. Let's call her frequency of throwing Rock u , and her frequency of throwing Paper v . Next, calculate the Expected Value of Player 1 for each of his possible options:

$$\begin{aligned}
EV_1(\text{Rock}) &= u \cdot (1) + v \cdot (-1) + (1 - u - v)(1) = 1 - 2v \\
EV_1(\text{Paper}) &= u \cdot (1) + v \cdot (1) + (1 - u - v)(-1) = -1 + 2u + 2v \\
EV_1(\text{Scissors}) &= u \cdot (-1) + v \cdot (1) + (1 - u - v)(1) = 1 - 2u
\end{aligned}$$

Setting these equal gives:

$$1 - 2v = -1 + 2u + 2v = 1 - 2u.$$

From $1 - 2v = 1 - 2u$ we see that $u = v$, and substituting that into the first equation gives

$$1 - 2v = -1 + 2v + 2v \Rightarrow 2 = 6v \Rightarrow u = v = 1/3.$$

This means that the solution to the game is the following: Player 1 plays Rock 1/3 of the time, Paper 5/12 of the time and Scissors the remaining 1/4 of the time. Player 2 plays each of her options with frequency 1/3.

The result may surprise you: after all, it seems as if Player 2 should make an extra effort to get her bonus. However, Player 1 has chosen his frequencies in such a way that no matter what Player 2 does, there is no way to profit from the bonus by changing frequencies. And thus the best Player 2 can do is make sure that Player 1 can't exploit her; this means she chooses frequencies in such a way that Player 1 can't benefit by switching to a pure strategy.

However, there is one remaining question: which side should you choose? To this end, we go back to the EV calculation. We know that the EV of all three moves is equal in the Nash Equilibrium, but how much is it? For Player 1, we find that $EV_1(\text{Rock}) = 1 - 2v = 1 - 2 \cdot 1/3 = 1/3$. And for Player 2, it is $EV_2(\text{Rock}) = 2 - 2p - 3q = 2 - 2 \cdot 1/3 - 3 \cdot 5/12 = 1/12$. This means that Player 1 earns, on average, 1/3 per round, and that Player 2 averages 1/12 per round. Thus both benefit from playing this game, but Player 1 benefits most.

Exercise 37. Consider the variation on RPS where everything is as in the original game, except for the following two outcomes: when both players throw Rock, Player 1 earns one point and Player 2 loses 1. And when Player 1 throws Rock but Player 2 throws Paper, Player 1 loses two points and Player 2 wins two points. Write down the payoff matrix, find the Nash Equilibrium and determine which player has the advantage in this variant.



III.3 THE AKQ GAME

The AKQ game is an extremely simplified form of poker. It's so simple that you probably wouldn't want to play it. Still, it contains several worthwhile lessons, which we need to understand before we can hope tackling more realistic scenarios.

Description

The game involves two players, I and II. Both get dealt one card randomly from the three-card deck A,K,Q. There is an amount M in the pot. Player I may check or bet one unit (the pot size is measured in terms of units). If player I checks, there is a showdown and the highest card wins the pot. If player I bets then player II may fold or call. In the first case, player I wins the pot, in the second case there is a showdown and the highest card wins the pot.

Exploitative Play

Whenever one of the players does not play an optimal strategy, he or she can be exploited by the other player. Usually, finding out how different strategies can be exploited or what the best counterstrategy gives us hints about what the optimal strategies will look like.

First, suppose that player I decides to only value bet his aces¹ and check down his kings and queens. How should player II respond? Well, in this case he should never call a bet since a bet means that player I has the ace. Thus player II folds to any bet. In this situation, both players break even; both win the pot 50% of the time, and no other money changes hands.

Can player I do better than breaking even? The problem is that he doesn't get paid off on his value bets. Player II has no reason to ever call when he's guaranteed to lose. Thus player I needs to give player II a reason to call. What could that reason be? Well, if player I would also bet weaker hands then player II should do something to prevent player I from stealing the pot with the weaker hand. In other words: player I will have to start bluffing!

This idea is extremely important: the main reason for bluffing is not to make a big profit by stealing pots; it is to get value out of your winning hands. If you don't bluff, you don't get paid off.

¹In poker, a *value bet* is a bet which you make because you think you very likely have the best hand and might get paid off by a worse hand.

How should player I bluff? There is no point in betting with a king, because player II will always call with an ace and will always fold a queen (you can't catch a bluff with the weakest possible hand). Betting with a king is a sucker play, whether you try to do it for value or as a bluff. Thus player I will have to use his queens to bluff.

So suppose now player I still bets all his aces, but now also bets his queens as a bluff. How should player II respond? If player II doesn't change his old strategy of folding to any bet, he will lose $2/3$ of the pots because he folds, and in addition he will lose those pots where player I has a K and he has a Q. This is not good. He will have to start calling in certain spots. Obviously, calling with an ace is always a good idea, since it always wins. So, how will player II do if he always calls with aces and folds all kings and queens? We could calculate the expectation for both players here (you may want to do that as an exercise) but one thing jumps out: if player I bluffs, then player II will gain one bet when he has the ace and calls, but loses the entire pot when he has a king. Does that mean that folding the king is a bad idea?

Suppose player II calls with aces and kings. Now bluffing the queen is an ineffective strategy for player I, because it never works. How should he adapt? Simple: by not bluffing and only value betting aces. Now he makes a profit, because he gets paid off if player II has a king. In turn, player II should stop calling with kings if player I never bluffs, but once that happens, player I will start bluffing again, and so on.

What we see here is that the players start oscillating between strategies. This means that the optimal strategy cannot be a pure strategy, but must involve a mix of pure strategies. To summarize the main observations so far:

- Player I can guarantee to break even by never betting (or by only betting his aces).
- Player I should always bet aces, since that is a weakly dominant strategy.
- However, if player I only value bets his aces then player II will always fold to a bet and player I never gets paid off.
- Thus, player I needs to bluff occasionally with his queens in order to keep player II from guessing.
- Player II should never call with a queen, since that is a dominated strategy.
- Player I will never bet a king, because it will never get called by a queen and always get called by an A. Thus betting a K is (weakly) dominated by not betting.

III.3.1 SOLVING THE AKQ GAME

From the previous observations we know now what type of strategies we're looking for: player I should always bet his aces and sometimes bluff his queens, while player 2 should always call with aces, and occasionally call with kings. The remaining questions to be answered are therefore: how often should player I bluff his queen hands? And how often should player II call a bet when holding a king?

In what follows we shall use the game-theoretic concept of a Nash Equilibrium in mixed strategies to solve these problems. The key idea is again that the optimal mix of two pure strategies makes the opponent indifferent between his/her pure strategies.

Let α be the percentage of times player I bluffs his queens; alternatively, α is the ratio of bluffs to value bets. Since player I value bets all his aces, he bets $\frac{1}{3} + \frac{1}{3}\alpha$ of his hands in total. This means that when player 1 makes a bet, the probability that it is a bluff is

$$P(\text{bluff}|\text{bet}) = \frac{\text{perc. of bluffing hands}}{\text{perc. of betting hands}} = \frac{\frac{1}{3}\alpha}{\frac{1}{3}\alpha + \frac{1}{3}} = \frac{\alpha}{\alpha + 1}.$$

As a consequence, we also find that the chance that a bet is a value bet is

$$P(\text{value bet}|\text{bet}) = \frac{1}{\alpha + 1}.$$

If α is the optimal bluffing frequency, then it makes player II indifferent between calling and folding. If player II calls a bluff, then he earns 1. If he calls a value bet, he loses 1. This translates as:

$$\begin{aligned} \text{EV}(\text{call}) &= \text{EV}(\text{call, player I has A}) + \text{EV}(\text{call, player I has Q}) \\ &= \frac{1}{\alpha + 1}(-1) + \frac{\alpha}{\alpha + 1}(1) \\ &= \frac{\alpha - 1}{\alpha + 1} \end{aligned}$$

On the other hand, when player II folds, he doesn't lose anything when player I has an A, but when player I has a Q he loses the pot of M chips. Therefore:

$$\begin{aligned} \text{EV}(\text{fold}) &= \text{EV}(\text{fold, player I has A}) + \text{EV}(\text{fold, player I has Q}) \\ &= \frac{1}{\alpha + 1}(0) + \frac{\alpha}{\alpha + 1}(-M) \\ &= -M \cdot \frac{\alpha}{\alpha + 1} \end{aligned}$$

We now set these two equal, and solve for α :

$$\begin{aligned} \text{EV}(\text{call}) = \text{EV}(\text{fold}) &\Leftrightarrow \frac{\alpha - 1}{\alpha + 1} = \frac{-M\alpha}{\alpha + 1} \\ &\Leftrightarrow \alpha - 1 = -M\alpha \\ &\Leftrightarrow \alpha + M\alpha = 1 \\ &\Leftrightarrow \alpha(1 + M) = 1 \\ &\Leftrightarrow \alpha = \frac{1}{1 + M} \end{aligned}$$

Next, call the fraction of times player II calls with his kings β . If this calling frequency is optimal, it must make player I indifferent between bluffing and checking his queens. Note that when player 1 has a queen and bets, we have

$$P(\text{get called by A—bluff}) = \frac{1}{2},$$

$$P(\text{get called by K—bluff}) = \frac{1}{2} \cdot \beta$$

and that the probability that he will win the pot uncontested is

$$P(\text{fold K—bluff}) = \frac{1}{2}(1 - \beta).$$

Now the expected value of checking is always 0. The expectation of bluffing is:

$$\begin{aligned} \text{EV}(\text{bluff}) &= \text{lose against A} + \text{lose against K} + \text{K folds} \\ &= \frac{1}{2}(-1) + \frac{1}{2}\beta(-1) + \frac{1}{2}(1 - \beta)(M) \\ &= \frac{1}{2}(-1 - \beta + M - M\beta) \\ &= \frac{1}{2}(-1 - \beta(1 + M) + M) \end{aligned}$$

Setting this equal to 0 and solving for β gives

$$\beta = 0 \Leftrightarrow -1 - \beta(1 + M) + M = 0 \Leftrightarrow \beta(1 + M) = M - 1 \Leftrightarrow \beta = \frac{M - 1}{1 + M}.$$

This means that we have found the optimal (equilibrium) strategies in the AKQ game:

Solution to AKQ game: Player I bets all his aces, never bets a king, and bluffs a fraction $\alpha = \frac{1}{1+M}$ of his queens. Player II always calls with aces, never calls with queens and calls with a fraction $\beta = \frac{M-1}{M+1}$ of kings. (Assuming $M > 1$.)

Observations

The first main lesson to be drawn from the above analysis is that the frequency with which you bluff depends greatly on the size of the pot. When, for example, the pot is 2 chips, then the optimal bluffing frequency is 1/3. When the pot is larger, say 9 chips, that frequency goes down to 1/10. So we can see that the larger the pot, the less often you should bluff. Similarly, when the pot is large, you should call frequently to avoid being bluffed out of it. After all, you only invest one chip, and losing a large pot would be a bad thing. You only have to be right one in many times for this to be correct.

There is one other interesting point: the smaller the pot, the more frequently player I should bluff. When the pot is 1.1 chips, the bluffing frequency is close to 1/2. However, when the pot is smaller than 1, things change. In that case, there is no point in bluffing anymore, because player 2 can simply respond by switching to the pure strategy of only calling with aces. Then the times player 1 picks up the pot with a bluff don't make up for the fact that his value bets don't get paid off anymore. This is an extension of the concept we saw earlier: when the pot is very small, you don't have to worry about

getting bluffed, so you don't have to pay off your opponent's value bets. By contrast, a big pot gives your opponent leverage, and forces you to pay off his value bets frequently.

We can compute the expectation of this game for both players. For player I, that is done as follows: his expectation of a bluff is 0. (After all, the calling frequency was designed to achieve exactly that.) His expectation of checking a king is 0 as well, since half the time he'll win and half the time he'll lose the pot. Thus we need only find the expectation of a value bet. Since player II pays off $\frac{1}{2}\beta$ of times, this gives

$$\text{EV}(\text{value bet}) = \frac{1}{2}\beta \cdot (1)$$

Since $\beta = \frac{M-1}{M+1}$, we get that the EV of the game for player 1 is

$$\begin{aligned} \text{EV}_1 &= \text{prob. of having an A} \cdot \text{prob. of getting paid} \cdot (1) \\ &= \frac{1}{3} \cdot \frac{1}{2}\beta \\ &= \frac{1}{6} \frac{M-1}{M+1}. \end{aligned}$$

It follows that as the pot grows larger, the expectation for player I tends to $\frac{1}{6}$. We stress that this value is what player I earns as a result of the actions taken. The initial pot is thought of as provided by a third party. This value is usually referred to as the *ex-showdown* value.

III.3.2 POSSIBLE EXTENSIONS

What should be the next step? The AKQ-game is often called a *half-street* game, because the action never gets back to the first player. One possible extension of the game therefore would be to allow player II to make a bet when checked to, in which case player I has to decide whether to call or fold. This variant is called *Kuhn poker*, after Harald Kuhn, an influential game theorist. Many of the same ideas we saw in the AKQ-game play are role here as well. For example, when player I has a king and checks, we now are playing essentially the AKQ-game, but now with the roles reversed: player II either value bets his ace, or bluffs his Q.

A next possible extension would be to allow player II to raise. Or we could drop the restriction that the bet size is always one unit, and make it a no-limit or a pot-limit game. Each of these extensions can be solved manually. The interested reader is referred to the book *The Mathematics of Poker*, by B. Chen and J. Ankenman.

Exercise 38. Suppose that in the original AKQ game, we would allow player II to raise when player I bets. If he raises, then player I can decide to call or fold. Try to identify possible situations where player II could profitably raise. Make sure to think ahead and anticipate how player I will respond to a raise given his possible holdings.

Exercise 39. Suppose that we have the same rules as in the original AKQ-game, but now we use the cards AKQJ. Can you determine the optimal strategies for both players? (Hint: this is not so easy, so attempt only when you have nothing to do for the rest of the evening.)

An extra complication arises when we want to consider games with more than 2 players. One important factor is that of *implicit collusion*. An example of this is the following: suppose in a 3-way

pot, player A has a strong made hand, player B has a strong draw, and player C has a weak made hand. Player A bets, and player B correctly calls. Player C is behind and doesn't have the correct odds to call (although he cannot know this for sure). Normally, player A would want him to call (after all, your profit comes from your opponents mistakes). However, in this situation player A is not the only one who benefits from an incorrect call: player B gets part of these profits when he hits his draw. One can show that player A would rather see player C fold in this case. Implicitly player B and player C make money at the expense of player A, even though player C makes a mistake. For a more precise analysis, see the wikipedia article on Morton's Theorem.

III.3.3 EXPLOITIVE VS. EQUILIBRIUM PLAY

Suppose you're playing a game which you have solved completely. That is, you know the equilibrium strategy (or strategies) from game-theoretic analysis. Your opponent, however, doesn't seem to know much about game theory, and you observe after a couple of rounds that he diverges from equilibrium strategy quite a bit. How do you respond?

At first, you may think that you want to keep playing the equilibrium strategy. After all, this strategy is optimal in a sense, no? Moreover, your strategy is designed to make your opponent indifferent between his options, so does it really matter in what proportion he plays them?

To the first point: a strategy is an equilibrium when it is not exploitable, i.e. when your opponent cannot improve his payoffs by changing his strategy. However, when one player diverges from that strategy, we are no longer in a Nash Equilibrium. And therefore it may be possible for the other player to improve her payoffs by changing her strategy too.

As for the second point: it is true that your opponent is indifferent between the available pure strategies when you are playing the optimal mix. However, once your opponent diverges from optimal play, you are no longer indifferent between your pure strategies. Thus you want to switch to whatever pure strategy gives you the highest payoff.

Exploitive Play: As soon as your opponent diverges from game-theoretically optimal (equilibrium) play, you can take advantage by switching to a pure strategy which yields a higher payoff.

Ideally, you play an equilibrium strategy until you see that your opponent isn't playing optimally, so that you can switch to a maximally exploitive strategy. However, there is an important practical issue: when you switch to a pure strategy, sooner or later your opponent will notice that. Then he will likely make adjustments, and start exploiting you. For example, in a Rock-Paper-Scissors game, when you notice that your opponent plays Rock 42% of the time, your maximally exploitive strategy is to play paper all the time. But when your opponent starts noticing, he will start playing Scissors a lot more. You, in turn, will respond by playing Rock more, and so on.

This means that in practical play, you don't want to make it too obvious that you are exploiting a mistake by your opponent. You want to exploit it, but in a way which doesn't make it clear how. If your opponent plays Rock too much, you'll play a bit more paper, but not so much that your opponent

will pick up on it. And once you feel that he does start to realize that you're playing more paper, you should be ready to switch to playing a bit more rock, since he will respond by playing more scissors.

This is very much how poker works as well. Players look for tendencies in their opponent's play, and try to figure out ways of exploiting them. At some point, adjustments are being made, and the story repeats itself. Poker is extra complicated, because there are various aspects in a game plan which may be exploited. For example, one player may try to exploit that the other raises too much before the flop, while the second tries to exploit the fact that the first folds too much to big bets on later streets.



III.4 STRATEGIC MOVES

In Chapter I, we studied sequential and simultaneous games, and optimal strategies for them. Earlier in this chapter, we added to that picture by allowing for mixed strategies. We now add a new dimension to game theory by considering actions which are not part of the game itself, but which change the game in one way or another. Such actions are called *strategic moves*. They include actions which are very common, and even essential, to many strategic situations: threats, promises, commitments. In this section we introduce such moves, and learn how they can help us better understand the dynamics of a variety of everyday conflicts.

III.4.1 BINDING ONESELF

It is often believed that having several strategic options is better than having only one. Sometimes this is true, but there are many situations where the mere fact that we have a certain option available hurts us, and where flexibility is our greatest enemy. Consider someone who wants to lose weight. Every day, he faces many possible options concerning his nutrition: buy a salad or buy candy; drink juice or drink beer, and so on. If the unhealthy options were not available, then dieting would be a lot more successful. Along similar lines, some people have themselves banned from certain websites or from casinos, in order to guarantee that they cannot access them and spend too much time or money there. Thus the option which is most tempting right now is eliminated, and what is left is the option that is best in the long run.

In these examples, you may argue, the problem is a lack of will power and not a lack of strategic insight. While one can argue about that, the idea should be clear: sometimes you're in a better position when you eliminate certain options. Another example makes the same point: suppose you're buying a house and negotiating the price. You really want the house, and it would be a big disappointment if you couldn't get it. You've made up your mind that \$410,000 is a reasonable price, and that you don't want to go over that price (even though you may be able to). You've just offered \$400,000, and the seller counters with \$430,000. Your turn: you offer \$410,000, and tell the seller that that's your final offer. Take it or leave it! Now if things work out, the seller realizes that you don't want to pay more and agrees to the price. But what if the seller isn't fully convinced and simply asks for \$420,000? Now your tactics have backfired: if you make a new offer (\$413,000, say) then you've lost credibility. The problem here is that the seller is not convinced that you won't pay more. If it was clear that you

couldn't possibly pay more than that, your position would be much stronger. Not being able to accept deals which are less favorable to you is a great strategic asset in negotiations!

This leads to the question of how we can bind ourselves in such a way that we won't be able to back out of a commitment. If we can successfully do that and convince our opponent that we are committed to a certain course of action, then he will have to adapt to that reality, and we've managed to change the game to our advantage. This is the idea expressed in Thomas Schelling's work on strategic moves:

“The power to bind an adversary is the power to bind oneself. ”

III.4.2 COMMITMENTS, THREATS, PROMISES

One of the simplest strategic moves is a commitment. When you make a commitment (just as in the colloquial usage of the word) you state that you will, at some future point in the game, take a certain course of action. Of course, this is only informative when that course of action may not be the best one at that point. In the negotiation example, you committed yourself to rejecting any offer higher than your final offer. As became clear there, the problem with this is that once push comes to shove, you may be tempted to back out of your commitment. Another problem is that your opponent must be convinced that you're committed, otherwise he will simply ignore your statement. In other words, a commitment has to be *credible* for it to have the desired effect.

Threats and promises are similar to commitments. A *threat* is a statement of the form:

Threat: “If you make (or don't make) move X, then I will respond with move Y (which is bad for you and is not good for me either).”

A *promise* has the form:

Promise: “If you make (or don't make) move X, then I will respond with move Y (which is good for you but is not for me).”

Thus when you make a threat, you are trying to influence your opponent's actions by indicating that if he doesn't do what you want (or does what you don't want), then you will punish him for that by choosing a response which hurts him. And when you make a promise, you are trying to influence his actions by indicating that when he does what you want him to do (or doesn't do what you don't want) then you will reward him by choosing a move which is good for him.


There are three essential points to make here: First, threats and promises are *conditional*; they are only triggered by a certain move of your opponent. By contrast, a commitment is unconditional in that it is not dependent on what your opponent does.

Second, the course of action which is threatened (or promised) must have a cost. Carrying out the threat (or fulfilling the promise) is not the best move at that stage of the game. For example, if a parent threatens to punish a misbehaving child and the child keeps misbehaving, then the parent faces a choice: follow up on the threat and punish the child, or let it slide. For most parents, punishing a child is not what they enjoy doing. Thus they would prefer not to punish the child. Similarly, if they promised the child a toy if she behaves well, then following up on that promise carries the price of having to buy the toy.

If, punishing the child is preferable to the parents anyway, then we don't speak of a threat anymore: it is simply a *warning*. Similarly, if the parents get most enjoyment out of giving toys to their child, then it is not a promise, but an *assurance*.

Third, the distinction between threats and promises may be blurry. If I'm blackmailing you and tell you that I'm going to destroy the incriminating pictures if you pay me, then that can be considered either a promise (because destroying is good for you) or a threat (because the alternative is bad for you). However, there is one key feature of threats which sets them apart: when successful, a threat costs nothing. A promise on the other hand, will have to be fulfilled and paid for when it works the way you intended it.

Let us consider a game which we looked at in the first chapter: the monopoly game. In one of the variations, the game tree was as follows:



The standard analysis gives a rollback equilibrium where the first player enters and the second player shares the market. However, the second player would much prefer the first not to enter, and thus he may consider making a threat: "If you enter the market, I will start a price war, even though that hurts my own profit margin". If player 1 believes the threat, then he will choose to stay out. However, he might think that player 2 is merely bluffing. After all, if he enters, it is in the best interest of player 2 to share the market. You see, whether a threat is effective or not depends entirely on whether the other player believes you'll follow up on it.

Another aspect which should not be ignored is clarity. If you want a threat or promise to be effective, then you better make sure that your opponent gets the message loud and clear. In some cases, that can be problematic. If communication lines are compromised or when both parties are not on speaking terms, then getting the message across can be difficult. In other cases there are other obstructions. For example, antitrust laws forbid threatening competitors with price wars.

There is another classification of conditional strategic moves, which is orthogonal to the one discussed above. Some moves are intended to have a *deterrent* effect, while others are intended to have a *compellent* effect. A threat can be deterrent or compellent. A deterrent threat aims at preventing a certain action

by your opponent. For example, a sign saying “Trespassers will be shot” is supposed to deter people from entering by threatening death. By contrast, a compellent threat is intended to induce a certain action. For example, if a teacher tells his students that he will give them an F in the course when they don’t hand in their assignments in time, then he is making a threat which is supposed to compel them to do what he wants.

Similarly, promises can be classified as compellent or deterrent. An example of a compellent promise is when a politician tells the electorate that he will lower taxes when elected. And an example of a deterrent promise is when someone who is being robbed tells the robber that he will get all money and valuables, as long as he doesn’t hurt him.

Again, sometimes it depends on the formulation of the situation and one’s perspective on it whether a move is compellent or deterrent. Nevertheless, it is useful to understand what the purported function of a threat or promise is.

III.4.3 CREDIBILITY

You can make threats and promises all day long, but what matters is whether your opponents take them seriously. In other words, they are only effective insofar as they are credible. Thus one is well-advised to think about this before making a threat.

Suppose Mr. F and Kimberly are playing a poker game, and Mr. F, who is holding a weak hand, wants to discourage Kimberly from betting. (He’s hoping for a cheap showdown, but isn’t comfortable to invest any more money in the hand.) In order to do so, he looks menacingly at his opponent as she contemplates a bet, and he grabs his chips in order to convince her that he is determined to call. Is this threat credible? Unless Kimberly really believes that Mr. F is going to call, it isn’t. On the contrary, she will often think that Mr. F would not discourage her if he wants her to bet, or when he plans on calling anyway. Kimberly is well-advised to bet, since there is a good chance Mr. F will fold. (Of course, if Mr. F is sophisticated he may feign this behavior in order to induce a bluff. Kimberly in turn has to ask herself on what level Mr. F is functioning and whether he is capable of such a move.)

How do we make our threats and promises credible? There are several methods.

Changing Payoffs

Recall that one of the main problems with following up on threats and keeping promises is that they are not the best move available; thus there is always the temptation to back out and make the better move instead. But what if we were able to change the payoffs so that the threatened course of action became the one with the highest payoff? In other words, what if we could change the threat into a warning?

One way to achieve this is by using *contracts*. If I make a promise, and you’re afraid I’m going to back out of it, we set up a contract in which I declare to pay you a large sum of money should I break my promise. If this contract is enforceable, then I’ve managed to change my payoffs: all of a sudden not following up on my promise is a very bad option. It is important that the contract cannot be breached (or that this can only be done at high cost), and also that it is not in the interest of the parties involved that the contract can be renegotiated.

Another way of changing payoffs is by getting a third party involved. For example, if I promise to

give you a certain amount of money when you do something I want, you may worry about me not paying you. If I give the money in advance to a neutral third party, then I cannot back out of my promise any longer.

Recall the ultimatum game from Chapter I. The second player is at the mercy of the first, and has to accept whatever player I leaves for him. Would a threat work here? What if player II tells player I that he will reject any offer below 50-50? The problem of course is that the threat is hollow. Player I knows full well that player II is not going to follow up on it. How should player II make the threat credible? By giving money to a third party, and telling that third party that if he accepts any offer below 50-50, then that third party can keep the money. If the amount is equal to or more than half of the stakes of the game, then it becomes bad for player II to accept an unfair deal. Player I recognizes this, and is now forced to make a fair division.

In some cases the game at hand will be repeated, may be part of a bigger game, or there may be other games going on at the same time. In this case, threats and promises may be made more credible by putting one's reputation on the line. If a negotiator breaks a promise in one of his negotiations, then he damages his reputation, which will hurt him in his next negotiation. Similarly, a politician may deceive the public by making false promises on the eve of the election, but next time around the public will (hopefully) punish him for that. A parent who does not follow up on a threat of punishing a misbehaving child is going to experience increasingly larger discipline problems, since the child knows that she can safely ignore these threats.

In all of these examples, the price of a damaged reputation changes the payoffs. In some cases, it is clear that this change reverses the order of the preferences, so that threats or promises become credible.

Irreversibility

Changing the payoffs leaves all your strategic options open, but changes your preferences among them. More radical is to make certain moves impossible. This technique is called *burning your bridges*, or *burning your ships*. The latter refers to the move made by conquistador Hernán Cortés, who conquered the Aztec empire. After landing on the coast of what is now Mexico, he ordered his men to burn the ships²; that sent a clear signal both to his own men as well as to the opposing forces that they were there to stay, and that they would be forced to fight to the death.

A similar thing happens when a player pushes all-in in a poker game. Once he does that, it is overly clear that he will go to showdown and cannot be made to fold any longer. This puts the pressure on the opponent. Of course, whether this is a good move or not depends on the situation, but at the very least it is a clear case of credible commitment.

If your opponent is threatening to burn his own bridges you may want to prevent that. This is well-known among military strategists, who will recommend that one should always leave the enemy a way to escape³. Similarly, in a debate or a negotiation process one may want to leave the opponent a way of making concessions without losing face.

As we saw earlier, communication plays an important role: if you don't get a threat or a promise

²Historical accounts differ in whether he burnt some of them or all of them.

³This dictum already appears in Sun Tzu's *Art of War*.

across clearly, then it is unlikely to work. However, there is more to the story than just that. For example, consider the Battle of the Sexes game. You could try to make the other person believe that you are committed to going to your favorite restaurant, for example by calling her and saying so. However, nothing prevents your partner from trying exactly the same thing on you. Then you're back to the same old stalemate. Ideally, you want to be the person dissing out the threats and promises, while being immune to those of the opposing party. How could you achieve such immunity? Well, by cutting off communication. If you can't reach me, you can't threaten me. You simply switch off your cell phone. Of course, for this to work, it is essential that the other party knows that you can't be reached.

In Stanley Kubrick's movie *Dr. Strangelove* the Soviets have built a *Doomsday Device*. In this case, the device is a network of buried nuclear bombs, which are programmed to detonate as soon as there is an attack on Soviet soil. It is programmed in such a way, that there is no possibility for humans to switch it off: if tampered with, it will denonate automatically.

The doomsday device is a credible deterrent: given the fact that nuclear catastrophe is inevitable once one tries to attack, any attempt to do so would be suicide. The credibility stems from the fact that human intervention is no longer possible. No sane person would retaliate for a minor attack by destroying the entire planet, and if, say the Soviet premier was to operate the machine one could deduce that there was a very high chance that it would never be used. But by leaving the outcome beyond control of human beings, the threat is completely credible.

Even though credibility is not one of them, there are some problems with doomsday devices. Again, it is imperative that everyone is fully aware of their existence. In the movie, the Soviets activate the device on Friday, planning to announce it to the world on Monday. But during the weekend, an attack takes place. The other problem is that the device may be triggered because of an event which should not be regarded as a hostile attack. Accidents happen, and these should not set off the device.

If you find destroying the planet a bit drastic, consider this example. Hostile takeovers are a common phenomenon in the business world. How can a company protect itself against such corporate raiders? One method is the so-called *poison pill*; this is a legal commitment (usually a contract with the shareholders) that guarantees that, should a certain percentage of shares end up in the hands of one individual shareholder, the company will pay out a large sum of money to all of the share holders. In effect, this would cost so much that it would all but ruin the company, thereby making it unattractive to attempt a takeover. Even though this construction has proved successful in some cases, there are various ways around it if it is not set up in a truly irrevokable manner. As with all contracts, there is always the possibility of renegotiating them, or simply waiting until they expire.

Salami Tactics

Salami is an Italian dry-cured sausage, which is best enjoyed thinly sliced. Eating an entire salami without slicing it would be revolting, but each individual slice is easily consumed.

The symbolism here is that sometimes winning the game in one big step may be impossible, but if you can move in many small, easy steps you may get there in the end. For example, consider trade negotiations. Instead of trying to agree on all facets of international trade at once, one often observes two countries negotiate specific issues one at a time. For example, the U.S. and South Korea may negotiate the import of U.S. beef, or the import of Korean cars. Splitting the problem into small, manageable subproblems tends to have a better chance of success.

Salami tactics can also be used to neutralize a large threat. For example, if a parent tells her son that he will be grounded for two weeks if he makes a mess of his room, then even though this threat may be credible, it can be frustrated. The problem here is that a minor transgression does not warrant invoking the punishment. Thus the son can make little bits of mess without fearing punishment, and in the end there is a total mess while at no particular point in time the parent felt justified in following up on the threat.

Dictators are also well-aware of this: no superpower is likely to start a war over a minor incident. Thus Kim Jong-Il could sink a South Korean ship from time to time, or launch a missile, knowing that this would not result in an attack by the enemy.

There are several lessons: first, threats should be proportional. If you only have a big gun, then it is clear that you're unlikely to use it, and hence you're inviting the opposition to use salami tactics. Second, if you make a threat, there should be very clear and unambiguous criteria as to what counts as a breach. If your opponent can operate within a grey area the pressure is back on you. Third, salami tactics are not always negative. In some cases moving in small steps can help gain mutual trust and make the overall goal more attainable.

III.4.4 IRRATIONALITY

One final strategic idea is to create and exploit an aura of unpredictability and irrationality. One of the first examples of this idea at work appears in the Iliad: the Greeks are putting together an army to attack Troy, and the Greek leaders send messengers to Ithaca in order to recruit Ulysses, one of the great Greek heroes. Ulysses, however, has seen enough war and wishes to stay home. In order to escape his fate, he feigns insanity: once the messengers arrive at his house, they see him plow the beach instead of the farmland. This charade almost works; unfortunately for Ulysses though, one of the messengers devises a clever *screening test*: in order to find out whether Ulysses is really insane, he takes his little son and places him in front of the plow. At this point, Ulysses has to give up his act, as continuing it would mean killing his own son.

One advantage of appearing to be irrational is that it becomes very difficult for your opponents to figure out what your payoffs are, thereby making it harder to come up with a good strategy. A rational leader will try to protect his citizens from harm, and will not endanger himself and his country by aggressive acts towards other countries. But many dictators make a point of creating some doubts in the minds of the rest of the world as to whether they are fully rational or not.

On a smaller scale, pretending not to understand, not to hear, or not to care about certain things can be effective as well, provided others believe you. You can argue with the parking officer that you didn't see the sign. When a schoolyard bully picks on certain kids, it's because he knows they won't fight back. Once a kid convinces him that he's willing to put up a fight even if he loses and gets beaten up, the bully won't bother him. And at the poker table, you often see a maniac who raises every pot. Sometimes, this really is someone who just enjoys messing around. But other times, this is a calculated strategy: by creating the impression that he truly does not care about losing, other players know that it will cost them to see a river, and this often stops them from playing back.

The term *brinkmanship* refers to the art of gradually increasing risk in a situation where you cannot directly achieve your goals. A classic example is seen in many movies, where a cop deliberately creates a dangerous situation in order to force a confession from a criminal. The key feature is that the situation

is risky for both parties, and that the increase in risk is gradual. Thus in effect, this is a game of chicken: the person employing the strategy hopes that the opponent has a lower tolerance for risk, and hence will cave in at some point. Needless to say, this strategy can go very wrong, either because of bad luck or because you've underestimated your opponent. Use at your own risk!



III.5 THE BOTTOM LINE

We have spent a lot of time studying how to recognize which bet bets have positive expectation and which don't, so that we can identify profitable situations and gain a strategic advantage. But even if you find yourself in the situation where you have a nice edge you have some thinking to do. To make the point concrete, imagine that your bankroll is \$1,000, and that you're trying to increase this as much as possible. I'm offering you a favorable bet: we flip a coin, and when it comes Heads I pay you 1.2 times whatever you wagered, while you lose your wager when it comes Tails. Clearly, this bet has positive expectation: you expect to make $0.1 \cdot X$, where X is the amount wagered. To make things even better, I'll let you play this game as many times as you like.

You're in a good spot here and you definitely want to play this game over and over again. However, the question is this: what fraction of your available \$1,000 do you want to risk on each bet? Clearly, risking it all on one bet is foolish: half of the time you go broke and you lose out on the possibility of capitalizing on the opportunity. On the other hand, you don't want to be too timid either: if you only bet \$1 at a time, it will take ages before you start showing a nice profit. Somewhere in between these extremes, there is an optimal amount to bet, which balances the two aspects of avoiding ruin and maximizing bankroll growth.

It is this problem we address in this final section, where we study the Kelly Criterion and other concepts relating to long-term bankroll management. These ideas are of importance not only to poker players but also to investors.

III.5.1 RISK OF RUIN

We consider a situation where a gambler places a series of bets until he either achieves a certain monetary goal or goes broke. We shall refer to the parties involved as "Gambler" and "House". The Gambler starts with an initial bankroll of 1,000 dollars. (All results will easily be adapted to other starting amounts.)

Scenario 1

In this scenario the gambler plays an even money bet against the house with a 50% chance of winning each individual bet, and aims to double his bankroll. His strategy is to bet a fixed amount $\$1,000/N$ each time. For example, if he decides to bet \$50 each time, we have $N = 20$.

By symmetry, it is clear that the gambler's chances of succeeding are exactly 50%. How long do we expect the gambling session to take? This depends on the bet size, but one thing is certain: it will end sooner or later. The reason, informally speaking, is that if you repeat a random experiment such as a coinflip often enough, then eventually one will see a sequence of N identical outcomes.

We can be more specific though: if $N = 1$, things will be over in one round. If $N = 2$, then things are more complicated. If the gambler wins the first two bets (probability $1/4$) or the house wins the first two (probability $1/4$) then the game is over. But in the remaining two cases we're back to square 1. In the latter case (which occurs with probability $1/2$) we again have that in two out of four cases the game ends within two move, and that in the remaining two cases the game starts over again. Thus with probability $(1/2)^k$, the game is over in $2k$ rounds. For larger values of N similar reasoning gives that with probability $(1/N)^k$, the game is over in Nk rounds.

Scenario 2

We consider the same betting situation, but now the Gambler has the goal of reaching a target bankroll of \$4,000. Again, the betsize is a fixed $1,000/N$. Now, the gambler's chances of success are smaller. There is a 50% chance that he manages to get up to \$2,000, as seen in the previous scenario. Now he is half way there: with probability $1/2$ has will manage to double up again. Thus the chances of success are $1/4$ in total.

More generally, when playing a fair even money bet starting with a bankroll of K betting units and playing until a target of $L > K$ is reached (or until bankruptcy) the chances of success are $p(\text{success}) = K/L$.

Sometimes this is reformulated slightly: when the gambler starts with K units, and the casino has K' units, then the chances of the gambler breaking the casino are

$$p(\text{breaking the house}) = \frac{K}{K + K'}$$

while the chances of the gambler going broke are

$$p(\text{ruin}) = \frac{K'}{K + K'}$$

Scenario 3

Again the same rules as before, but now our gambler plays on forever (or until he is broke). This seems silly, but in practice few people know how to quit when ahead! What are the chances that the gambler goes broke eventually? We can imagine that the gambler plays against an infinitely wealthy house. Then by the previous formulas

$$p(\text{ruin}) = 1.$$

Scenario 4

Let us now assume that the gambler has found a favourable game, where his odds of winning each individual bet are p . Again, we suppose that the bet is even money. Let the gambler start with K

Gambler	House	Chance of success:
5	5	0.732
5	25	0.635
10	10	0.881
10	50	0.866
30	30	0.997
30	150	0.997
100	100	≈ 1
100	1000	≈ 1

Figure III.5: Risk of Ruin, $p = 0.55$

betting units, and let the house have K' .

For convenience, write $r = \frac{1-p}{p}$ for the loss-win ratio on the individual bet. Then the chance that the gambler goes broke is

$$p(\text{ruin}) = \frac{r^K - r^{K+K'}}{1 - r^{K+K'}},$$

and the chance of the gambler breaking the house is

$$p(\text{breaking the house}) = \frac{1 - r^K}{1 - r^{K+K'}}.$$

For example, if the gambler has a $2/3$ chance of winning each individual bet, then $r = \frac{1/3}{2/3} = \frac{1}{2}$. If he starts with 3 betting units ($K = 3$) and attempts to double up then the chance of ruin is

$$p = \frac{(1/2)^3 - (1/2)^6}{1 - (1/2)^6} = \frac{1/8 - 1/64}{1 - 1/64} = 1/9$$

and hence the chance of doubling up is $8/9$.

When the chances of winning are $p = 0.55$, Figure III.5 gives the chances of breaking the house for various starting bankrolls:

One obvious conclusion is that the larger your starting bankroll, the smaller the chance of going broke because of bad luck. This of course has significant implications for poker players, sports bettors, and investors.

Exercise 40. What are the gambler’s chances of success in roulette when she is always betting \$10 on red starting with \$50 and having the goal of winning \$100?

Exercise 41. Suppose the gambler has a 60% chance of winning each bet and gets paid even money. The gambler starts with \$1000 and bets \$10 each round. However, the gambler spends all profits above \$1,500. Explain why the gambler goes broke sooner or later.

III.5.2 THE KELLY CRITERION

As we saw in the previous section, even when one has an edge there may still be a chance of getting unlucky and going broke. This means that one has to be careful about investing too large a part of one’s

bankroll in a single betting opportunity. The *Kelly Criterion*, developed by John Kelly, tells us how to maximize bankroll growth without running the risk of going broke. We first explain the formula, and then we discuss some important restrictions.

Kelly Criterion: Given a bet with probability p of winning, and which lays B to 1 odds, the fraction F of the available bankroll which should be wagered is

$$F = \frac{Bp + p - 1}{B}$$

For example, consider an even money bet ($B = 1$) which has a 70% chance of being won ($p = 0.7$). Then the formula tells us to invest

$$F = \frac{1 \cdot (0.7) + (0.7) - 1}{1} = 0.4 = 40\%$$

of our bankroll.

When $p = 0.4$, $B = 2$, then we get

$$F = \frac{2 \cdot (0.4) + (0.4) - 1}{2} = 0.1 = 10\%$$

so we should invest 10% of our bankroll.

When $p = 0.3$, $B = 2$, we find

$$F = \frac{2 \cdot (0.3) + (0.3) - 1}{2} = -0.1$$

which is a negative number. This means that we shouldn't invest at all. This makes sense, because this bet has negative expectation.

We can deduce that $F = 0$ precisely when $B = \frac{1-p}{p}$. In this case, the gambler does not have an edge (has neutral expectation), so the formula recommends to refrain from betting.

Note that in the case of an even money bet ($B = 1$) we can simplify the formula to

$$F = 2p - 1.$$

There are several important warnings that come with this formula. The first is that it only applies to situations where the same bet is repeated a large number of times. For one-shot events, the formula is meaningless, since it is designed to maximize long-term bankroll growth.

Next, the formula presupposes that the gambler can wager any fraction of the bankroll. This is not always realistic: sometimes we cannot buy 7.4 shares, or play in a 4.3-8.6 poker game. In these cases, it is usually recommended to err on the side of caution, i.e. to invest slightly less than the Kelly bet.

Also, it is important to keep in mind that the actual size of the bet changes after each round of betting. When we lose a bet, our bankroll has diminished, and so our next bet will be smaller than the previous one. Similarly, when we win, our bankroll is larger, and so our bet will increase accordingly.

In theory, one cannot go broke when using the Criterion, because one will never bet the entire bankroll. However, in practice things may be different. After a long run of bad luck one may be down to one's last dollar, and it may be impossible to bet the required fraction of that.

Kelly betting has been used successfully in many situations, including investing. In practice, many investors will recommend using an approach which is more conservative than the Kelly Criterion, for example, by betting half of what the Criterion prescribes. This is called *half-Kelly*. This reduces risk and variance, but also reduces the long-term growth.

It is not easy to apply the Kelly Criterion directly to poker, because there are a lot of variables to be taken into account. Ideally, one would like a formula which tells exactly how many buy-ins are needed for a given game. Many such formulas have been proposed, but it is important to keep in mind that the number we are after depends on

- Winrate
- Variance
- Possibility to drop down in stakes if needed
- Tolerance of risk

Of course, if you don't have an edge in a game, you shouldn't play. The larger your winrate, the greater your edge, and the more you would like to invest. However, games differ greatly when it comes to variance (which is a measure of the swings due to short-term luck). In an aggressive short-handed Heads-Up NL game, or a PLO game, swings of 10-20 buy-ins are quite common for winning players. A bankroll should comfortably sustain such swings.

It also matters whether your bankroll is replacable or not. If you are a professional player who doesn't have the possibility of redepositing, then you need to be much more conservative with your bankroll than when you have a day job and merely play as a hobby. Also, some people are more comfortable taking risks than others, and when developing your bankroll management strategy you should take this into account by asking which part of your bankroll you're comfortable with losing. A last important factor is whether you add all of your winnings to your bankroll, or whether you occasionally take out some of these winnings. When you always withdraw to the point where you have 20 buy-ins left, it is a mathematical certainty that at some point you will hit a bad run of luck which will force you to drop down in stakes or redeposit.

While we do not attempt here to evaluate some of the proposed formulas, we do give a few rules of thumb. For lower stakes NL Hold'em games, a player with a decent edge should be able to comfortably play with 30 buy-ins. This number goes up when the games get more aggressive, when the edge becomes smaller, or when the variance in the game goes up for other reasons. For someone who plays Heads-Up NL and plays an aggressive style, 50 buy-ins is usually considered the bare minimum, and many advocate 100 buy-ins. For PLO these numbers are even higher. While in a lower stakes game a good player can play with 30-50 buy-ins, the swingy nature of the aggressive higher stakes games make 100 buy-ins a must.

III.5.3 ON FREE LUNCHES

We conclude with a little puzzle. There is a lot of evidence for the claim that there is no way to predict the fluctuations in the stock market⁴. This seems to surprise many who are used to reading about investment strategies, top hedge fund managers, highly successful funds, et cetera. But if you have millions of people play a game of chance, then it is inevitable that a couple of them will have incredibly good results. Those are the individuals that we look back on and think of as being more clever than the rest. You could compare the situation to giving everyone in Canada one dollar, pairing up all these people and let them flip a coin for their dollar. The winners continue, until we only have a few millionaires left. Are these millionaires great at coin flipping? No, there's nothing special about them; it was inevitable from the beginning that we would end up with a few winners.

So suppose this is correct, and that the fluctuations in stock prices are just random walks. Let's pick a stock. Each day, the price of the stock goes up or down by a certain percentage, but there is no way we can predict in advance what direction it will go. If the stock is equally likely to go up or down then there is no point in buying it, is there? Wrong! Even if you are convinced that the short-term fluctuations are completely random and that over the long term the stock is expected to remain constant, you can still make money off it.

Here's how. We suppose for simplicity that every day the price of the stock either doubles or halves. This is a bit extreme, but the calculations below can easily be adapted to whatever figure you think is realistic. So on day 1 the price might be \$64, then on day 2 it goes down to \$32, on day 3 it goes down again to \$16, then up to \$32 again, and so on. Clearly there is no point in investing your entire bankroll on this stock: after all, we expect to make no profit in the long run. But here's the trick: every day, you take half of whatever your total worth is at the moment, and invest it in the stock. The other half you keep under your mattress. This is repeated every day, so that you're always recalibrating to make sure that your money is evenly divided between the stock and the mattress.

In the following table we see what happens to our bankroll when we follow this scheme for a couple of days:

	Stock Price	Total Bankroll
Day 1	\$64	\$1,000
Day 2	\$32	\$750
Day 3	\$16	\$562.50
Day 4	\$32	\$842.75
Day 5	\$16	\$632.81
Day 6	\$32	\$949.22
Day 7	\$64	\$1,423.83
Day 8	\$32	\$1067.87
Day 9	\$64	\$1,601.81
Day 10	\$128	\$2,402.71
Day 11	\$64	\$1,802.03
Day 12	\$32	\$1,351.52

We begin with \$1,000 in our bankroll. We split it in two halves, and buy \$500 worth of stock. At the

⁴Of course, having insider knowledge is an undeniable edge, but using it is illegal, so we don't consider it here.

end of the day, the stock has lost half of its value. Now our net worth is \$750. We split it in half, and invest \$375 in stock. Again it loses half its value, making our net worth \$562.50. And so on; as we continue recalibrating our investment at the end of every day, we make money out of thin air: after 10 days the stock price is back to where it started, but we made about \$800 on our initial investment of \$1,000. Even on the last day, where the stock price loses half its value, we're still up a solid 35%. Compare that to the 50% you would have lost had you invested all your money in this stock!

Let's be a bit more precise about how much money this is going to make us. Each time the stock goes up, our bankroll increases by a factor $3/2$. Indeed, we divide it in half, one half doubles and the other half stays the same, giving three halves. When the stock goes down, our bankroll decreases by a factor of $3/4$. Suppose now that the price goes down first, and then up the next day. Then our bankroll first gets multiplied by $3/4$, and then by $3/2$. The result is that our bankroll grew by a factor $3/4 \cdot 3/2 = 9/8$. Since each time we're simply multiplying our bankroll with a certain factor, the order of events is irrelevant. At the end, all that matters is how many times the stock went up and how many times it went down. If it went up k times and down n times, then our initial bankroll has been multiplied by a factor $(3/2)^k \cdot (3/4)^n$. In the special case where $k = n$ (then the stock price is back to its original value) we find that our bankroll has increased by a factor $(9/8)^k$. This is good news: our bankroll is growing exponentially, and we've effectively found ourselves a bank which pays 12.5% interest every other day. Even if at the end of the year the stock price were to tumble to a few pennies we're still going to be astronomically rich. For example, if the stock were to drop to a price of $\$ \frac{1}{512}$, or one-fifth of a cent, then our bankroll would still be an agreeable \$11.9 billion dollar.

What's the catch? Well, it's not in the math, which is pretty straightforward. The assumption that stock prices follow a random walk is also not implausible. Ideally you want not just a stock whose price behaves randomly, but whose price always increases or decreases by a large factor. It's possible to exploit smaller changes, but it takes much longer. But the main problem is your financial institution, which charges you a transaction fee each day. If it were not for these fees, a volatile but randomly behaving stock price would be all you'd need.

There's one interesting thing to notice here: this scheme works *because* the stock price behaves randomly. As soon as the price follows a certain trend, then using this system is a bad idea. Clearly if you have reason to believe that the price will go down then you don't want to invest a penny, let alone half your bankroll. And if you know for sure that the price will go up, you want to make the most of that and invest more than half.

Exercise 42. Suppose we have a stock whose price increases or decreases randomly each day by a fixed amount. For example, the initial price could be \$100, and each day it goes up or down \$10. Will the same scheme work? Would it be more or less profitable than when the price increases or decreases by 10% each day?

Exercise 43. Explain the paradoxical nature of the above scheme. How is it possible that at the end of the year the stock price may be decimated but we've earned millions of dollars?

RESOURCES AND SUGGESTED READING

Game theory is very popular and widespread, and hence textbooks, online courses, research papers and websites dedicated to the subject are in abundance. Similarly, there is a wealth of information about probability theory, poker and related matters. Below you find a selection of useful books, sites and papers, together with a brief description of the contents.



GAME THEORY

- www.gametheory.net is a site dedicated to all facets of game theory. Here you will find lots of book reviews, interactive tools for solving games, and a list of movies and tv shows featuring game theory.
- Avinash Dixit and Barry Nalebuff, *The Art of Strategy*, is a highly readable and entertaining introduction to game theory and its applications. This text does not require much mathematical background.
- William Poundstone, *Prisoner's Dilemma*, Anchor Books NY, 1992. This highly recommended book not only details the origins of the famous dilemma, but also contains a lot of biographical information about John von Neumann and the social and political background against which the development of game theory took place.
- Thomas Schelling, *The Strategy of Conflict*, Harvard University Press, 1980 (2nd ed.). This classic work on strategic moves and their applications to nuclear deterrence and more is not a light read, but contains a wealth of ideas.
- John Williams, *The Compleat Strategyst* (freely available at www.rand.org). Older but well-written introduction to game theory, with an emphasis on various 2-player simultaneous games.
- Colin Camerer, *Behavioral Game Theory*, Princeton University Press, 2003. This book describes a large number of experiments in game theory, such as the various experiments with the ultimatum game.

- Robert Axelrod, *The Evolution of Cooperation*. Informative text which describes and analyzes the results of the prisoner's dilemma tournaments, and from there treats the connections between game theory and evolution.
- Yale Courses. This is a video recording of a complete course on game theory taught by Ben Polak.
- http://tuvalu.santafe.edu/~wbarthur/Papers/El_Farol.html. This paper discusses a classic game called the El-Farol Bar Problem.
- http://euclid.trentu.ca/math/bz/pirates_gold.pdf. Generalizations and analysis of the pirates game.
- <http://www.amsta.leeds.ac.uk/~pmt6jrp/personal/blotto.html>. The Colonel Blotto game is an easy game to explain, but is as of yet unsolved. This text explains the game and presents some results from computer simulations.
- <http://www.ams.org/samplings/feature-column/fcarc-rationality>. This AMS paper discusses some aspects of the role of rationality in game theory.



PROBABILITY THEORY

- Florence Nightingale David, *Games, Gods and Gambling*, Dover Publications Inc, 1998 (reproduction). Excellent text on the history of gambling and probability.
- Leonard Mlodinow, *The Drunkard's Walk*, Vintage Books, 2008. Highly entertaining informal introduction to probability and statistics from a historical perspective.
- www.random.org. Experience randomness yourself by having random sequences generated while you wait.
- Wolfram Mathematica is a piece of mathematics software which can help you understand various concepts, problems and examples. A demonstration version can be downloaded for free at www.wolfram.org, and after installation various instructive examples can be obtained. Particularly useful are the demonstrations about coinflips, gambler's ruin, dice probabilities and the birthday problem.



POKER AND OTHER GAMES

- The Theory of Poker, *David Sklansky*, TwoPlusTwo Publishing, 2005 (4th ed.). In its day, this was the first book to analyse the mathematical aspects of the game, including some game theory applications. Still highly valuable.

-
- The Mathematics of Poker, *Bill Chen and Jerrod Ankenman*, ConJelCo LCC, 2006. Detailed text on the use of probability and game theory in poker. Includes analysis of several toy poker games such as the AKQ game, and much more. Strongly recommended for those who want to take their poker math to the next level.
 - http://en.wikipedia.org/wiki/Poker_probability. This site contains lots of information on poker odds.
 - <http://www.cardplayer.com>. News and information about poker and other games. Includes detailed introductions to poker in almost all of its forms, instructional articles, and odds calculators.
 - <http://twodimes.net/poker/>. This tool allows you to calculate your equity against various hands or hand ranges. Very useful for those looking to improve their understanding of the game.
 - James McManus, *Cowboys Full*, Picador, 2009. Great read about how poker developed from an average parlor game into an immensely popular spectator sport.
 - William Poundstone, *Fortune's Formula*, Hill and Wang, 2005. The story of the Kelly Criterion, and how it was used by various ingenious people to beat the casinos and the stock market. Highly recommended.
 - Edward Thorp, *Beat the Dealer*, Vintage Books NY, 1966 (2nd ed.). Entertaining and informative first-hand account from the father of card counting about his system, his successes and the numerous ways in which the casinos tried to thwart him.
 - Thomas Bass, *The Eudaemonic Pie*, iUniverse.com Inc, 2000 (2nd ed.). The compelling story of a group of bright science students who develop a computer to beat blackjack.



PSYCHOLOGY

- Scott Plous, *The Psychology of Judgment and Decision Making*, McGraw-Hill, 1993. Excellent introduction to psychological aspects of decision making, which discusses various models of decision making, utility theory, heuristics and biases, framing, common fallacies, and much more.
- Daniel Kahneman, Paul Slovic and Amos Tversky, *Judgment under Uncertainty: Heuristics and Biases*, Cambridge University Press, 2008 (first printed in 1982). Collection of research papers on various aspects of heuristics and biases.
- Dan Ariely, *Predictably Irrational*, Harper Collins, 2008. Popular exposition of how and why we tend towards irrational behavior.



MISCELLANEOUS

- Mikal Aasved, *The Psychodynamics and Psychology of Gambling*, Charles Thomas, 2002. Comprehensive overview of psychological theories about gambling and addiction. First volume in a 3-part series.
- Mikal Aasved, *The Biology of Gambling*, Charles Thomas, 2002. Overview of biological approaches to gambling and addiction.
- Mikal Aasved, *The Sociology of Gambling*, Charles Thomas, 2002. Discussion of sociological theories of gambling behavior.
- www.worldrps.com. Everything you wanted to know and more about the fascinating game of Rock, Paper, Scissors.
- <http://www.transience.com.au/pearl3.html>. Play NIM online.
- <http://www.letsmakeadeal.com/index.htm>. Find out more about the Monty Hall Problem.
- http://en.wikipedia.org/wiki/Joseph_Jagger. Article on the man who broke the bank in Monte Carlo.
- http://en.wikipedia.org/wiki/Unlawful_Internet_Gambling_Enforcement_Act. Internet gambling is currently illegal in the US and in several provinces in Canada. This article explains the history and contents of the US legislation.
- <http://www.rand.org/>. Official website of the RAND corporation. Historical information, free downloads of books and articles, information on current affairs and even some job opportunities.
- www.nobelprize.org. Official site of the Nobel Prize organization. You can find several interesting lectures by well-known game theorists and psychologists, such as Thomas Schelling and Daniel Kahneman.
- <http://www.pacm.princeton.edu/video.shtml>. Although the quality isn't great, this video of a joint presentation by John Nash and Harald Kuhn touches upon several concepts from the course.
- www.olg.ca. Find out what the Ontario Lottery and Gaming commission has to say about gambling in Ontario. Includes various interesting statistics on how much we spend on gambling, and about the plans for introducing online gambling in Ontario.
- www.wizardofodds.com. Site dedicated to odds and strategies for casino games, hosted by a reputable gambling scholar.

INDEX

- 21flag, 13
- additivity, 61
- AKQ game, 92
- anchoring, 74
- arms race, 36
- Arrow, Kenneth, 36
- assurance, 100
- astralagus, 41
- Aumann, Robert, 37
- availability, 75

- backward induction, 9
- battle of the sexes game, 32, 103
- Bayes' Theorem, 71
- Bernouilli, Nicolaus, 16
- best response, 33
- bias, 73, 74
- Birthday problem, 40, 78
- bluffing, 92
- bluffing frequency, 94
- bluffing odds, 62
- Book on Games of Chance, 43
- Borel, Emile, 34
- Boy or Girl problem, 40, 77
- break-even point, 62, 63
- brinkmanship, 105
- burning bridges, 102

- calling frequency, 94
- Camerer, Colin, 19
- Cardano's Law, 44, 70
- Cardano, Girolamo, 42
- Caro, Mike, 20

- centipede game, 19
- chess, 3
- chicken game, 29, 88
- Cicero, Marcus Tullius, 42
- Cold War, 36
- Collbohm, Franklin, 35
- commitment, 99
- complementation, 65
- conditional probability, 70
- contract, 101
- cooperate, 23
- coordination game, 30
- credibility, 101

- Dantzig, George, 36
- defect, 23
- dictator game, 18
- difference, 65
- Doomsday Device, 103
- Doping test problem, 40, 78
- Dresher, Melvin, 22, 36
- Durrrr Million Dollar Challenge, 49
- Dwan, Tom, 49

- edge, 4
- equilibrium, 9
- Eudaemon Collective, 46
- event, 44, 64
- ex-showdown value, 96
- Expected Value, 56, 58

- Farmer, James Doyne, 46
- Fermat, Pierre, 42
- Flood, Merrill, 22

- focal point, 31
- free rider dilemma, 24
- Galileo, Galilei, 47
- Gambler's Fallacy, 72
- gambling
 - history of, 41
- game, 3
 - half-street, 96
 - sequential, 3
 - simultaneous, 22
- game theory, 34
 - behavioral, 18
 - experimental, 18
- game tree, 4
- Grand Duke's problem, 47
- Harsanyi, John, 35
- heuristic, 73
 - anchoring, 74
 - availability, 75
 - representativeness, 75
- heuristics, 74
- house edge, 61
- IEDS, 29
- implicit collusion, 96
- implied odds, 62
- independence, 72
- indifference, 85
- information
 - complete, 5
- information theory, 46
- intersection, 65
- investment game, 3, 8
- irrationality, 104
- Jagger, Joseph, 46
- Kahneman, Daniel, 16, 73
- Kelly Criterion, 108
- Kelly, John, 108
- Kriegspiel, 35
- Kuhn poker, 96
- Kuhn, Harald, 96
- leaf, 4
- levels of thinking, 22
- Luce, R. Duncan, 36
- matching pennies game, 84
- Minimax Theorem, 34
- monopoly game, 11, 100
- Monty Hall problem, 40, 76
- Morgenstern, Oskar, 34
- Morton's Theorem, 97
- Nash Equilibrium, 33, 35
 - in mixed strategies, 88
- Nash, John, 35–37
- NIM, 13
- node, 4
- O.J. Simpson trial, 41, 78
- Oppenheimer, Robert, 35
- outcome (of a random experiment), 44
 - favorable, 44
 - possible, 44
- outs, 49
 - clean, 49
 - tainted, 49
- Packard, Norman, 46
- Pascal's Wager, 57
- Pascal, Blaise, 42, 56
- path, 6
- payoff matrix, 23
- payoffs, 4, 101
 - cardinal, 16
 - ordinal, 16
- Pensées, 56
- Pirate game, 14
- play, 6
 - exploitative, 92
- poison pill, 103
- poker
 - psychology, 20
- poker hand
 - drawing, 48
 - made, 48
- Pot Limit Omaha, 49
- pot odds, 61
- Poundstone, William, 36

- prior probability, 71
- prisoner's dilemma, 22
- probability
 - laws of, 67
- probability theory, 64
 - history of, 41
- Project Manhattan, 35
- promise, 99
 - compellent, 101
 - deterrent, 101
- prospect theory, 16
- pruning, 7

- RAND corporation, 35
- random experiment, 42, 44
- random variable, 44
- Rapoport, Anatol, 36
- rationality, 21
- representativeness, 75
- reputation, 102
- risk of ruin, 105
- rollback equilibrium, 7
- roulette, 45, 59

- salami tactics, 103
- sample point, 64
- sample space, 44, 64
- satisficing, 74
- Schelling point, 31
- Schelling, Thomas, 31, 36, 99
- Schoonmaker, Alan, 20
- screening test, 104
- Selten, Reinhard, 35
- Shannon, Claude, 46
- Shapley, Lloyd, 36
- Shubik, Martin, 36
- Simon, Herbert, 74
- stag hunt, 32
- strategic move, 32, 37, 98
 - (un)conditional, 99
- strategy, 5
 - dominant
 - strictly, 26
 - weakly, 26
 - dominated
 - strictly, 26
 - exploitive, 97
 - mixed, 85
 - optimal, 97
 - profile, 25
 - pure, 25
- sucker bet, 12

- Texas Holdem, 48
- Thorp, Edward, 46
- threat, 99
 - compellent, 101
 - deterrent, 101
- tragedy of the commons, 24
- Tucker, Albert, 22
- Tversky, Amos, 16, 73
- Tzu, Sun, 35

- ultimatum game, 16, 102
- union, 65
- util, 15
- utility, 15
- utility theory, 16

- valence effect, 75
- value
 - of a game, 87
- Venn diagram, 65
- Von Neumann, John, 34, 36

- warning, 100
- Williams, John, 36
- winner's curse, 12
- World Government, 36

- ZECK, 13