

CHAPTER 1

Life Is Poker, Not Chess

Pete Carroll and the Monday Morning Quarterbacks

One of the most controversial decisions in Super Bowl history took place in the closing seconds of Super Bowl XLIX in 2015. The Seattle Seahawks, with twenty-six seconds remaining and trailing by four points, had the ball on second down at the New England Patriots' one-yard line. Everybody expected Seahawks coach Pete Carroll to call for a handoff to running back Marshawn Lynch. Why wouldn't you expect that call? It was a short-yardage situation and Lynch was one of the best running backs in the NFL.

Instead, Carroll called for quarterback Russell Wilson to pass. New England intercepted the ball, winning the Super Bowl moments later. The headlines the next day were brutal:

- *USA Today*: "What on Earth Was Seattle Thinking with Worst Play Call in NFL History?"
- *Washington Post*: "'Worst Play-Call in Super Bowl History' Will Forever Alter Perception of Seahawks, Patriots"
- FoxSports.com: "Dumbest Call in Super Bowl History Could Be Beginning of the End for Seattle Seahawks"
- *Seattle Times*: "Seahawks Lost Because of the Worst Call in Super Bowl History"
- The *New Yorker*: "A Coach's Terrible Super Bowl Mistake"

Although the matter was considered by nearly every pundit as beyond debate, a few outlying voices argued that the play choice was sound, if not brilliant. Benjamin Morris's analysis on FiveThirtyEight.com and Brian Burke's on Slate.com convincingly argued that the decision to throw the ball was totally defensible, invoking clock-management and end-of-game considerations. They also pointed out that an interception was an extremely unlikely outcome. (Out of sixty-six passes attempted from an opponent's one-yard line during the season, zero had been intercepted. In the previous fifteen seasons, the interception rate in that situation was about 2%.)

Those dissenting voices didn't make a dent in the avalanche of criticism directed at Pete Carroll. Whether or not you buy into the contrarian analysis, most people didn't want to give Carroll the credit for having thought it through, or having *any* reason at all for his call. That raises the question: Why did *so* many people *so* strongly believe that Pete Carroll got it *so* wrong?

We can sum it up in four words: the play didn't work.

Take a moment to imagine that Wilson completed the pass for a game-winning touchdown. Wouldn't the headlines change to "Brilliant Call" or "Seahawks Win Super Bowl on Surprise Play" or "Carroll Outsmarts Belichick"? Or imagine the pass had been incomplete and the Seahawks scored (or didn't) on a third- or fourth-down running play. The headlines would be about those other plays. What Pete Carroll called on second down would have been ignored.

Carroll got unlucky. He had control over the quality of the play-call decision, but not over how it turned out. It was exactly because he didn't get a favorable result that he took the heat. He called a play that had a high percentage of ending in a game-winning touchdown or an incomplete pass (which would have allowed two more plays for the Seahawks to hand off the ball to Marshawn Lynch). He made a good-quality decision that got a bad result.

Pete Carroll was a victim of our tendency to equate the quality of a decision with the quality of its outcome. Poker players have a word for this: "resulting." When I started playing poker, more experienced players warned me about the dangers of resulting, cautioning me to resist the temptation to change my strategy just because a few hands didn't turn out well in the short run.

Pete Carroll understood that his universe of critics was guilty of resulting. Four days after the Super Bowl, he appeared on the *Today* show and acknowledged, “It was the worst *result* of a call ever,” adding, “The call would have been a great one if we catch it. It would have been just fine, and nobody would have thought twice about it.”

Why are we so bad at separating luck and skill? Why are we so uncomfortable knowing that results can be beyond our control? Why do we create such a strong connection between results and the quality of the decisions preceding them? How can we avoid falling into the trap of the Monday Morning Quarterback, whether it is in analyzing someone else’s decision or in making and reviewing the decisions in our own lives?

The hazards of resulting

Take a moment to imagine your best decision in the last year. Now take a moment to imagine your worst decision.

I’m willing to bet that your best decision preceded a good result and the worst decision preceded a bad result.

That is a safe bet for me because resulting isn’t just something we do from afar. Monday Morning Quarterbacks are an easy target, as are writers and bloggers providing instant analysis to a mass audience. But, as I found out from my own experiences in poker, resulting is a routine thinking pattern that bedevils all of us. Drawing an overly tight relationship between results and decision quality affects our decisions every day, potentially with far-reaching, catastrophic consequences.

When I consult with executives, I sometimes start with this exercise. I ask group members to come to our first meeting with a brief description of their best and worst decisions of the previous year. I have yet to come across someone who doesn’t identify their best and worst *results* rather than their best and worst decisions.

In a consulting meeting with a group of CEOs and business owners, one member of the group identified firing the president of his company as his worst decision. He explained, “Since we fired him, the search for a replacement has been awful. We’ve had two different people on the job. Sales

are falling. The company's not doing well. We haven't had anybody come in who actually turns out to be as good as he was."

That sounds like a disastrous result, but I was curious to probe into why the CEO thought the decision to fire his president was so bad (other than that it didn't work out).

He explained the decision process and the basis of the conclusion to fire the president. "We looked at our direct competitors and comparable companies, and concluded we weren't performing up to their level. We thought we could perform and grow at that level and that it was probably a leadership issue."

I asked whether the process included working with the president to understand his skill gaps and what he could be doing better. The company had, indeed, worked with him to identify his skill gaps. The CEO hired an executive coach to work with him on improving his leadership skills, the chief weakness identified.

In addition, after executive coaching failed to produce improved performance, the company considered splitting the president's responsibilities, having him focus on his strengths and moving other responsibilities to another executive. They rejected that idea, concluding that the president's morale would suffer, employees would likely perceive it as a vote of no confidence, and it would put extra financial pressure on the company to split a position they believed one person could fill.

Finally, the CEO provided some background about the company's experience making high-level outside hires and its understanding of the available talent. It sounded like the CEO had a reasonable basis for believing they would find someone better.

I asked the assembled group, "Who thinks this was a bad decision?" Not surprisingly, everybody agreed the company had gone through a thoughtful process and made a decision that was reasonable given what they knew at the time.

It sounded like a bad result, not a bad decision. The imperfect relationship between results and decision quality devastated the CEO and adversely affected subsequent decisions regarding the company. The CEO had identified the decision as a mistake solely because it didn't work out. He obviously felt a lot of anguish and regret because of the decision. He stated very clearly that he thought he should have known that the decision to fire the

president would turn out badly. His decision-making behavior going forward reflected the belief that he made a mistake. He was not only resulting but also succumbing to its companion, hindsight bias. Hindsight bias is the tendency, after an outcome is known, to see the outcome as having been inevitable. When we say, “I should have known that would happen,” or, “I should have seen it coming,” we are succumbing to hindsight bias.

Those beliefs develop from an overly tight connection between outcomes and decisions. That is typical of how we evaluate our past decisions. Like the army of critics of Pete Carroll’s decision to pass on the last play of the Super Bowl, the CEO had been guilty of resulting, ignoring his (and his company’s) careful analysis and focusing only on the poor outcome. The decision didn’t work out, and he treated that result as if it were an inevitable consequence rather than a probabilistic one.

In the exercise I do of identifying your best and worst decisions, I never seem to come across anyone who identifies a bad decision where they got lucky with the result, or a well-reasoned decision that didn’t pan out. We link results with decisions even though it is easy to point out indisputable examples where the relationship between decisions and results isn’t so perfectly correlated. No sober person thinks getting home safely after driving drunk reflects a good decision or good driving ability. Changing future decisions based on that lucky result is dangerous and unheard of (unless you are reasoning this out while drunk and obviously deluding yourself).

Yet this is exactly what happened to that CEO. He changed his behavior based on the quality of the result rather than the quality of the decision-making process. He decided he drove better when he was drunk.

Quick or dead: our brains weren’t built for rationality

The irrationality displayed by Pete Carroll’s critics and the CEO should come as no surprise to anyone familiar with behavioral economics. Thanks to the work of many brilliant psychologists, economists, cognitive researchers, and neuroscientists, there are a number of excellent books that explain why

humans are plagued by certain kinds of irrationality in decision-making. (If you are unaware of these books, see the Selected Bibliography and Recommendations for Further Reading.) But here's a summary.

To start, our brains evolved to create certainty and order. We are uncomfortable with the idea that luck plays a significant role in our lives. We recognize the existence of luck, but we resist the idea that, despite our best efforts, things might not work out the way we want. It feels better for us to imagine the world as an orderly place, where randomness does not wreak havoc and things are perfectly predictable. We evolved to see the world that way. Creating order out of chaos has been necessary for our survival.

When our ancestors heard rustling on the savanna and a lion jumped out, making a connection between “rustling” and “lions” could save their lives on later occasions. Finding predictable connections is, literally, how our species survived. Science writer, historian, and skeptic Michael Shermer, in *The Believing Brain*, explains why we have historically (and prehistorically) looked for connections even if they were doubtful or false. Incorrectly interpreting rustling from the wind as an oncoming lion is called a type I error, a false positive. The consequences of such an error were much less grave than those of a type II error, a false negative. A false negative could have been fatal: hearing rustling and always assuming it's the wind would have gotten our ancestors eaten, and we wouldn't be here.

Seeking certainty helped keep us alive all this time, but it can wreak havoc on our decisions in an uncertain world. When we work backward from results to figure out why those things happened, we are susceptible to a variety of cognitive traps, like assuming causation when there is only a correlation, or cherry-picking data to confirm the narrative we prefer. We will pound a lot of square pegs into round holes to maintain the illusion of a tight relationship between our outcomes and our decisions.

Different brain functions compete to control our decisions. Nobel laureate and psychology professor Daniel Kahneman, in his 2011 best-selling *Thinking, Fast and Slow*, popularized the labels of “System 1” and “System 2.” He characterized System 1 as “fast thinking.” System 1 is what causes you to hit the brakes the instant someone jumps into the street in front of your car. It encompasses reflex, instinct, intuition, impulse, and automatic processing. System 2, “slow thinking,” is how we choose, concentrate, and expend mental energy. Kahneman explains how System 1 and System 2 are

capable of dividing and conquering our decision-making but work mischief when they conflict.

I particularly like the descriptive labels “reflexive mind” and “deliberative mind” favored by psychologist Gary Marcus. In his 2008 book, *Kluge: The Haphazard Evolution of the Human Mind*, he wrote, “Our thinking can be divided into two streams, one that is fast, automatic, and largely unconscious, and another that is slow, deliberate, and judicious.” The first system, “the reflexive system, seems to do its thing rapidly and automatically, with or without our conscious awareness.” The second system, “the deliberative system . . . deliberates, it considers, it chews over the facts.”

The differences between the systems are more than just labels. Automatic processing originates in the evolutionarily older parts of the brain, including the cerebellum, basal ganglia, and amygdala. Our deliberative mind operates out of the prefrontal cortex.

Colin Camerer, a professor of behavioral economics at Caltech and leading speaker and researcher on the intersection of game theory and neuroscience, explained to me the practical folly of imagining that we could just get our deliberative minds to do more of the decision-making work. “We have this thin layer of prefrontal cortex made just for us, sitting on top of this big animal brain. Getting this thin little layer to handle more is unrealistic.” The prefrontal cortex doesn’t control most of the decisions we make every day. We can’t fundamentally get more out of that unique, thin layer of prefrontal cortex. “It’s already overtaxed,” he told me.

These are the brains we have and they aren’t changing anytime soon.*
Making more rational decisions isn’t just a matter of willpower or consciously handling more decisions in deliberative mind. Our deliberative capacity is already maxed out. We don’t have the option, once we recognize the problem, of merely shifting the work to a different part of the brain, as if you hurt your back lifting boxes and shifted to relying on your leg muscles.


Both deliberative and reflexive mind are necessary for our survival and advancement. The big decisions about what we want to accomplish recruit the deliberative system. Most of the decisions we execute on the way to achieving those goals, however, occur in reflexive mind. The shortcuts built into the automatic processing system kept us from standing around on the savanna, debating the origin of a potentially threatening sound while its source

devoured us. Those shortcuts keep us alive, routinely executing the thousands of decisions that make it possible for us to live our daily lives.

We need shortcuts, but they come at a cost. Many decision-making missteps originate from the pressure on the reflexive system to do its job fast and automatically. No one wakes up in the morning and says, “I want to be closed-minded and dismissive of others.” But what happens when we’re focused on work and a fluff-headed coworker approaches? Our brain is already using body language and curt responses to get rid of them without flouting conventions of politeness. We don’t deliberate over this; we just do it. What if they had a useful piece of information to share? We’ve tuned them out, cut them short, and are predisposed to dismiss anything we do pick up that varies from what we already know.

Most of what we do daily exists in automatic processing. We have habits and defaults that we rarely examine, from gripping a pencil to swerving to avoid an auto accident. The challenge is not to change the way our brains operate but to figure out how to work within the limitations of the brains we already have. Being aware of our irrational behavior and wanting to change is not enough, in the same way that knowing that you are looking at a visual illusion is not enough to make the illusion go away. Daniel Kahneman used the famous Müller-Lyer illusion to illustrate this.

MÜLLER-LYER ILLUSION



Which of these three lines is longest? Our brain sends us the signal that the second line is the longest, but you can see from adding the measurement lines that they are the same length.



We can measure the lines to confirm they are the same length, but we can't make ourselves unsee the illusion.

What we can do is look for practical work-arounds, like carrying around a ruler and knowing when to use it to check against how your brain processes what you see. It turns out that poker is a great place to find practical strategies to get the execution of our decisions to align better with our goals.

Understanding how poker players think can help us deal with the decision challenges that bedevil us in our workplaces, financial lives, relationships—even in deciding whether or not passing the ball was a brilliant play.

Two-minute warning

Our goal is to get our reflexive minds to execute on our deliberative minds' best intentions. Poker players don't need to know the underlying science to understand the difficulty of reconciling the two systems. Poker players have to make multiple decisions with significant financial consequences in a compressed time frame, and do it in a way that lassoes their reflexive minds to align with their long-term goals. This makes the poker table a unique laboratory for studying decision-making.

Every poker hand requires making at least one decision (to fold your starting cards or play them), and some hands can require up to twenty decisions. During a poker game in a casino card room, players get in about thirty hands per hour. An average hand of poker takes about two minutes to complete, including the time it takes for the dealer to gather, shuffle, and deal the cards between hands. Poker sessions typically last for several hours, with many decisions in every hand. This means a poker player makes hundreds of decisions per session, all of which take place at breakneck speed.

The etiquette and rules of the game discourage players from slowing down the game to deliberate, even when huge financial consequences ride on the decision. If a player takes extra time, another player can “call the clock” on them. This gives the deliberating player all of *seventy seconds* to now make up their mind. That is an eternity in poker time.

Every hand (and therefore every decision) has immediate financial consequences. In a tournament or a high-stakes game, each decision can be worth more than the cost of an average three-bedroom house, and players have to make those decisions more quickly than we decide what to order in a restaurant. Even at lower stakes, most or all of the money a player has on the table is potentially at stake in every decision. Poker players, as a result, must become adept at in-the-moment decision-making or they won’t survive in the profession. That means finding ways to execute their best intentions (deliberated in advance) within the constraints of the speed expected at the table. Making a living at poker requires interpolating between the deliberative and reflexive systems. The best players must find ways to harmonize otherwise irresolvable conflicts.

In addition, once the game is over, poker players must learn from that jumbled mass of decisions and outcomes, separating the luck from the skill, the signal from the noise, and guarding against resulting. That’s the only way to improve, especially when those same under-pressure situations will recur in a variety of forms.

Solving the problem of how to execute is even more important than innate talent to succeed in poker. All the talent in the world won’t matter if a player can’t execute; avoiding common decision traps, learning from results in a rational way, and keeping emotions out of the process as much as possible. Players with awe-inspiring talent clean up on their best nights but go broke plenty of other nights if they haven’t confronted this challenge. The poker

players who stand the test of time have a variety of talents, but what they share is the ability to execute in the face of these limitations.

We all struggle to execute our best intentions. Poker players have the same struggle, with the added challenges of time pressure, in-your-face uncertainty, and immediate financial consequences. That makes poker a great place to find innovative approaches to overcoming this struggle. And the value of poker in understanding decision-making has been recognized in academics for a long time.

Dr. Strangelove

It's hard for a scientist to become a household name. So it shouldn't be surprising that for most people the name John von Neumann doesn't ring a bell.

That's a shame because von Neumann is a hero of mine, and should be to anyone committed to making better decisions. His contributions to the science of decision-making were immense, and yet they were just a footnote in the short life of one of the greatest minds in the history of scientific thought. (And, not coincidentally, he was a poker player.)

After a twenty-year period in which he contributed to practically every branch of mathematics, this is what he did in the last ten years of his life: played a key role on the Manhattan Project, pioneered the physics behind the hydrogen bomb, developed the first computers, figured out the optimal way to route bombers and choose targets at the end of World War II, and created the concept of mutually assured destruction (MAD), the governing geopolitical principle of survival throughout the Cold War. Even after being diagnosed with cancer in 1955 at the age of fifty-two, he served in the first civilian agency overseeing atomic research and development, attending meetings, though in great pain, in a wheelchair for as long as he was physically able.

Despite all he accomplished in science, somehow von Neumann's legacy in popular culture is as one of the models for the title character in Stanley Kubrick's apocalyptic comedy, *Dr. Strangelove*: a heavily accented, crumpled, wheelchair-bound genius whose strategy of relying on mutually assured destruction goes awry when an insane general sends a single bomber on an

unauthorized mission that could trigger the automated firing of all American and Soviet nuclear weapons.

In addition to everything else he accomplished, John von Neumann is also the father of game theory. After finishing his day job on the Manhattan Project, he collaborated with Oskar Morgenstern to publish *Theory of Games and Economic Behavior* in 1944. The Boston Public Library's list of the "100 Most Influential Books of the Century" includes *Theory of Games*. William Poundstone, author of a widely read book on game theory, *Prisoner's Dilemma*, called it "one of the most influential and least-read books of the twentieth century." The introduction to the sixtieth-anniversary edition pointed out how the book was instantly recognized as a classic. Initial reviews in the most prestigious academic journals heaped it with praise, like "one of the major scientific achievements of the first half of the twentieth century" and "ten more such books and the progress of economics is assured."

Game theory revolutionized economics, evidenced by at least eleven economics Nobel laureates connected with game theory and its decision-making implications, including John Nash (a student of von Neumann's), whose life story was chronicled in the Oscar-winning film *A Beautiful Mind*. Game theory has broad applications outside economics, informing the behavioral sciences (including psychology and sociology) as well as political science, biomedical research, business, and numerous other fields.

Game theory was succinctly defined by economist Roger Myerson (one of the game-theory Nobel laureates) as "the study of mathematical models of conflict and cooperation between intelligent rational decision-makers." Game theory is the modern basis for the study of the bulk of our decision-making, addressing the challenges of changing conditions, hidden information, chance, and multiple people involved in the decisions. Sound familiar?

Fortunately, you don't need to know any more than this about game theory to understand its relevance. And the important thing for this book is that John von Neumann modeled game theory on a stripped-down version of poker.

Poker vs. chess

In *The Ascent of Man*, scientist Jacob Bronowski recounted how von Neumann described game theory during a London taxi ride. Bronowski was a chess enthusiast and asked him to clarify. “You mean, the theory of games like chess?”

Bronowski quoted von Neumann’s response: “‘No, no,’ he said. ‘Chess is not a game. Chess is a well-defined form of computation. You may not be able to work out the answers, but in theory there must be a solution, a right procedure in any position. Now, real games,’ he said, ‘are not like that at all. Real life is not like that. Real life consists of bluffing, of little tactics of deception, of asking yourself what is the other man going to think I mean to do. And that is what games are about in my theory.’”

The decisions we make in our lives—in business, saving and spending, health and lifestyle choices, raising our children, and relationships—easily fit von Neumann’s definition of “real games.” They involve uncertainty, risk, and occasional deception, prominent elements in poker. Trouble follows when we treat life decisions as if they were chess decisions.

Chess contains no hidden information and very little luck. The pieces are all there for both players to see. Pieces can’t randomly appear or disappear from the board or get moved from one position to another by chance. No one rolls dice after which, if the roll goes against you, your bishop is taken off the board. If you lose at a game of chess, it must be because there were better moves that you didn’t make or didn’t see. You can theoretically go back and figure out exactly where you made mistakes. If one chess player is more than just a bit better than another, it is nearly inevitable the better player will win (if they are white) or, at least, draw (if they are black). On the rare occasions when a lower-ranked grand master beats a Garry Kasparov, Bobby Fischer, or Magnus Carlsen, it is because the higher-ranked player made identifiable, objective mistakes, allowing the other player to capitalize.

Chess, for all its strategic complexity, isn’t a great model for decision-making in life, where most of our decisions involve hidden information and a much greater influence of luck. This creates a challenge that doesn’t exist in chess: identifying the relative contributions of the decisions we make versus luck in how things turn out.

Poker, in contrast, is a game of *incomplete information*. It is a game of decision-making under conditions of uncertainty over time. (Not coincidentally, that is close to the definition of game theory.) Valuable information remains hidden. There is also an element of luck in any outcome. You could make the best possible decision at every point and still lose the hand, because you don't know what new cards will be dealt and revealed. Once the game is finished and you try to learn from the results, separating the quality of your decisions from the influence of luck is difficult.

In chess, outcomes correlate more tightly with decision quality. In poker, it is much easier to get lucky and win, or get unlucky and lose. If life were like chess, nearly every time you ran a red light you would get in an accident (or at least receive a ticket). If life were like chess, the Seahawks would win the Super Bowl every time Pete Carroll called that pass play.

But life is more like poker. You could make the smartest, most careful decision in firing a company president and still have it blow up in your face. You could run a red light and get through the intersection safely—or follow all the traffic rules and signals and end up in an accident. You could teach someone the rules of poker in five minutes, put them at a table with a world champion player, deal a hand (or several), and the novice could beat the champion. That could never happen in chess.

Incomplete information poses a challenge not just for split-second decision-making, but also for learning from past decisions. Imagine my difficulty as a poker player in trying to figure out if I played a hand correctly when my opponents' cards were never revealed. If the hand concluded after I made a bet and my opponents folded, all I know is that I won the chips. Did I play poorly and get lucky? Or did I play well?

If we want to improve in any game—as well as in any aspect of our lives—we have to learn from the results of our decisions. The quality of our lives is the sum of decision quality plus luck. In chess, luck is limited in its influence, so it's easier to read the results as a signal of decision quality. That more tightly tethers chess players to rationality. Make a mistake and your opponent's play points it out, or it is capable of analysis afterward. There is always a theoretically right answer out there. If you lose, there is little room to off-load losing to any other explanation than your inferior decision-making. You'll almost never hear a chess player say, "I was robbed in that

game!” or, “I played perfectly and caught some terrible breaks.” (Walk the hallways during a break in a poker tournament and you’ll hear a lot of that.)

That’s chess, but life doesn’t look like that. It looks more like poker, where all that uncertainty gives us the room to deceive ourselves and misinterpret the data. Poker gives us the leeway to make mistakes that we never spot because we win the hand anyway and so don’t go looking for them, or the leeway to do everything right, still lose, and treat the losing result as proof that we made a mistake. Resulting, assuming that our decision-making is good or bad based on a small set of outcomes, is a pretty reasonable strategy for learning in chess. But not in poker—or life.

Von Neumann and Morgenstern understood that the world doesn’t easily reveal the objective truth. That’s why they based game theory on poker. Making better decisions starts with understanding this: uncertainty can work a lot of mischief.

A lethal battle of wits

In one of the more famous scenes in *The Princess Bride*, the Dread Pirate Roberts (the love-besotted Westley) catches up to Vizzini, the mastermind who kidnapped Princess Buttercup. Having vanquished Fezzik the Giant in a battle of strength and having outdueled swordsman Inigo Montoya, the Dread Pirate Roberts proposes he and Vizzini compete in a lethal battle of wits, which provides a great demonstration of the peril of making decisions with incomplete information. The pirate produces a packet of deadly iocane powder and, hiding two goblets of wine from view, he empties the packet, and puts one goblet in front of himself and the other in front of Vizzini. Once Vizzini chooses a goblet, they will both drink “and find out who is right and who is dead.”

“It’s all so simple,” Vizzini scoffs. “All I have to do is deduce, from what I know of you, the way your mind works. Are you the kind of man who would put the poison into his own glass, or into the glass of his enemy?” He provides a dizzying series of reasons why the poison can’t (or must) be in one cup, and then in the other. His rant accounts for cleverness, anticipating cleverness, iocane’s origin (the criminal land of Australia), untrustworthiness,

anticipating untrustworthiness, and dueling presumptions related to Westley defeating the giant and the swordsman.

While explaining all this, Vizzini diverts Westley's attention, switches the goblets, and declares that they should drink from the goblets in front of them. Vizzini pauses for a moment and, when he sees Westley drink from his own goblet, confidently drinks from the other.

Vizzini roars with laughter. "You fell victim to one of the classic blunders. The most famous is 'Never get involved in a land war in Asia,' but only slightly less well known is this: 'Never go in against a Sicilian when death is on the line.'"

In the midst of laughing, Vizzini falls over, dead. Buttercup says, "To think, all that time it was your cup that was poisoned."

Westley tells her, "They were both poisoned. I've spent the last two years building up immunity to iocane powder."

Like all of us, Vizzini didn't have all the facts. He considered himself a genius without equal: "Let me put it this way. Have you ever heard of Plato, Aristotle, Socrates? *Morons*." But, also like all of us, he underestimated the amount and effect of what he didn't know.

Suppose someone says, "I flipped a coin and it landed heads four times in a row. How likely is that to occur?"

It feels like that should be a pretty easy question to answer. Once we do the math on the probability of heads on four consecutive 50-50 flips, we can determine that would happen 6.25% of the time ($.50 \times .50 \times .50 \times .50$).

That's making the same mistake as Vizzini. The problem is that we came to this answer without knowing anything about the coin or the person flipping it. Is it a two-sided coin or three-sided or four? If it is two-sided, is it a two-headed coin? Even if the coin is two-sided (heads and tails), is the coin weighted to land on heads more often than tails (but maybe not always)? Is the flipper a magician who is capable of influencing how the coin lands? This information is all incomplete, yet we answered the question as if we had examined the coin and knew everything about it. We never considered that both goblets might be poisoned. ("Inconceivable" would have been Vizzini's term, had he been able to comment on his own death.)

Now if that person flipped the coin 10,000 times, giving us a sufficiently large sample size, we could figure out, with some certainty, whether the coin is fair. Four flips simply isn't enough to determine much about the coin.

We make this same mistake when we look for lessons in life's results. Our lives are too short to collect enough data from our own experience to make it easy to dig down into decision quality from the small set of results we experience. If we buy a house, fix it up a little, and sell it three years later for 50% more than we paid, does that mean we are smart at buying and selling property, or at fixing up houses? It could, but it could also mean there was a big upward trend in the market and buying almost any piece of property would have made just as much money. Or maybe buying that same house and not fixing it up at all might have resulted in the same (or even better) profit. A lot of previously successful house flippers had to face that real possibility between 2007 and 2009.

When someone asks you about a coin they flipped four times, there is a correct answer: "I'm not sure."

"I'm not sure": using uncertainty to our advantage

Just as we have problems with resulting and hindsight bias, when we evaluate decisions solely on how they turn out, we have a mirror-image problem in making prospective decisions. We get only one try at any given decision—one flip of the coin—and that puts great pressure on us to feel we have to be certain before acting, a certainty that necessarily will overlook the influences of hidden information and luck.

Famed novelist and screenwriter William Goldman (who wrote *The Princess Bride*, as well as *Misery* and *Butch Cassidy and the Sundance Kid*) reflected on his experiences working with actors like Robert Redford, Steve McQueen, Dustin Hoffman, and Paul Newman at the height of their successful careers. What did it mean to be a "movie star"? He quoted an actor who explained the type of characters he wanted to play: "I don't want to be the man who learns. I want to be the man who *knows*."

We are discouraged from saying "I don't know" or "I'm not sure." We regard those expressions as vague, unhelpful, and even evasive. But getting comfortable with "I'm not sure" is a vital step to being a better decision-maker. We have to make peace with *not knowing*.

Embracing “I’m not sure” is difficult. We are trained in school that saying “I don’t know” is a bad thing. Not knowing in school is considered a failure of learning. Write “I don’t know” as an answer on a test and your answer will be marked wrong.

Admitting that we don’t know has an undeservedly bad reputation. Of course, we want to encourage acquiring knowledge, but the first step is understanding what we don’t know. Neuroscientist Stuart Firestein’s book *Ignorance: How It Drives Science* champions the virtue of recognizing the limits of our knowledge. (You can get a taste of the book by watching his TED Talk, “The Pursuit of Ignorance.”) In the book and the talk, Firestein points out that in science, “I don’t know” is not a failure but a necessary step toward enlightenment. He backs this up with a great quote from physicist James Clerk Maxwell: “Thoroughly conscious ignorance is the prelude to every real advance in science.” I would add that this is a prelude to every great decision that has ever been made.

What makes a decision great is not that it has a great outcome. A great decision is the result of a good process, and that process must include an attempt to accurately represent our own state of knowledge. That state of knowledge, in turn, is some variation of “I’m not sure.”

“I’m not sure” does not mean that there is no objective truth. Firestein’s point is, in fact, that acknowledging uncertainty is the first step in executing on our goal to get closer to what is objectively true. To do this, we need to stop treating “I don’t know” and “I’m not sure” like strings of dirty words.

What if we shifted our definition of “I don’t know” from the negative frame (“I have no idea” or “I know nothing about that,” which feels like we lack competence or confidence) to a more neutral frame? What if we thought about it as recognizing that, although we might know something about the chances of some event occurring, we are still not sure how things will turn out in any given instance? That is just the truth. If we accept that, “I’m not sure” might not feel so bad.

What good poker players and good decision-makers have in common is their comfort with the world being an uncertain and unpredictable place. They understand that they can almost never know exactly how something will turn out. They embrace that uncertainty and, instead of focusing on being sure, they try to figure out how *unsure* they are, making their best guess at the chances that different outcomes will occur. The accuracy of those guesses

will depend on how much information they have and how experienced they are at making such guesses. This is part of the basis of all bets.

To be sure, an experienced poker player is more likely to make a better guess than a novice player at determining the chances they will win or lose a hand. The experienced player knows the math better and is better able to narrow down what their opponents' cards might be based on how players behave with certain types of hands. They will also be better at figuring out the choices their opponents are likely to make with those cards. So, yes, more experience will allow the player to narrow down the possibilities. None of that experience, however, makes it possible for a poker player to know how any given hand will turn out.

This is true in any field. An expert trial lawyer will be better than a new lawyer at guessing the likelihood of success of different strategies and picking a strategy on this basis. Negotiating against an adversary whom we have previously encountered gives us a better guess at what our strategy should be. An expert in any field will have an advantage over a rookie. But neither the veteran nor the rookie can be sure what the next flip will look like. The veteran will just have a better guess.

It is often the case that our best choice doesn't even have a particularly high likelihood of succeeding. A trial lawyer with a tough case could be choosing among strategies that are all more likely to fail than to succeed. The goal of a lawyer in that situation is to identify the different possible strategies, figure out their best guess of the chance of success for each unpromising alternative, and pick the least awful one to maximize the quality of the outcome for their client. That's true in any business. Start-ups have very low chances of succeeding but they try nonetheless, attempting to find the best strategy to achieve the big win, even though none of the strategies is highly likely to create success for the company. This is still worthwhile because the payoff can be so large.

There are many reasons why wrapping our arms around uncertainty and giving it a big hug will help us become better decision-makers. Here are two of them. First, "I'm not sure" is simply a more accurate representation of the world. Second, and related, when we accept that we can't be sure, we are less likely to fall into the trap of black-and-white thinking.

Imagine you're stepping on a traditional medical scale. It has two weight bars, one with notches at fifty-pound intervals and the other with notches at

one-pound intervals. This allows the user to measure their weight down to the pound. What would happen if your doctor used a scale with only one bar with just two notches, one at fifty pounds and one at five hundred pounds, with no way to measure anything in between? Good luck getting medical advice after the person weighing you writes one or the other on your chart. You could only be morbidly obese or severely underweight. It would be impossible to make good decisions about your weight with such a poor model.

The same holds true for just about all of our decisions. If we misrepresent the world at the extremes of right and wrong, with no shades of grey in between, our ability to make good choices—choices about how we are supposed to be allocating our resources, what kind of decisions we are supposed to be making, and what kind of actions we are supposed to be taking—will suffer.

The secret is to make peace with walking around in a world where we recognize that we are not sure and that's okay. As we learn more about how our brains operate, we recognize that we don't perceive the world objectively. But our goal should be to *try*.

Redefining wrong

When I attend charity poker tournaments, I will often sit in as the dealer and provide a running commentary at the final table. The atmosphere at these final tables is fun and raucous. Everyone running the event has had a long night and is breathing a sigh of relief. There is typically a big crowd around the table including friends and families of the players, rooting them on (or vocally rooting against them). If people have been drinking, then . . . people have been drinking. Everyone is having a good time.

When players have put all their chips in the pot, there is no more betting on the hand. After an all-in situation, the players in the hand turn their cards faceup on the table so that everyone can see them before I deal the remaining cards. This makes it fun for the audience, because they get to see each player's position in the hand and the drama mounts. With the cards faceup, I can determine the likelihood each player will win the hand, and announce the percentage of the time each hand will win in the long run.

At one such tournament, I told the audience that one player would win 76% of the time and the other would win 24% of the time. I dealt the remaining cards, the last of which turned the 24% hand into the winner. Amid the cheers and groans, someone in the audience called out, “Annie, you were wrong!”

In the same spirit that he said it, I explained that I wasn’t. “I said that would happen 24% of the time. That’s not zero. You got to see part of the 24%!”

A few hands later, almost the same thing happened. Two players put all of their chips in the pot and they turned their cards faceup. One player was 18% to win and the other 82% to win the hand. Again, the player with the worse hand when they put in their chips hit a subsequent lucky card to win the pot.

This time that same guy in the crowd called out, “Look, it was the 18%!” In that aha moment, he changed his definition of what it meant to be wrong. When we think in advance about the chances of alternative outcomes and make a decision based on those chances, it doesn’t automatically make us wrong when things don’t work out. It just means that one event in a set of possible futures occurred.

Look how quickly you can begin to redefine what it means to be wrong. Once we start thinking like this, it becomes easier to resist the temptation to make snap judgments after results or say things like “I knew it” or “I should have known.” Better decision-making and more self-compassion follow.

The public-at-large is often guilty of making black-and-white judgments about the “success” or “failure” of probabilistic thinking. When the UK voted to leave the European Union (“Brexit”) in July 2016, it was an unlikely result. Betting shops had set odds heavily favoring a vote for Remain. That does not mean the betting shops had an opinion that Remain would win the day. The goal of the bookmaker is to make sure the amount of money bet on either side is equal, so that the losers essentially pay the winners while the bookmaker just takes their fee. They aim to have no stake in the outcome and adjust the odds accordingly. The bookmaker’s odds reflect the market’s view, essentially our collective best guess of what is fair.

That didn’t stop even sophisticated people from resulting, declaring after the vote came in Leave that the bookmakers made a mistake. The chief strategist at one Swiss bank told the *Wall Street Journal*, “I can’t remember

any time when the bookies were so wrong.” One of America’s most famous lawyers and professors, Alan Dershowitz, made this same error. Asserting in September 2016 that the Clinton-Trump election was too difficult to make any predictions about, he said, “Think about the vote on Brexit. Virtually all the polls—including exit polls that asked voters how they voted—got it wrong. The financial markets got it wrong. The bookies got it wrong.”

Just like my spectator, Dershowitz missed the point. Any prediction that is not 0% or 100% can’t be wrong solely because the most likely future doesn’t unfold. When the 24% result happened at the final table of the charity tournament, that didn’t reflect inaccuracy about the probabilities as determined before that single outcome. Long shots hit some of the time. Blaming the oddsmakers or the odds themselves assumes that once something happens, it was bound to have happened and anyone who didn’t see it coming was wrong.

The same thing happened after Donald Trump won the presidency. There was a huge outcry about the polls being wrong. Nate Silver, the founder of FiveThirtyEight.com, drew a lot of that criticism. But he never said Clinton was a sure thing. Based on his aggregation and weighting of polling data, he had Trump between 30% and 40% to win (approximately between two-to-one and three-to-two against) in the week before the election. An event predicted to happen 30% to 40% of the time will happen a lot.

Being a poker player, I’ve played out more two-to-one shots in my tournament career than I could possibly count. A lot of those have been situations where the tournament was on the line for me. If I lose the hand, I’m out of the tournament. If I win, I earn a huge pot, maybe even winning the entire tournament. I know viscerally how likely 60–40 and 70–30 favorites are to lose (and, of course, the opposite). When people complained that Nate Silver did his job poorly because he had Clinton favored, I thought, “Those people haven’t gotten all their chips in a pot with a pair against a straight draw and lost.” Or, more likely, they’ve had those things happen throughout their lives and didn’t realize that’s what 30% or 40% feels like.

Decisions are bets on the future, and they aren’t “right” or “wrong” based on whether they turn out well on any particular iteration. An unwanted result doesn’t make our decision wrong if we thought about the alternatives and probabilities in advance and allocated our resources accordingly, as my client the CEO and Pete Carroll both did. It would be absurd for me, after making a

big bet on the best possible starting hand (a pair of aces) and losing, to spend a lot of time thinking that I was wrong to make the decision to play the hand in the first place. That would be resulting.

When we think probabilistically, we are less likely to use adverse results alone as proof that we made a decision error, because we recognize the possibility that the decision might have been good but luck and/or incomplete information (and a sample size of one) intervened.

Maybe we made the best decision from a set of unappealing choices, none of which were likely to turn out well.

Maybe we committed our resources on a long shot because the payout more than compensated for the risk, but the long shot didn't come in this time.

Maybe we made the best choice based on the available information, but decisive information was hidden and we could not have known about it.

Maybe we chose a path with a very high likelihood of success and got unlucky.

Maybe there were other choices that might have been better *and* the one we made wasn't wrong or right but somewhere in between. The second-best choice isn't wrong. By definition, it is more right (or less wrong) than the third-best or fourth-best choice. It is like the scale at the doctor's office: there are a lot more choices other than the extremes of obesity or anorexia. For most of our decisions, there will be a lot of space between unequivocal "right" and "wrong."

When we move away from a world where there are only two opposing and discrete boxes that decisions can be put in—right or wrong—we start living in the continuum between the extremes. Making better decisions stops being about wrong or right but about calibrating among all the shades of grey.

Redefining wrong is easiest in situations where we know the mathematical facts in advance. In the charity-tournament final-table example with the players' cards faceup, or when I get all my chips in with the best possible starting hand, the hidden information is removed. We can make a clear calculation. If we have that unquestionably right and make an allocation of resources (a bet) on the calculation, we can more naturally get to "I wasn't wrong just because it didn't turn out well and I shouldn't change my behavior." When the chances are known, we are tethered more tightly to a

rational interpretation of the influence of luck. It feels a little more like chess that way.

There is no doubt it is harder to get there when we add in hidden information on top of the influence of luck. Untethered from seeing what the coin actually looks like, we are more likely to anchor ourselves to the way things turned out as the sole signal for whether we were right or wrong. We are more likely to declare, “I told you so!” or “I should have known!” When we start doing that, compassion goes out the window. Just ask Pete Carroll.

Redefining wrong allows us to let go of all the anguish that comes from getting a bad result. But it also means we must redefine “right.” If we aren’t wrong just because things didn’t work out, then we aren’t right just because things turned out well. Do we win emotionally to making that mindset trade-off?

Being right feels really good. “I was right,” “I knew it,” “I told you so”—those are all things that we say, and they all feel very good to us. Should we be willing to give up the good feeling of “right” to get rid of the anguish of “wrong”? Yes.

First, the world is a pretty random place. The influence of luck makes it impossible to predict exactly how things will turn out, and all the hidden information makes it even worse. If we don’t change our mindset, we’re going to have to deal with being wrong a lot. It’s built into the equation.

Poker teaches that lesson. A great poker player who has a good-size advantage over the other players at the table, making significantly better strategic decisions, will still be losing over 40% of the time at the end of eight hours of play. That’s a whole lot of wrong. And it’s not just confined to poker.

The most successful investors in start-up companies have a majority of bad results. If you applied to NASA’s astronaut program or the NBC page program, both of which have drawn thousands of applicants for a handful of positions, things will go your way a minority of the time, but you didn’t necessarily do anything wrong. Don’t fall in love or even date anybody if you want only positive results. The world is structured to give us lots of opportunities to feel bad about being wrong if we want to measure ourselves by outcomes. Don’t fall for it!

Second, being wrong hurts us more than being right feels good. We know from Daniel Kahneman and Amos Tversky’s work on loss aversion, part of prospect theory (which won Kahneman the Nobel Prize in Economics in

2002), that losses in general feel about two times as bad as wins feel good. So winning \$100 at blackjack feels as good to us as losing \$50 feels bad to us. Because being right feels like winning and being wrong feels like losing, that means we need two favorable results for every one unfavorable result just to break even emotionally. Why not live a smoother existence, without the swings, especially when the losses affect us more intensely than the wins?

• • •

Are you ready to really wrap your arms around uncertainty, like great decision-makers do? Are you ready to embrace this redefinition of wrong, and to recognize you are always guessing and that those guesses drive how you place your resources? Getting comfortable with this realignment, and all the good things that follow, starts with recognizing that you've been betting all along.

CHAPTER 3

Bet to Learn: Fielding the Unfolding Future

Nick the Greek, and other lessons from the Crystal Lounge

When I first started playing poker, I lived in Columbus, Montana, population 1,200. The nearest poker game was forty miles away in downtown Billings, in the basement of a bar called the Crystal Lounge. Every day, I drove those forty miles, arriving by early afternoon. I would play until the evening, and drive home.

The game was filled with characters out of a clichéd vision of Montana: ranchers and farmers killing time in the off-season, filling the basement room with smoke wafting over the brims of their cowboy hats. It was 1992 but it could have just as easily been 1952 from the décor and the grizzled countenances of the locals. The only thing suggesting that John Wayne wasn't going to mosey on in was a handful of misfits, including me (a woman and the youngest player by a few decades, on the lam from defending my dissertation at the University of Pennsylvania) and a player named "Nick the Greek."

If your name is Nick, you come from Greece, and you gamble, they're going to call you Nick the Greek, as sure as they'll call you Tiny if you weigh more than 350 pounds. (And, yes, there was a guy named Tiny, real name Elwood, who regularly played in the game.) This Nick the Greek was of the small-time Billings variety. He was the general manager of the hotel across

the street, having gotten transferred from Greece by the hotel chain. He left work for a couple of hours every afternoon like clockwork to play in the game.

Nick the Greek had formed an unusual set of beliefs that guided his poker decisions. I knew this because he described them to me and the other players, at length, using the results of particular hands to punctuate his points. He fixated on the relatively common belief that the element of surprise was important in poker. (Don't be predictable, mix up your play—that sort of stuff.) Then he jacked it up on steroids. To him, a starting pair of aces, the mathematically best two cards you can get dealt, was the worst hand because everyone predictably played it.

“They always expect you to have aces. You get killed with that hand.”

By that logic, he explained, the very best two starting cards to play were the *mathematically weakest* two cards you could receive, a seven and a deuce of different suits—a hand almost any player avoids.

“I bet you never saw it coming,” he would say when he turned over that hand and won a pot. And because he played seven-deuce all the time, occasionally the stars would line up and he'd win. I also remember times when he threw away a pair of aces, faceup, at the first sign of a bet. (Never mind that he was compromising a vital element of his strategy of subterfuge by constantly *showing and telling* us he was doing this. Given that he had such an entrenched set of beliefs, it's not surprising that he didn't see the incongruity.)

Nick the Greek, needless to say, rarely came out ahead. Yet he never changed his strategy, often complaining about his bad luck when he lost, though never in a bitter way. He was a friendly guy, pleasant to play with—the perfect poker opponent. I tried to time my daily arrivals so I'd be in the game when he made his afternoon appearance.

One day, Nick the Greek didn't show up to the game. When I asked where he was, another player muttered, confidentially (though it seemed everyone in the game already knew about this), “Oh, he got sent back.”

“Sent back?”

“Yeah, to Greece. They deported him.”

I can't say that Nick the Greek's deportation was the result of his wacky poker beliefs, but I have my suspicions. Other players speculated that he went

broke, or dipped into the till at the hotel, or lost his work visa because he was playing poker every day on company time.

I can say that Nick the Greek lost a lot of money based on his beliefs—or, more accurately, because he ignored lots of feedback that his strategy was a losing one. He eventually went broke because he didn't recognize learning opportunities as they arose.

If Nick the Greek were unique in his resistance to learning from the outcomes he was having at the poker table, I suppose he would just be a footnote for me, a funny story of a guy unique in his ability to hold tight to his strategy despite that strategy resulting in a lot of losing. But, while an extreme case to be sure, Nick the Greek wasn't all that unique. And that was a puzzle for me. I was taught, as all psychology students are, that learning occurs when you get lots of feedback tied closely in time to decisions and actions. If we took that at face value, poker would be an ideal learning environment. You make a bet, get an immediate response from opponents, and win or lose the hand (with real-money consequences), all within minutes.

So why was Nick the Greek, who had been playing for years, unable to learn from his mistakes? Why was a novice like me cleaning up in the game? The answer is that while experience is necessary to becoming an expert, it's not sufficient.

Experience can be an effective teacher. But, clearly, only some students listen to their teachers. The people who learn from experience improve, advance, and (with a little bit of luck) become experts and leaders in their fields. I benefited from adopting the learning habits of some of the phenomenal poker players I was exposed to along the way. We can all benefit from those practical strategies to become better decision-makers. Thinking in bets can help us get there.

But before getting to the solutions, we must first understand the problem. What are the obstacles in our way that make learning from experience so difficult? We all clearly have a desire to reach our long-term goals. Listening to what our outcomes have to teach us is necessary to do that. So what is systematically getting in the way?

Outcomes are feedback

We can't just "absorb" experiences and expect to learn. As novelist and philosopher Aldous Huxley recognized, "Experience is not what happens to a man; it is what a man does with what happens to him." There is a big difference between getting experience and becoming an expert. That difference lies in the ability to identify when the outcomes of our decisions have something to teach us and what that lesson might be.

Any decision, whether it's putting \$2 on Count de Change at the racetrack or telling your kids they can eat whatever they want, is a bet on what will likely create the most favorable future for us. The future we have bet on unfolds as a series of outcomes. We bet on staying up late to watch the end of a football game and we sleep through our alarm, wake up tired, get to work late, and get reprimanded by the boss. Or we stay up late and any of the myriad other outcomes follows, including waking up perfectly on time and making it to work early. Whichever future actually unfolds, when we decide to stay up late to see the end of the game, we are making a bet that we will be happier in the future for having seen the final play. We bet on moving to Des Moines and we find our dream job, meet the love of our life, and take up yoga. Or, like John Hennigan, we move there, hate it within two days, and have to buy our way home for \$15,000. We bet on firing a division president or calling a pass play, and the future unfolds as it does. We can represent this like so:

BETTING ON A FUTURE


BELIEF → BET → (SET OF OUTCOMES)

As the future unfolds into a set of outcomes, we are faced with another decision: Why did something happen the way it did?

How we figure out what—if anything—we should learn from an outcome becomes another bet. As outcomes come our way, figuring out whether those outcomes were caused mainly by luck or whether they were the predictable

result of particular decisions we made is a bet of great consequence. If we determine our decisions drove the outcome, we can feed the data we get following those decisions back into belief formation and updating, creating a learning loop:

LEARNING LOOP 1



We have the opportunity to learn from the way the future unfolds to improve our beliefs and decisions going forward. The more evidence we get from experience, the less uncertainty we have about our beliefs and choices. Actively using outcomes to examine our beliefs and bets closes the feedback loop, reducing uncertainty. This is the heavy lifting of how we learn.

Ideally, our beliefs and our bets improve with time as we learn from experience. Ideally, the more information we have, the better we get at making decisions about which possible future to bet on. Ideally, as we learn from experience we get better at assessing the likelihood of a particular outcome given any decision, making our predictions about the future more accurate. As you may have guessed, when it comes to how we process experience, “ideally” doesn’t always apply.

Learning might proceed in a more ideal way if life were more like chess than poker. The connection between outcome quality and decision quality would be clearer because there would be less uncertainty. The challenge is that any single outcome can happen for multiple reasons. The unfolding future is a big data dump that we have to sort and interpret. And the world doesn’t connect the dots for us between outcomes and causes.

If a patient comes into a doctor’s office with a cough, the doctor must work backward from that one symptom, that one outcome of a possible

disease process, to decide among the multiple reasons the patient might have that cough. Is it because of a virus? Bacteria? Cancer? A neurological disorder? Because a cough looks roughly the same whether it is from cancer or a virus, working backward from the symptom to the cause is difficult. The stakes are high. Misdiagnose the cause, and the patient might die. That is why doctors require years of training to properly diagnose patients.

When the future coughs on us, it is hard to tell why.

Imagine calls to a customer by two salespeople from the same company. In January, Joe pitches the company's products and gets \$1,000 in orders. In August, Jane calls on the same customer and gets \$10,000 in orders. What gives? Was it because Jane is a better salesperson than Joe? Or was it because the company updated its product line in February? Did a low-cost competitor go out of business in April? Or is the difference in their success due to any of a variety of other unconsidered reasons? It's hard to know why because we can't go back in time and run the controlled experiment where Joe and Jane switch places. And the way the company sorts this outcome can affect decisions on training, pricing, and product development.

This problem is top of mind for poker players. Most poker hands end in a cloud of incomplete information: one player bets, no one calls the bet, the bettor is awarded the pot, and no one is required to reveal their hidden cards. After those hands, the players are left guessing why they won or lost the hand. Did the winner have a superior hand? Did the loser fold the best hand? Could the player who won the hand have made more money if they chose a different line of play? Could the player who lost have made the winner forfeit if they chose to play the hand differently? In answering these questions, none of the players knows what cards their opponents actually held, or how the players would have reacted to a different sequence of betting decisions. How poker players adjust their play from experience determines their future results. How they fill in all those blanks is a vitally important bet on whether they get better at the game.

We are good at identifying the “-ER” goals we want to pursue (better, smarter, richer, healthier, whatever). But we fall short in achieving our “-ER” because of the difficulty in executing all the little decisions along the way to our goals. The bets we make on when and how to close the feedback loop are part of the execution, all those in-the-moment decisions about whether something is a learning opportunity. To reach our long-term goals, we have to

improve at sorting out when the unfolding future has something to teach us, when to close the feedback loop.

And the first step to doing this well is in recognizing that things sometimes happen because of the other form of uncertainty: luck.

Luck vs. skill: fielding outcomes


The way our lives turn out is the result of two things: the influence of skill and the influence of luck. For the purposes of this discussion, any outcome that is the result of our decision-making is in the skill category. If making the same decision again would predictably result in the same outcome, or if changing the decision would predictably result in a different outcome, then the outcome following that decision was due to skill. The quality of our decision-making was the main influence over how things turned out. If, however, an outcome occurs because of things that we can't control (like the actions of others, the weather, or our genes), the result would be due to luck. If our decisions didn't have much impact on the way things turned out, then luck would be the main influence.*

When a golfer hits a tee shot, where the ball lands is the result of the influence of skill and luck, whether it is a first-time golfer or Rory McIlroy. The elements of skill, those things directly in the golfer's control that influence the outcome, include club choice, setup, and all the detailed mechanics of the golf swing. Elements of luck include a sudden gust of wind, somebody yelling their name as they swing, the ball landing in a divot or hitting a sprinkler head, the age of the golfer, the golfer's genes, and the opportunities they received (or didn't receive) up to the moment of the shot.

An outcome like losing weight could be the direct result of a change in diet or increased exercise (skill), or a sudden change in our metabolism or a famine (luck). We could get in a car crash because we didn't stop at a red light (skill) or because another driver ran a red light (luck). A student could do poorly on a test because they didn't study (skill) or because the teacher is mean (luck). I can lose a hand of poker because I made poor decisions, applying the skill elements of the game poorly, or because the other player got lucky.

Chalk up an outcome to skill, and we take credit for the result. Chalk up an outcome to luck, and it wasn't in our control. For any outcome, we are faced with this initial sorting decision. That decision is a bet on whether the outcome belongs in the "luck" bucket or the "skill" bucket. This is where Nick the Greek went wrong.

We can update the learning loop to represent this like so:



Think about this like we are an outfielder catching a fly ball with runners on base. Fielders have to make in-the-moment game decisions about where to throw the ball: hit the cutoff man, throw behind a base runner, throw out an advancing base runner. Where the outfielder throws after fielding the ball is a bet.

We make similar bets about where to "throw" an outcome: into the "skill bucket" (in our control) or the "luck bucket" (outside of our control). This initial fielding of outcomes, if done well, allows us to focus on experiences that have something to teach us (skill) and ignore those that don't (luck). Get this right and, with experience, we get closer to whatever "-ER" we are striving for: better, smarter, healthier, happier, wealthier, etc.

It is hard to get this right. Absent omniscience, it is difficult to tell why anything happened the way it did. The bet on whether to field outcomes into the luck or skill bucket is difficult to execute because of ambiguity.

Working backward is hard: the SnackWell's Phenomenon

In the nineties, millions of people jumped on the SnackWell's bandwagon. Nabisco developed these devil's food cookies as a leading product to take advantage of the now-discredited belief that fat, not sugar, makes you fat. Foods made with less fat were, at the time, considered healthier. With the blessing of the U.S. government, companies swapped in sugar for fat as a flavoring ingredient. SnackWell's came in a green package, the color associated with "low fat" and, therefore, "healthy"—like spinach!

For all those people trying to lose weight or make healthier snacking choices, SnackWell's were a delicious godsend. SnackWell's eaters bet their health on substituting these cookies for other types of snacks like, say, cashews, which are high in fat. You could ingest sugar-laden SnackWell's by the box, because sugar wasn't the enemy. Fat was the enemy, and the packaging screamed "LOW FAT!"

Of course, we know now that obesity rose significantly during the low-fat craze. (Michael Pollan used the phrase "SnackWell's Phenomenon" in describing people increasing their consumption of something that has less of a bad ingredient.) As those SnackWell's eaters gained weight, it wasn't easy for them to figure out why. Should the weight gain be fielded into the skill bucket, used as feedback that their belief about the health value of SnackWell's was inaccurate? Or was the weight gain due to bad luck, like a slow metabolism or something else that wasn't their fault or at least didn't have to do with their choice to eat SnackWell's? If the weight gain got fielded into the luck bucket, it wouldn't be a signal to alter the choice to eat SnackWell's.

Looking back *now*, it seems obvious how the weight gain should have been fielded. But it is only obvious once you *know* that SnackWell's are an unhealthy choice. We have the benefit of twenty years of new research, more and better-quality information about what causes weight gain. The folks on the low-fat bandwagon had only their weight gain to learn from. The cards remained concealed.

Working backward from the way things turn out isn't easy. We can get to the same health outcome (weight gain) by different routes. One person might

choose SnackWell's; another might choose Oreos (also a Nabisco product, developed by the same person who invented SnackWell's); a third might choose lentils and kale. If all three people gain weight, how can any of them figure it out for sure?

Outcomes don't tell us what's our fault and what isn't, what we should take credit for and what we shouldn't. Unlike in chess, we can't simply work backward from the quality of the outcome to determine the quality of our beliefs or decisions. This makes learning from outcomes a pretty haphazard process. A negative outcome could be a signal to go in and examine our decision-making. That outcome could also be due to bad luck, unrelated to our decision, in which case treating that outcome as a signal to change future decisions would be a mistake. A good outcome could signal that we made a good decision. It could also mean that we got lucky, in which case we would be making a mistake to use that outcome as a signal to repeat that decision in the future.

When Nick the Greek won with a seven and a deuce, he fielded that outcome into the skill bucket, taking credit for his brilliant strategy. When he lost with that hand—a much more common occurrence—he wrote it off as bad luck. His fielding error meant he never questioned his beliefs, no matter how much he lost. We're all like Nick the Greek sometimes. Uncertainty—luck and hidden information—gave him the leeway to make fielding errors about why he was losing. We all face uncertainty. And we all make fielding errors.

Rats get tripped up by uncertainty in a way that should appear very familiar to us. Classical stimulus-response experiments have shown that the introduction of uncertainty drastically slows learning. When rats are trained on a *fixed* reward schedule (for example, a pellet for every tenth press of a lever), they learn pretty fast to press that lever for food. If you withdraw the reward, the lever-pressing behavior is quickly extinguished. The rats figure out that no more food is on its way.

But when you reward the rats on a variable or intermittent reinforcement schedule (a pellet that comes *on average* every tenth lever press), that introduces uncertainty. The average number of lever presses for the reward is the same, but the rat could get a reward on the next press or not for thirty presses. In other words, the rats are rewarded the way humans usually are: having no way to know with certainty what will happen on the next try. When

you withdraw the reward from those rats, the lever-pressing behavior extinguishes only after a very long time of fruitless lever pushing, sometimes thousands of tries.

We might imagine the rats thinking, “I bet the next lever press will get me a pellet. . . . I’ve just been getting unlucky . . . I’m due.” Actually, we don’t even have to imagine this. We can hear it if we listen to what people say while they play slot machines. Slot machines operate on a variable-payoff system. It’s no wonder that, despite those machines being among the worst bets in the casino, the banks of slots in a casino are packed. In the end, our rat brains dominate.

If this all doesn’t seem difficult enough, outcomes are rarely all skill or all luck. Even when we make the most egregious mistakes and get appropriately negative outcomes, luck plays a role. For every drunk driver who swerves into a ditch and flips his car, there are several who swerve harmlessly across multilane highways. It might feel like the drunk driver in the ditch deserved that outcome, but the luck of the road conditions and presence or absence of other drivers also played a role. When we do everything right, like drive through a green light perfectly sober and live to tell the tale, there is also an element of luck. No one else simultaneously ran a red light and hit us. There *wasn’t* a patch of ice on the road to make us lose control of our vehicle. We *didn’t* run over a piece of debris and blow a tire.

When we field our outcomes as the future unfolds, we always run into this problem: the way things turn out could be the result of our decisions, luck, or some combination of the two. Just as we are almost never 100% wrong or right, outcomes are almost never 100% due to luck or skill. Learning from experience doesn’t offer us the orderliness of chess or, for that matter, folding and sorting laundry. Getting insight into the way uncertainty trips us up, whether the errors we make are patterned (hint: they are) and what motivates those errors, should give us clues for figuring out achievable strategies to calibrate the bets we make on our outcomes.

“If it weren’t for luck, I’d win every one”

Just as with motivated reasoning, our fielding errors aren’t random. They are, borrowing from psychologist and behavioral economist Dan Ariely,^{*} “predictably irrational.” The way we field outcomes is predictably patterned: we take credit for the good stuff and blame the bad stuff on luck so it won’t be our fault. The result is that we don’t learn from experience well.

“Self-serving bias” is the term for this pattern of fielding outcomes. Psychologist Fritz Heider was a pioneer in studying how people make luck and skill attributions about the results of their behavior. He said we study our outcomes like scientists, but like “naïve scientists.” When we figure out why something happened, we look for a plausible reason, but one that also fits our wishes. Heider said, “It is usually a reason that flatters us, puts us in a good light, and it is imbued with an added potency by the attribution.”

Our capacity for self-deception has few boundaries. Look at the reasons people give for their accidents on actual auto insurance forms: “I collided with a stationary truck coming the other way.” “A pedestrian hit me and went under my car.” “The guy was all over the road. I had to swerve a number of times before I hit him.” “An invisible car came out of nowhere, struck my car, and vanished.” “The pedestrian had no idea which direction to run, so I ran over him.” “The telephone pole was approaching. I was attempting to swerve out of its way when it struck my car.”^{*}

Stanford law professor and social psychologist Robert MacCoun studied accounts of auto accidents and found that in 75% of accounts, the victims blamed someone else for their injuries. In multiple-vehicle accidents, 91% of drivers blamed someone else. Most remarkably, MacCoun found that in *single*-vehicle accidents, 37% of drivers still found a way to pin the blame on someone else.

We can’t write this off to lack of self-knowledge by a few bad drivers. John von Neumann, considered a terror on the roads of Princeton, New Jersey, once offered this explanation after smashing his car: “I was proceeding down the road. The trees on the right were passing me in orderly fashion at 60 MPH. Suddenly, one of them stepped out in my path. Boom!” Et tu, JvN? Et tu?

This predictable fielding error is probably the single most significant problem for poker players. I watched it firsthand with Nick the Greek at the Crystal Lounge. When he lost with seven-deuce, it was because he got unlucky that time. When he won playing those cards, it was because his plan of “surprise attack” was so brilliant. Off-loading the losses to luck and onboarding the wins to skill meant he persisted in overestimating the likelihood of winning with seven-deuce. He kept betting on a losing future.

And this is not confined to Nick the Greeks of the small-time Billings variety. Phil Hellmuth, the biggest winner in World Series of Poker history (fourteen championship bracelets and counting), famously fell prey to this fielding error. After getting eliminated from a televised poker tournament, Hellmuth said on camera to ESPN, “If it weren’t for luck, I’d win every one.” This line has become legend in the poker world. (It was even the basis for a song, “I’d Win Everytime [If It Wasn’t for Luck]” in *All In: The Poker Musical*, a show based on Phil’s life.) When the ESPN episode aired, there was a collective gasp in the poker community. After all, Phil is saying that if the luck element were eliminated from poker—if he were playing chess—his poker skills are so superior that he would win every tournament he entered. So, clearly, any negative outcomes are due to luck, and any positive outcomes are the result of his superior skill.

Poker players may have gasped, but the only difference between Phil and everyone else is that he said this *out loud and on television*. Most of us just have the sense to keep the sentiment to ourselves, especially when the cameras and mics are on. But, trust me, we are all vulnerable to the exact same thinking.

I, certainly, am not exempt. When I was playing poker, I did my share of taking credit for winning and complaining about bad luck when I lost. It’s a fundamental urge. I’ve been conscious of that tendency in all areas of my life. Remember, we can know it’s a visual illusion, but that doesn’t keep us from still seeing it.

Self-serving bias has immediate and obvious consequences for our ability to learn from experience.* Blaming the bulk of our bad outcomes on luck means we miss opportunities to examine our decisions to see where we can do better. Taking credit for the good stuff means we will often reinforce decisions that shouldn’t be reinforced and miss opportunities to see where we could have done better. To be sure, some of the bad stuff that happens is

mainly due to luck. And some of the good stuff that happens is mainly due to skill. I just know that's not true all the time. 100% of our bad outcomes aren't because we got unlucky and 100% of our good outcomes aren't because we are so awesome. Yet that is how we process the future as it unfolds.

The predictable pattern of blaming the bad stuff on the world and taking credit for the good stuff is by no means limited to poker or car accidents. It's everywhere.

When Chris Christie participated in a debate before the Republican presidential primary in Iowa in early 2016, he played the part of behavioral psychologist in his attack on Hillary Clinton's response to the tragic outcome in Benghazi: "She refuses to be held accountable for anything that goes wrong. If it had gone right, believe me, she would have been running around to be able to take credit for it." Whether or not the accusation is correct, Christie certainly got the human tendency right: we take credit for good things and deflect blame for bad things. Ironically, only minutes earlier, he offered a pretty competitive example of the bias in himself. When asked by the moderator if the GOP should take a chance on nominating him in light of Bridgegate, he answered, "Sure, because there's been three different investigations that have proven that I knew nothing." Then he added, "And let me tell you something else. I inherited a state in New Jersey that was downtrodden and beaten by liberal Democratic policies, high taxes, high regulation. And this year, in 2015, New Jersey had the best year of job growth that our state has ever had in the last fifteen years. That's because we've put conservative policies in place."

That's a pretty fast pivot from "that bad outcome isn't my fault" to "and let me tell you the good outcome I can take credit for."

I described this pattern during an address at a meeting of the International Academy of Trial Lawyers (IATL). One lawyer in the audience rushed to tell me after the speech about a senior partner he trained under when he was first out of law school. "You won't believe how on-point this is, Annie. I assisted this partner at several trials and at the end of each day, he would analyze the witness testimony the same way. If the witness helped our case, he would say, 'You see how well I prepared that witness? When you know how to prepare a witness, you get the results you want.' If the witness hurt our case, he would tell me, 'That guy refused to listen to me.' It never varied."

I'm betting any parent of a school-aged child knows this. On occasions when my kids have done poorly on a test, it seems to have never been because they didn't study. "The teacher doesn't like me. Everybody did poorly. The teacher put material on the test we didn't cover in class. You can ask anyone!"

Self-serving bias is a deeply embedded and robust thinking pattern. Understanding why this pattern emerges is the first step to developing practical strategies to improve our ability to learn from experience. These strategies encourage us to be more rational in the way we field outcomes, fostering open-mindedness in considering all the possible causes of an outcome, not just the ones that flatter us.

All-or-nothing thinking rears its head again

Black-and-white thinking, uncolored by the reality of uncertainty, is a driver of both motivated reasoning and self-serving bias. If our only options are being 100% right or 100% wrong, with nothing in between, then information that potentially contradicts a belief requires a total downgrade, from right all the way to wrong. There is no "somewhat less sure" option in an all-or-nothing world, so we ignore or discredit the information to hold steadfast in our belief.

Both of these biases cause us to see our outcomes through the equivalent of a funhouse mirror. The reflection distorts reality, maximizing the appearance of our skill in good outcomes. For bad outcomes, it makes skill all but disappear, while luck looms giant.

Just as with motivated reasoning, self-serving bias arises from our drive to create a positive self-narrative. In that narrative, taking credit for something good is the same as saying we made the right decision. And being right feels good. Likewise, thinking that something bad was our fault means we made a wrong decision, and being wrong feels bad. When our self-image is at stake, we treat our fielding decisions as 100% or 0%: right versus wrong, skill versus luck, our responsibility versus outside our control. There are no shades of grey.

Fielding outcomes with the goal of promoting our self-narrative and doing it in an all-or-nothing fashion alters our ability to make smart bets about why the future unfolded in a particular way. Learning from experience is difficult—and sometimes impossible—with this kind of biased, broad-brush thinking. Outcomes are rarely the result of our decision quality alone or chance alone, and outcome quality is not a perfect indicator of the influence of luck or skill. When it comes to self-serving bias, we act as if our good outcomes are perfectly correlated to good skill and our bad outcomes are perfectly correlated to bad luck.* Whether it is a poker hand, an auto accident, a football call, a trial outcome, or a business success, there are elements of luck and skill in virtually any outcome.

That the motivation to update our self-image in a positive way underlies self-serving bias gives us an idea of where we might look for solutions to overcome the bias. Maybe we could stop clinging to ego, giving up on that need to have a positive narrative of our lives. Maybe we could still drive a positive narrative but, instead of updating through credit and blame, we could get off on striving to be more objective and open-minded in assessing the influence of luck and skill on our outcomes. Maybe we could put in the time and hard work to retrain the way we process results, moving toward getting our positive self-image updates from accurate fielding, from truthseeking.

Or maybe we could detour around the obstacles altogether by finding a work-around that doesn't require us to address self-serving bias at all.

People watching

One could conceivably argue that maybe self-serving bias isn't such a big deal because we can learn from other people's experience. Maybe the solution that has evolved is to compensate for the obstacles in learning from our own experience by watching other people do stuff. There are more than seven billion other people on the planet who do stuff all the time. As Yogi Berra said, "You can observe a lot by watching."

Watching is an established learning method. There is an entire industry devoted to collecting other people's outcomes. When you read the *Harvard Business Review* or any kind of business or management case study, you're

trying to learn from others. An important element of medical education is watching doctors perform medical procedures or other caregivers do their jobs up close. They watch, then they assist . . . and then, hopefully, they've learned. Who would want a surgeon who says, "This will be the first time I've seen inside a live human body"? It's possible we could follow that example, going to school on the experiences of the people around us.

In poker, the bulk of what goes on is watching. An experienced player will choose to play only about 20% of the hands they are dealt, forfeiting the other 80% of the hands before even getting past the first round of betting. That means about 80% of the time is spent just watching other people play. Even if a poker player doesn't learn all that efficiently from the outcomes of the hands they play themselves, there is still a whole lot to be learned from watching what happens to the other players in the game. After all, there is four times as much watching everyone else as there is playing a hand yourself.

Not only are all those other people's outcomes plentiful, they are also *free* (aside from any ante). When a poker player chooses to play a hand, they are putting their own money at risk. When a poker player is just watching the game, they get to sit back while other people put money at risk. That's an opportunity to learn at no extra cost.

When any of us makes decisions in life away from the poker table, we always have something at risk: money, time, health, happiness, etc. When it's someone else's decision, we don't have to pay to learn. They do. There's a lot of free information out there.

Unfortunately, learning from watching others is just as fraught with bias. Just as there is a pattern in the way we field our own outcomes, we field the outcomes of our peers predictably. We use the same black-and-white thinking as with our own outcomes, but now we flip the script. Where we blame our own bad outcomes on bad luck, when it comes to our peers, bad outcomes are clearly their fault. While our own good outcomes are due to our awesome decision-making, when it comes to other people, good outcomes are because they got lucky. As artist and writer Jean Cocteau said, "We must believe in luck. For how else can we explain the success of those we don't like?"

When it comes to watching the bad outcomes of other people, we load the blame on them, quickly and heavily. One of the most famous moments in

baseball history, in which 40,000 people at a baseball stadium and tens of millions around the world instantly blamed an otherwise typical fan for keeping the Chicago Cubs out of the World Series, demonstrates this. The incident is known as the Bartman play.

In 2003, the Chicago Cubs were one game from reaching their first World Series since 1945. They led three games to two in the series against the Florida Marlins, and led Game Six with one out in the top of the eighth inning. A Marlins hitter lifted a foul fly toward the left-field stands. At the base of the sloping wall separating spectators from the field, Cubs left fielder Moises Alou reached his glove over his head for the ball as several spectators on the other side of the wall also reached for it. One of the 40,000 spectators at Wrigley Field, Steve Bartman, deflected the ball, which then bounced off the railing, landing at the feet of another spectator. Alou stormed away, expressing anger at the fan who kept him from attempting to catch the ball.

The Cubs had a 3–0 lead at the time Bartman, those folks near him, and Moises Alou reached for the foul ball. Had Alou caught it, the Cubs would have been four outs from reaching the World Series. Here's how the future unfolded after Bartman touched the ball: The Cubs lost the game, and Game Seven too, failing again to reach the World Series. Bartman got the blame, first from the 40,000 people present (who pointed and chanted “A**hole!,” threw beer and garbage at him, and yelled death threats), then from the millions of Cubs fans, not just on sports and news shows replaying the incident at the time but for more than a decade. One person who actually assaulted Bartman while he was surrounded by stadium security said, “I wanted to expose him for ruining what could have been a once-in-a-lifetime experience.”

Steve Bartman had a bad outcome. He reached for the ball, and the Cubs eventually lost. Was that due to his bad decision-making or bad luck? To be sure, he made the decision to reach for the ball, so there was some skill in that. He also, however, had an overwhelming amount of bad luck. Almost uniformly, people discounted that bad luck, blaming Bartman for the loss of that game and the series itself.

Alex Gibney's ESPN documentary on the Bartman play, *Catching Hell*, displays the double standard, showing the replay from numerous angles and interviewing spectators and members of the media from the scene. Gibney noted (and the footage clearly showed), “There are a lot of people who went

for that ball.” One fan, who lunged right next to Bartman for the ball, couldn’t deny that he had reached for the ball, the same as Bartman. “I went for the ball. There’s no way . . . I obviously went for the ball.” Yet, he tried to claim that, unlike Bartman, he never would have interfered with the play: “Once I saw [Moises Alou’s] glove, I had no interest in the ball.” That fan, at once, took credit for his good outcome in not having touched the ball and placed blame squarely on Bartman’s decision-making rather than bad luck.

While everybody in the vicinity behaved the same way, Bartman was the unlucky one who touched the baseball. But the fans didn’t see it as bad luck. They saw it as his fault. Worse yet, none of the subsequent things that happened in the game—all clearly outside Bartman’s control—mitigated his responsibility. Remember that after that ball landed in the stands, the Cubs were in the same position they were in before the Bartman play. They still had the Marlins down to their last five outs, with a 3–0 lead in the game, their ace pitcher hurling a shutout, and a 3–2 lead in a best-of-seven series. That batter who hit the foul ball was still at the plate, with a 3–2 count. The Marlins went on to score *eight runs* in the inning, seven of them after (two batters later) the Cubs shortstop, Alex Gonzalez, bobbled an inning-ending double-play ball for an error.

Bartman had a lot of bad luck to end up with that result, much of which had to do with the play of the team, over which Bartman clearly had no control. The fans, nevertheless, laid all the blame on Bartman rather than on, say, Gonzalez. Nearby fans yelled, “Rot in hell! Everyone in Chicago hates you! You suck!” As he walked through the concourse, people yelled, “We’re gonna kill you! Go to prison!” “Put a twelve-gauge in his mouth and pull the trigger!”

It would be nice if a story line developed that gave Steve Bartman credit for the Cubs winning the World Series in 2016. After all, Steve Bartman was a key player in the chain of events responsible for a moribund franchise hiring baseball-turnaround specialist Theo Epstein as president of baseball operations and Joe Maddon as manager. Not surprisingly, that narrative hasn’t gotten any traction.*

We see this pattern of blaming others for bad outcomes and failing to give them credit for good ones all over the place. When someone else at work gets a promotion instead of us, do we admit they worked harder than and deserved it more than we did? No, it was because they schmoozed the boss.

If someone does better on a test at school, it was because the teacher likes them more. If someone explains the circumstances of a car accident and how it wasn't their fault, we roll our eyes. We assume the cause was their bad driving.

When I started in poker, I followed the same pattern (and still fight the urge in every area of my life to this day). Even though I was quick to take credit for my own successes and blame bad luck for my losses, I flipped this when evaluating other players. I didn't give other players enough credit for winning (or, viewed another way, give other players credit for my losing) and I was quick to blame their losses on their poor play.

When I first started, my brother gave me a list of cards that were good to play, written on a napkin while we were eating in a coffee shop at Binion's Horseshoe Casino. I clutched this napkin like I was Moses clutching the Ten Commandments. When I saw people win with hands that were not on my brother's list, I dismissed it as good luck since they clearly did not know how to play. I was so closed-minded to the thought that they might deserve credit for winning that I didn't even bother to describe these hands to my brother to ask him if there might be a good reason they would play a hand off the list.

As my understanding of the game grew over time, I realized the hands on that list weren't the only hands you could ever play. For one thing, only playing hands on the list meant you would never bluff. My brother gave me that list to keep me, a total novice, out of trouble. He gave me that list because the hands he was telling me to play would limit the mistakes a total beginner might make. I didn't know that, depending on all sorts of factors, sometimes there are hands that wouldn't be on the list that would be perfectly okay to play.

Even though I had access to the list writer, I never asked my brother why these guys were playing these off-list cards. My biased assessment of why they were winning slowed my learning down considerably. I missed out on a lot of opportunities to make money because I dismissed other players as lucky when I might have been learning from watching them. To be sure, some of those people shouldn't have been playing those hands and were actually playing poorly. But, as I figured out almost a year into playing, not all of them.

The systematic errors in the way we field the outcomes of our peers comes at a real cost. It doesn't just come at the cost of reaching our goals but

also at the cost of compassion for others.

Other people's outcomes reflect on us

We all want to feel good about ourselves in the moment, even if it's at the expense of our long-term goals. Just as with motivated reasoning and self-serving bias, blaming others for their bad results and failing to give them credit for their good ones is under the influence of ego. Taking credit for a win lifts our personal narrative. So too does knocking down a peer by finding them at fault for a loss. That's *schadenfreude*: deriving pleasure from someone else's misfortune. *Schadenfreude* is basically the opposite of compassion.

Ideally, our happiness would depend on how things turn out for us regardless of how things turn out for anyone else. Yet, on a fundamental level, fielding someone's bad outcome as their fault feels good to us. On a fundamental level, fielding someone's good outcome as luck helps our narrative along.

This outcome fielding follows a logical pattern in zero-sum games like poker. When I am competing head-to-head in a poker hand, I *must* follow this fielding pattern to square my self-serving interpretation of my own outcomes with the outcomes of my opponent. If I win a hand in poker, my opponent loses. If I lose a hand in poker, my opponent wins. Wins and losses are symmetrical. If I field my win as having to do with my skillful play, then my opponent in the hand must have lost because of their less skillful play. Likewise, if I field my loss as having to do with luck, then my opponent must have won due to luck as well. Any other interpretation would create cognitive dissonance.

Thinking about it this way, we see that the way we field other people's outcomes is just part of self-serving bias. Viewed through this lens, the pattern begins to make sense.

But this comparison of our results to others isn't confined to zero-sum games where one player directly loses to the other (or where one lawyer loses to opposing counsel, or where one salesperson loses a sale to a competitor, etc.). We are really in competition for resources with everyone. Our genes are

competitive. As Richard Dawkins points out, natural selection proceeds by competition among the phenotypes of genes so we literally evolved to compete, a drive that allowed our species to survive. Engaging the world through the lens of competition is deeply embedded in our animal brains. It's not enough to boost our self-image solely by our own successes. If someone we view as a peer is winning, we feel like we're losing by comparison. We benchmark ourselves to them. If their kids are doing better in school than ours, what are we doing wrong with our kids? If their company is in the news because it is about to go public, what's wrong with us that we're just inching forward in our work?

We think we know the ingredients for happiness. Sonja Lyubomirsky, a psychology professor at the University of California, Riverside, and popular author on the subject of happiness, summarized several reviews of the literature on the elements we commonly consider: "a comfortable income, robust health, a supportive marriage, and lack of tragedy or trauma." Lyubomirsky noted, however, that "the general conclusion from almost a century of research on the determinants of well-being is that objective circumstances, demographic variables, and life events are correlated with happiness less strongly than intuition and everyday experience tell us they ought to be. By several estimates, all of these variables put together account for no more than 8% to 15% of the variance in happiness." What accounts for most of the variance in happiness is how we're doing *comparatively*. (The breadth and depth of all that research on happiness and its implications is important, but it's beyond what we need to understand our issue with sorting others' outcomes. I encourage you to read Lyubomirsky's work on the subject, Daniel Gilbert's *Stumbling on Happiness*, and Jonathan Haidt's *The Happiness Hypothesis*, cited in the Selected Bibliography and Recommendations for Further Reading.)

A consistent example of how we price our own happiness relative to others comes from a version of the party game "Would You Rather . . . ?" When you ask people if they would rather earn \$70,000 in 1900 or \$70,000 now, a significant number choose 1900. True, the average yearly income in 1900 was about \$450. So we'd be doing phenomenally well compared to our peers from 1900. But no amount of money in 1900 could buy Novocain or antibiotics or a refrigerator or air-conditioning or a powerful computer we could hold in one hand. About the only thing \$70,000 bought in 1900 that it

couldn't buy today was the opportunity to soar above most everyone else. We'd rather lap the field in 1900 with an average life expectancy of only forty-seven years than sit in the middle of the pack now with an average life expectancy of over seventy-six years (and a computer in our palm).

A lot of the way we feel about ourselves comes from how we think we compare with others. This robust and pervasive habit of mind impedes learning. Luckily, habits can be changed, whether the habit is biting your nails or decrying your terrible luck when you lose. By shifting what it is that makes us feel good about ourselves, we can move toward a more rational fielding of outcomes and a more compassionate view of others. We can learn better and be more open-minded if we work toward a positive narrative driven by engagement in truthseeking and striving toward accuracy and objectivity: giving others credit when it's due, admitting when our decisions could have been better, and acknowledging that almost nothing is black and white.

Reshaping habit

Phil Ivey is one of those guys who can easily admit when he could have done better. Ivey is one of the world's best poker players, a player almost universally admired by other professional poker players for his exceptional skill and confidence in his game. Starting in his early twenties, he built a reputation as a top cash-game player, a top tournament player, a top heads-up player, a top mixed-game player—a top player in every form and format of poker. In a profession where, as I've explained, most people are awash in self-serving bias, Phil Ivey is an exception.

In 2004, my brother provided televised final-table commentary for a tournament in which Phil Ivey smoked a star-studded final table. After his win, the two of them went to a restaurant for dinner, during which Ivey deconstructed every potential playing error he thought he might have made on the way to victory, asking my brother's opinion about each strategic decision. A more run-of-the-mill player might have spent the time talking about how great they played, relishing the victory. Not Ivey. For him, the opportunity to learn from his mistakes was much more important than

treating that dinner as a self-satisfying celebration. He earned a half-million dollars and won a lengthy poker tournament over world-class competition, but all he wanted to do was discuss with a fellow pro where he might have made better decisions.

I heard an identical story secondhand about Ivey at another otherwise celebratory dinner following one of his now ten World Series of Poker victories. Again, from what I understand, he spent the evening discussing in intricate detail with some other pros the points in hands where he could have made better decisions. Phil Ivey, clearly, has different habits than most poker players—and most people in any endeavor—in how he fields his outcomes.

Habits operate in a neurological loop consisting of three parts: the cue, the routine, and the reward. A habit could involve eating cookies: the cue might be hunger, the routine going to the pantry and grabbing a cookie, and the reward a sugar high. Or, in poker, the cue might be winning a hand, the routine taking credit for it, the reward a boost to our ego. Charles Duhigg, in *The Power of Habit*, offers the golden rule of habit change—that the best way to deal with a habit is to respect the habit loop: “To change a habit, you must keep the old cue, and deliver the old reward, but insert a new routine.”

When we have a good outcome, it cues the routine of crediting the result to our awesome decision-making, delivering the reward of a positive update to our self-narrative. A bad outcome cues the routine of off-loading responsibility for the result, delivering the reward of avoiding a negative self-narrative update. With the same cues, we flip the routine for the outcomes of peers, but the reward is the same—feeling good about ourselves.

The good news is that we can work to change this habit of mind by substituting what makes us feel good. The golden rule of habit change says we don’t have to give up the reward of a positive update to our narrative, nor should we. Duhigg recognizes that respecting the habit loop means respecting the way our brain is built.

Our brain is built to seek positive self-image updates. It is also built to view ourselves in competition with our peers. We can’t install new hardware. Working with the way our brains are built in reshaping habit has a higher chance of success than working against it. Better to change the part that is more plastic: the routine of what gives us the good feeling in our narrative and the features by which we compare ourselves to others.

At least as far back as Pavlov, behavioral researchers* have recognized the power of substitution in physiological loops. In his famous experiments, his colleague noticed that dogs salivated when they were about to be fed. Because they associated a particular technician with food, the presence of the technician triggered the dogs' salivation response. Pavlov discovered the dogs could learn to associate just about any stimulus with food, including his famous bell, triggering the salivary response.

We can work to change the bell we ring, substituting what makes us salivate. We can work to get the reward of feeling good about ourselves from being a good credit-giver, a good mistake-admitter, a good finder-of-mistakes-in-good-outcomes, a good learner, and (as a result) a good decision-maker. Instead of feeling bad when we have to admit a mistake, what if the bad feeling came from the thought that we might be missing a learning opportunity just to avoid blame? Or that we might be basking in the credit of a good result instead of, like Phil Ivey, recognizing where we could have done better? If we work toward that, we can transform the unproductive habits of mind of self-serving bias and motivated reasoning into productive ones. If we put in the work to practice this routine, we can field more of our outcomes in an open-minded, more objective way, motivated by accuracy and truthseeking to drive learning. The habit of mind will change, and our decision-making will better align with executing on our long-term goals.

There are people who, like Phil Ivey, have substituted the routine of truthseeking for the outcome-oriented instinct to focus on seeking credit and avoiding blame. When we look at the people performing at the highest level of their chosen field, we find that the self-serving bias that interferes with learning often recedes and even disappears. The people with the most legitimate claim to a bulletproof self-narrative have developed habits around accurate self-critique.

In sports, the athletes at the top of the game look at outcomes to spur further improvement. American soccer great Mia Hamm said, "Many people say I'm the best women's soccer player in the world. I don't think so. And because of that, someday I just might be." Such quotes can be discounted as a polite way of dealing with the media. There are plenty of contrary examples burned in our consciousness, like John McEnroe arguing line calls or golf pros who have turned missed putts into a ritual of staring at the offending line of the putt and tapping down phantom spike marks. Those aren't anything

more than performance tics. It's practically a ritual on the PGA Tour that if a player misses a makeable putt, he has to stare at the green like it was somehow at fault. What you don't see are practice rituals like Phil Mickelson's, in which he places ten balls in a circle, three feet from the hole. He has to sink all ten, and then repeat the process *nine more times*. Players of Phil Mickelson's caliber couldn't engage in such a demanding regimen if they actually assigned much blame to spike marks.

Changing the routine is hard and takes work. But we can leverage our natural tendency to derive some of our self-esteem by how we compare to our peers. Just as Duhigg recommends respecting the habit loop, we can also respect that we are built for competition, and that our self-narrative doesn't exist in a vacuum. Keep the reward of feeling like we are doing well compared to our peers, but change the features by which we compare ourselves: be a better credit-giver than your peers, more willing than others to admit mistakes, more willing to explore possible reasons for an outcome with an open mind, even, and especially, if that might cast you in a bad light or shine a good light on someone else. In this way we can feel that we are doing well by comparison because we are doing something unusual and hard that most people don't do. That makes us feel exceptional.

Once we start listening for it, we hear a chorus out in the world like I heard during breaks in poker tournaments: "things are going great because I'm making such good decisions"; "things went poorly because I got so unlucky." That's what the lawyer heard from his senior partner in each evening's postmortem of the trial. That's what we heard from Chris Christie in the 2016 Republican presidential debate. That's what I heard in every poker room I was ever in. There were times, and there still are, when I'm part of that chorus. Increasingly, though, I've learned to use that chorus to avoid self-serving bias rather than giving in to it. When I admitted mistakes, when I recognized the luck element in my successes, when I gave other players credit for making some good decisions, when I was eager to share a hand that I thought I played poorly because I might learn something from it, that chorus reminded me that what I was doing was hard, and that others weren't often doing it. Identifying learning opportunities that other players were missing made me feel good about myself, reinforcing my routine change.

Ideally, we wouldn't compare ourselves with others or get a good feeling when the comparison favors us. We might adopt the mindful practices of

Buddhist monks, observing the flow of inner thoughts, emotions, and bodily sensations without judging them as good or bad at all. That's a great goal, and I'm all for a regular mindful practice. It will, the research indicates, help improve quality of life and is worth pursuing. But getting all the way there is a tall order if we don't want to quit our day jobs and move to Tibet. It works against the way our brains evolved, against our competitive drive. As a parallel practice, the more practical and immediate solution is to work with what we've got, using that comparison to strengthen our focus on accuracy and truthseeking. Plus, we won't have to give up our lives and find a remote mountaintop to live on.

We need a mindset shift. We need a plan to develop a more productive habit of mind. That starts in deliberative mind and requires foresight and practice, but if it takes hold, it can become an established habit, running automatically and changing the way we reflexively think.

We can get to this mindset shift by behaving as if we have something at risk when we sort outcomes into the luck and skill buckets, because we do *have a lot at risk* on that fielding decision. Thinking in bets is a smart way to start building habits that achieve our long-term goals.

“Wanna bet?” redux

Treating outcome fielding as a bet can accomplish the mindset shift necessary to reshape habit. If someone challenged us to a meaningful bet on how we fielded an outcome, we would find ourselves quickly moving beyond self-serving bias. If we wanted to win that bet, we wouldn't reflexively field bad outcomes as all luck or good ones as all skill. (If you walked into a poker room and threw around words like “always” and “never,” you'd soon find yourself challenged to a bunch of bets. It's easy to win a bet against someone who takes extreme positions.)

Imagine getting into an accident at an intersection after losing control of your car on a patch of ice you couldn't see. Your first thought would likely be that you got unlucky. But what if you had to bet on that? Depending on the details, there would be a lot of alternatives to just plain bad luck you would now consider. Based on the weather, maybe you could have anticipated some

ice on the road. Maybe you were driving too fast for the weather conditions. Once the car started sliding, maybe you could have steered differently or maybe you pumped the brakes when you shouldn't have. Maybe you could have taken a safer route, choosing a main road that would have been salted. Maybe you should have kept the Mustang in the garage and taken the Suburban.

Some of the reasons we come up with may be easy to discount. And some may not. The key is that in explicitly recognizing that the way we field an outcome is a bet, we consider a greater number of alternative causes more seriously than we otherwise would have. That is truthseeking. This is what Phil Ivey does.

The prospect of a bet makes us examine and refine our beliefs, in this case the belief about whether luck or skill was the main influence in the way things turned out. Betting on what we believe makes us take a closer look by making explicit what is already implicit: we have a great deal at risk in assessing why anything turned out the way it did. That sure sounds like a bet worth taking seriously.

When we treat outcome fielding as a bet, it pushes us to field outcomes more objectively into the appropriate buckets because that is how bets are won. Winning feels good. Winning is a positive update to our personal narrative. Winning is a reward. With enough practice, reinforced by the reward of feeling good about ourselves, thinking of fielding outcomes as bets will become a habit of mind.

Thinking in bets triggers a more open-minded exploration of alternative hypotheses, of reasons supporting conclusions *opposite* to the routine of self-serving bias. We are more likely to explore the opposite side of an argument more often and more seriously—and that will move us closer to the truth of the matter.

Thinking in bets also triggers perspective taking, leveraging the difference between how we field our own outcomes versus others' outcomes to get closer to the objective truth. We know we tend to discount the success of our peers and place responsibility firmly on their shoulders for their failures. A good strategy for figuring out which way to bet would be to imagine if that outcome had happened to us. If a competitor closes a big sale, we know about our tendency to discount their skill. But if we imagine that we had been the one who closed the sale, we are more likely to find the things to

give them credit for, that they did well and that we can learn from. Likewise, when we close the big sale, let's spare a little of the self-congratulations and, instead, examine that great result the way we'd examine it if it happened to someone else. We'll be more likely to find the things we could have done even better and identify those factors that we had no control over. Perspective taking gets us closer to the truth because that truth generally lies in the middle of the way we field outcomes for ourselves and the way we field them for others. By taking someone else's perspective, we are more likely to land in that middle ground.

Once we start actively training ourselves in testing alternative hypotheses and perspective taking, it becomes clear that outcomes are rarely 100% luck or 100% skill. This means that when new information comes in, we have options beyond unquestioned confirmation or reversal. We can modify our beliefs along a spectrum because we know it *is* a spectrum, not a choice between opposites without middle ground.

This makes us more compassionate, both toward ourselves and others. Treating outcome fielding as bets constantly reminds us outcomes are rarely attributable to a single cause and there is almost always uncertainty in figuring out the various causes. Identifying a negative outcome doesn't have the same personal sting if you turn it into a positive by finding things to learn from it. You don't have to be on the defensive side of every negative outcome because you can recognize, in addition to things you can improve, things you did well and things outside your control. You realize that not knowing is okay.

Certainly, in exchange for losing the fear of taking blame for bad outcomes, you also lose the unadulterated high of claiming good outcomes were 100% skill. That's a trade you should take. Remember, losing feels about twice as bad as winning feels good; being wrong feels about twice as bad as being right feels good. We are in a better place when we don't have to live at the edges. Euphoria or misery, with no choices in between, is not a very self-compassionate way to live.

You also become more compassionate toward other people when you treat fielding outcomes as bets. When you look at the outcomes of others from their perspective, you have to ask yourself, "What if that had happened to me?" You come up with more compassionate assessments of other people where bad things aren't always their fault and good things aren't always luck.

You are more likely to walk in their shoes. Imagine how Bartman's life would have changed if more people worked to think this way.

The hard way

Thinking in bets is hard, especially initially. It has to start as a deliberative process, and will feel clunky, weird, and slow. Certainly, there will be times it doesn't make sense. Like if you don't get a promotion at work, you're probably going to wonder how you're supposed to feel better acknowledging that so-and-so was more deserving and that you could learn a lot from them. It takes work to avoid the temptation to blame it on the boss being a jerk who doesn't know how to evaluate talent.

That feeling is natural. I built my poker career out of these principles of learning and truthseeking, yet I still catch myself falling into the self-serving bias and motivated reasoning traps. Duhigg tells us that reshaping a habit requires time, preparation, practice, and repetition.

Look at other kinds of habit changes. If I had a habit of getting out of bed at midnight to eat a cookie, it takes work as well as will to change that habit. I have to identify the habit I want to change, figure out the routine to substitute, and practice that routine in deliberative mind until the habit is reshaped. I would need to stock apples in the house, and keep them more readily available than cookies. Then I need to actually eat the apple at midnight instead of reaching for the cookie, repeating this routine until it becomes a new habit. That takes work and willpower and time.

Despite the difficulties, striving for accuracy through probabilistic thinking is a worthwhile routine to pursue. For one thing, it won't always be so difficult. We have to start doing this with deliberation and effort, but it eventually becomes a habit of mind. Just like declaring uncertainty in your beliefs, it eventually goes from a somewhat goofy and awkward extra step to a habit integral to how you view the world around you.

To be sure, thinking in bets is not a miracle cure. Thinking in bets won't make self-serving bias disappear or motivated reasoning vanish into thin air. But it will make those things better. And a little bit better is all we need to transform our lives. If we field just a few extra outcomes more accurately, if

we catch just a few extra learning opportunities, it will make a huge difference in what we learn, when we learn, and how much we learn.

Poker's compressed version of real-world decision-making showed me how being a little better at decision-making could make a big difference. A poker game can consist of a few hundred hands. Every hand can require up to twenty decisions. If, over the course of a game, there were a hundred outcomes that provided learning opportunities and we caught ten of them, we would still be missing 90% of our chances to learn. We wouldn't have transcended the way our brains function and built ourselves a different brain, but we don't need to. If our opponents are people like Nick the Greek, they are missing almost every learning opportunity. We're obviously going to do better than they do just catching 10%. If another of our opponents is someone just like us, but who isn't working to transform their outcome-processing routines, maybe they (this prior version of ourselves) will pick up five opportunities. Again, we are missing 90%, and we are still going to clean up on an opponent who is trying to learn but doesn't know how.*

The benefits of recognizing just a few extra learning opportunities compound over time. The cumulative effect of being a little better at decision-making, like compounding interest, can have huge effects in the long run on everything that we do. When we catch that extra occasional learning opportunity, it puts us in a better position for future opportunities of the same type. Any improvement in our decision quality puts us in a better position in the future. Think of it like a ship sailing from New York to London. If the ship's navigator introduces a one-degree navigation error, it would start off as barely noticeable. Unchecked, however, the ship would veer farther and farther off course and would miss London by miles, as that one-degree miscalculation compounds mile over mile. Thinking in bets corrects your course. And even a small correction will get you more safely to your destination.

The first step is identifying the habit of mind that we want to reshape and how to reshape it. That first step is hard and takes time and effort and a lot of missteps along the way. So the second step is recognizing that it is easier to make these changes if we aren't alone in the process. Recruiting help is key to creating faster and more robust change, strengthening and training our new truthseeking routines.

CHAPTER 4

The Buddy System

“Maybe you’re the problem, do you think?”

When Lauren Conrad, star of MTV’s *The Hills*, appeared on the *Late Show with David Letterman* in October 2008, her interview took an unexpected turn. The first minute was standard talk-show banter for a twenty-two-year-old star of a successful reality series: the amount of drama in her life. Less than a minute later, Conrad asked Letterman if he was calling her an idiot.

She started off the interview by discussing her ongoing feud with her by-then-former roommate Heidi Montag and Heidi’s boyfriend Spencer Pratt. In case you’re not familiar, here is the backstory: Lauren and Heidi’s friendship ended when they came to blows at a birthday party after Lauren accused Heidi and Spencer of starting rumors that she had made a sex tape. In addition, Lauren developed friendships with Stephanie (Spencer’s sister) and Holly (Heidi’s sister), complicating the social and family encounters of everyone involved. Lauren tried without success to strengthen the relationship between her roommates, Audrina and Lo. This strained Lauren’s friendship with Audrina, who reestablished her friendship with Heidi. Brody Jenner was also in the mix, dating Lauren, questioning her date with a *Teen Vogue* model, dating someone else himself, arguing with Spencer about his friendship with Lauren, getting accused of starting the rumors about Lauren’s sex tape, etc.

All of that drama happening in Conrad’s life is what David Letterman was referring to when he interjected, “That raises the question, maybe you’re

CHAPTER 6

Adventures in Mental Time Travel

Let Marty McFly run into Marty McFly

Thanks to the success of the three *Back to the Future* movies, our go-to source on the rules of time travel is more likely to be Doc Brown than Dr. Stephen Hawking. The first rule, emphasized by the trilogy and repeated by nearly every time-travel movie since, is “Whatever you do, don’t meet up with yourself!” Doc Brown (Christopher Lloyd) explains to Marty McFly (Michael J. Fox) in *Back to the Future: Part II* (1989) that “the encounter could create a time paradox, the results of which could cause a chain reaction that would unravel the very fabric of the space-time continuum and destroy the entire universe. Granted, that’s the worst-case scenario. The destruction might, in fact, be very localized, limited to merely our own galaxy.”

“Don’t meet up with yourself” has become an unquestioned element of the “science” of time travel. In *Timecop* (1994), because “the same matter can’t occupy the same space at the same time,” Jean-Claude Van Damme’s character destroys the villain by pushing his past and future versions together. The villain turns into a liquefied blob and disappears from existence.

In real-life decision-making, when we bring our past- or future-self into the equation, the space-time continuum doesn’t unravel. Far from turning us into a liquefied blob, a visit from past or future versions of us helps present-us make better bets. When making decisions, isolating ourselves from thinking about similar decisions in the past and possible future consequences is frequently the very thing that turns us into a blob, mired by in-the-moment thinking where the scope of time is distorted. As decision-makers, we *want*

to collide with past and future versions of ourselves. Our capacity for mental time travel makes this possible. As is the case with accountability, such meetings can lead to better decisions: at the moment of the decision, accountability to our group can pop us briefly into the future to imagine the conversation about the decision we will have with that group. Running that conversation will often remind us to stay on a more rational path.

Just as we can recruit other people to be our decision buddies, we can recruit other versions of ourselves to act as our own decision buddies. We can harness the power of mental time traveling, operationalizing it, encouraging it, and figuring out ways to cause that collision of past, present, and future as much as possible. Present-us needs that help, and past-us and future-us can be the best decision buddies for the job.*

Poker players have unique decision challenges that get them thinking a lot about how to get this collision of past-, present- and future-self to occur at the moment of making a decision and executing on it. Because decisions in poker are made so quickly, players don't have the luxury of time in coordinating their rational, long-term, strategic plans with their decisions at the poker table. And all those decisions made under severe time constraints have immediate consequences expressed as an exchange of poker chips. The constant exchange of chips reminds players that there is risk in every decision. Of course, the direction in which the chips flow in the short term only loosely correlates with decision quality. You can win a hand after making bad decisions and lose a hand after making good ones. But the mere fact that chips are changing hands is a reminder that every decision has consequences—that all those execution decisions you make along the way really matter.

Away from the poker table, we don't feel or experience the consequences of most of the decisions we make right away. If we are winning or losing to a particular decision, the consequences may take time to reveal themselves. If we make a losing eating decision, like substituting SnackWell's for apples, there's no immediate outcome that lets us know there might have been a cost to that choice. If we repeat that kind of decision enough, there will be consequences, but they take time to play out. In business, if a leader ignores the ideas of an intern because "What does an intern possibly know?" it could take years for that intern to become a successful competitor before that mistake becomes obvious. If the trajectory of that business suffers because of

the poverty of new ideas, the owner of the business might never realize the effect of that attitude.

The best poker players develop practical ways to incorporate their long-term strategic goals into their in-the-moment decisions. The rest of this chapter is devoted to many of these strategies designed to recruit past- and future-us to help with all the execution decisions we have to make to reach our long-term goals. As with all the strategies in this book, we must recognize that no strategy can turn us into perfectly rational actors. In addition, we can make the best possible decisions and still not get the result we want. Improving decision quality is about *increasing* our chances of good outcomes, not *guaranteeing* them. Even when that effort makes a small difference—more rational thinking and fewer emotional decisions, translated into an increased probability of better outcomes—it can have a significant impact on how our lives turn out. Good results compound. Good processes become habits, and make possible future calibration and improvement.

Those methods involve a lot of mental time travel, and those poker players could teach Marty McFly and Doc Brown a thing or two.

Night Jerry

For all the scientific research on the battle between our immediate desires and long-term goals, a particularly succinct explanation comes from Jerry Seinfeld, on why he doesn't get enough sleep: "I stay up late at night because I'm Night Guy. Night Guy wants to stay up late. 'What about getting up after five hours of sleep?' 'That's Morning Guy's problem. That's not my problem. I'm Night Guy. I stay up as late as I want.' So you get up in the morning: you're exhausted, you're groggy. 'Oooh, I hate that Night Guy.' See, Night Guy always screws Morning Guy."

That's a good example of how we struggle in the present to take care of our future-self. Night Jerry is always going to want to stay up late and, if Morning Jerry has no say in the decision, Night Jerry will get his way regardless of what's in Jerry's longer-term best interest. When we make in-the-moment decisions (and don't ponder the past or future), we are more likely to be irrational and impulsive.*

This tendency we all have to favor our present-self at the expense of our future-self is called temporal discounting.* We are willing to take an irrationally large discount to get a reward now instead of waiting for a bigger reward later. An example of temporal discounting among adults includes a study from the military drawdown in the 1990s that led tens of thousands of military employees to choose lump-sum retirement payments at drastically discounted rates instead of guaranteed annuity payments. The men and women of the U.S. military took lump-sum payments worth \$2.5 billion, a 40% discount compared to the present value of the annuity payments they would have received. (For additional sources on temporal discounting, see the Selected Bibliography and Recommendations for Further Reading.)

When Night Jerry stays up late, it's because it benefits him *now*; he discounts the benefits that come later from going to bed. Saving for retirement is a temporal discounting problem: the gratification of spending discretionary income is immediate. Putting it away for retirement means we have to wait decades to get enjoyment from that money. We are built for temporal discounting, for using the resources that are available to us now as opposed to saving them for a future version of us that we aren't particularly in touch with in the moment of the decision. Time traveling can get us in touch with that future version of us. It can get future-us to remind present-us, "Hey, don't discount!" Or at least, "Don't discount so much!"

When we think about the past and the future, we engage deliberative mind, improving our ability to make a more rational decision. When we imagine the future, we don't just make it up out of whole cloth, inventing a future based on nothing that we have ever seen or experienced. Our vision of the future, rather, is rooted in our memories of the past. The future we imagine is a novel reassembling of our past experiences. Given that, it shouldn't be surprising that the same neural network is engaged when we imagine the future as when we remember the past. Thinking about the future is *remembering* the future, putting memories together in a creative way to imagine a possible way things might turn out. Those brain pathways include the hippocampus (a key structure for memory) and the prefrontal cortex, which controls System 2, deliberative decision-making. It is our cognitive control center.* By engaging those pathways, Night Jerry can access memories like oversleeping and missing appointments or dozing off during morning meetings that he can use to imagine how tired Morning Jerry will be

or what's going to happen to Morning Jerry's schedule when he doesn't want to get up or how his day will go when he can't pay attention.

Wouldn't it be great if Morning Jerry could travel back in time and tap Night Jerry on the shoulder to tell him to go to bed? As it turns out, there's an app for that.

From advancements in photo techniques and virtual reality, there is software that can show you a prediction of what you will look like decades into the future. If you feel, like most adults, bad about seeing your parents age, these images of future-you can be unsettling, like looking into a sadist-designed funhouse mirror. Fortunately, they've found ways to put this age-progression technology to more productive uses than just making us stare into the void of our own mortality.

Saving for retirement is a Night Jerry-versus-Morning Jerry problem. If Night Jerry isn't even thinking ahead to tomorrow morning, he's certainly not thinking ahead several decades to retirement. Retirement planning involves a series of decisions in which our present-self can act to the detriment or benefit of our future-self. When we set retirement goals, we are necessarily thinking about our future-self's goals—how much we need to save for that older version of us to live comfortably. Our spending decisions, however, don't seem to be particularly focused on what's best for our seventy-year-old-self. In fact, a quick Google search on the topic quickly reveals that our retirement savings are dangerously low. According to one study by the Center for Retirement Research at Boston College, “roughly half of today's working households will not be able to maintain their standard of living in retirement.” Depending on which estimate you read, the shortfall could be \$6.8 to \$14 *trillion*.

Several organizations and companies with an interest in encouraging retirement planning have resources that allow clients to “meet” their future-selves as they make retirement decisions. In the simplest versions of these tools, clients plug in their age, income, savings practices, and retirement goals. The apps then show the client the financial situation and lifestyle their future-self can expect, compared with the present.

Prudential Retirement, AARP, and others have versions of these apps that emphasize the consequences of retirement planning by *visually* introducing us to our future-self. Bank of America Merrill Lynch in 2012 (for web-access computing) and 2014 (for mobile devices) introduced Merrill

Edge, which includes a tool called “Face Retirement.” Customers upload a picture of themselves and get to see, according to the press release, “a life-like 3D animation of their future self, enabling them to envision every wrinkle they could encounter at retirement age—and beyond.” Night Jerry gets a glimpse of what Morning Jerry looks like without enough sleep to function.

This idea that seeing our aged future-self could help us make better allocation decisions is based, in part, on research by Jeremy Bailenson and Laura Carstensen of Stanford University’s Freeman Spogli Institute for International Studies. They used immersive virtual-reality technology in a lab setting to demonstrate how a visit from Morning Jerry will help Night Jerry make better decisions. Subjects entered a virtual-reality environment, after which they were asked to allocate \$1,000 among accounts for various purposes, one of which was a hypothetical retirement account. Subjects seeing a digital representation of their present-self in the mirror allocated on average \$73.90 to the retirement account. Other subjects looking in that mirror saw an age-progressed version of themselves. This latter group of subjects, on average, allocated \$178.10 to the retirement account. This is a startling example of how future-us can act as an effective decision buddy for present-us.

Bringing our future-self into the decision gets us started thinking about the future consequences of those in-the-moment decisions. Fundamentally, Morning Jerry and Night Jerry are living the same life, and getting Morning Jerry into Night Jerry’s face will remind him of that. Seeing our aged-self in the mirror, along with a spreadsheet showing us how future-us has to struggle to get by, is a persuasive reminder to put aside some discretionary spending money for retirement. It’s that tap on the shoulder from our future-self. “Hey, don’t forget about me. I’m going to exist and I’d like you to please take that into account.”

We’re not perfectly rational when we ponder the past or the future and engage deliberative mind, but we are more likely to make choices consistent with our long-term goals when we can get out of the moment and engage our past- and future-selves. We want Night Jerry and Morning Jerry colliding on the decision of when to get some sleep. We want all those Marty McFlys to get the additional perspective of all the other Marty McFlys. And we want our aged, wrinkly self colliding with us when we decide between spending

more money now on something like a nicer car versus saving more money for retirement.

Moving regret in front of our decisions

Philosophers agree that regret is one of the most intense emotions we feel, but they have argued about whether it is productive or useful. Nietzsche said that remorse was “adding to the first act of stupidity a second.” Thoreau, on the other hand, praised the power of regret: “Make the most of your regrets; never smother your sorrow, but tend and cherish it till it comes to have a separate and integral interest. To regret deeply is to live afresh.”

The problem isn't so much whether regret is an unproductive emotion. It's that regret occurs after the fact, instead of before. As Nietzsche points out, regret can do nothing to change what has already happened. We just wallow in remorse about something over which we no longer have any control. But if regret occurred before a decision instead of after, the experience of regret might get us to change a choice likely to result in a bad outcome. Then we could embrace Thoreau's view and harness the power of regret because it would serve a valuable purpose. It would be helpful, then, if we could get regret to do some time traveling of its own, moving before our decisions instead of after them. That way, regret might be able to keep us from making a bad bet. In addition, it wouldn't, as Nietzsche implied, rear its head later by causing us to make a remorse-fueled second mistake.

Morning Jerry regretted Night Jerry's decision to stay up late, but it was too late for him to do anything about it. When we see the retirement-savings shortfall in this country, there is no doubt that many of the future, retired versions of us are going to regret the financial allocation decisions that younger-us made, after it is too late to fix it. The age-progression imaging works to address this issue of regret occurring when it's already too late. By giving us a look at the retirement-aged version of ourselves, it gives us a chance to experience some regret that we didn't plan adequately for retirement *before* we've made inadequate plans. That was one of the purposes of my loss limit in poker. Because of the loss-limit agreement I had made with myself and my group, I ran the conversation in my head that I'd be

forced to have when I explained why I kept playing beyond my limit. It gave me a chance to regret the decision before I bought more chips.

One of our time-travel goals is to create moments like that, where we can interrupt an in-the-moment decision and take some time to consider the decision from the perspective of our past and future. We can then create a habit routine around these decision interrupts to encourage this perspective taking, asking ourselves a set of simple questions at the moment of the decision designed to get future-us and past-us involved. We can do this by imagining how future-us is likely to feel about the decision or by imagining how we might feel about the decision today if past-us had made it. The approaches are complementary; whether you choose to travel to the past or travel to the future depends solely on what approach you find most effective.

Business journalist and author Suzy Welch developed a popular tool known as 10-10-10 that has the effect of bringing future-us into more of our in-the-moment decisions. “Every 10-10-10 process starts with a question. . . . [W]hat are the consequences of each of my options in ten minutes? In ten months? In ten years?” This set of questions triggers mental time travel that cues that accountability conversation (also encouraged by a truthseeking decision group). We can build on Welch’s tool by asking the questions through the frame of the past: “How would I feel today if I had made this decision ten minutes ago? Ten months ago? Ten years ago?” Whichever frame we choose, we draw on our past experiences (including similar decisions we may have regretted) in answering the questions, recruiting into the decision those less-reactive brain pathways that control executive functioning.

In poker, because the decisions are all made in the moment and the consequences are big and immediate, routines like 10-10-10 are a survival skill. I recognized in poker that in the same way that I was not the best judge of how I was playing after losing a certain amount of money, I was also not the best judge of the quality of my poker after about six to eight hours of play. Just as we can convince ourselves we are sober enough to drive, it is easy for poker players to convince themselves that they are alert enough to keep playing after many hours of intense, intellectually taxing work. In my more rational moments, away from the tables, I knew I would be better off if I played just six to eight hours per session. When I reached that point in a session and considered continuing past that time limit, I could use a 10-10-

10-like strategy to recruit my past- and future-self: How have I felt when I kept playing in the past? How has it generally worked out? When I look back, do I feel I was playing my best? This routine of asking myself these questions helped mitigate the in-the-moment risk that, as I was losing my mental edge, I might try to convince myself that the game was so great that I had to keep playing.

Moving regret in front of a decision has numerous benefits. First, obviously, it can influence us to make a better decision. Second, it helps us treat ourselves (regardless of the actual decision) more compassionately after the fact. We can anticipate and prepare for negative outcomes. By planning ahead, we can devise a plan to respond to a negative outcome instead of just reacting to it. We can also familiarize ourselves with the likelihood of a negative outcome and how it will feel. Coming to peace with a bad outcome in advance will feel better than refusing to acknowledge it, facing it only after it has happened.

After-the-fact regret can consume us. Like all emotions, regret initially feels intense but gets better with time. Time-travel strategies can help us remember that the intensity of what we feel now will subside over time. And that helps reduce the emotion we feel in the moment, making it less likely that we will prove Nietzsche right and add a second act of stupidity to the first.

A flat tire, the ticker, and a zoom lens

Imagine you are standing on a narrow strip of concrete on the shoulder of the highway. Behind you is your car, hazard lights flashing. The rear tire on the driver's side is shredded. It is fully dark and the drizzle has turned into a cold, heavy downpour. You've called roadside assistance, twice, and both times (after long hold times) spoken with operators who've told you someone will arrive "as soon as they get there after responding to your call." You decide to change the tire yourself, only to discover you have no jack. You're soaked to the skin and cold.

How does it feel? It likely feels like the worst moment of your life. You are likely bemoaning how unlucky you are, wondering why these things

always happen to you. You are miserable and you can't imagine feeling any other way.*


That's how it feels in the moment. But if the flat tire had happened a year ago, do you think it would have an effect on your happiness today, or your overall happiness over the past year? Not likely. It likely wouldn't cause your overall happiness to tick up or down. It would probably have faded to a funny story (or a story you try to make sound funny) told at cocktail parties.

In our decision-making lives, we aren't that good at taking this kind of perspective—at accessing the past and future to get a better view of how any given moment might fit into the scope of time. It just feels how it feels in the moment and we react to it. We want to create opportunities to take the broader perspective prior to making decisions driven by the magnified feelings we have in the moment. A 10-10-10 strategy does that, getting us to imagine the decision or outcome in the perspective of the past and the future.


The flat tire isn't as awful as it seems in the moment. This kind of time-travel strategy calms down the in-the-moment emotions we have about an event, so we can get back to using the more rational part of our brain. Recruiting past-us and future-us in this way activates the neural pathways that engage the prefrontal cortex, inhibiting emotional mind and keeping events in more rational perspective. This discourages us from magnifying the present moment, blowing it out of proportion and overreacting to it.

This overestimation of the impact of any individual moment on our overall happiness is the emotional equivalent of watching the ticker in the financial world. We make a long-term stock investment because we want it to appreciate over years or decades. Yet there we are, watching a downward tick over a few minutes, consumed by imagining the worst. What's the volume? Is it heavier than usual? Better check the news stories. Better check the message boards to find out what rumors are circulating.

A stock like Berkshire Hathaway reveals why ticker watching isn't a particularly productive endeavor when you are investing for the long run. Look at this chart of Berkshire's performance since 1964:




Now zoom in on a random day in late January 2017. The upticks and downticks look large and potentially frightening. You can imagine sitting at the low point around 11:30, feeling like your losses are spiraling.



If you zoomed in on the performance of Berkshire Hathaway stock during the banking crisis, September 2008 to March 2009, you would feel

terrible most *days*:

BERKSHIRE HATHAWAY (SEPT. 2008—MARCH 2009)



Yet we know from the first chart, the big picture, that all those minute-to-minute and even day-to-day changes had little effect on the investment's general upward trajectory.

Our problem is that we're ticker watchers of our own lives. Happiness (however we individually define it) is not best measured by looking at the ticker, zooming in and magnifying moment-by-moment or day-by-day movements. We would be better off thinking about our happiness as a long-term stock holding. We would do well to view our happiness through a wide-angle lens, striving for a long, sustaining upward trend in our happiness stock, so it resembles the first Berkshire Hathaway chart.

Mental time travel makes that kind of perspective possible. We can use our past- and future-selves to pull us out of the moment and remind us when we're watching the ticker, looking at our lives through that lens on extreme zoom.

When we view these upticks and downticks under the magnification of that in-the-moment zoom lens, our emotional responses are, similarly,

amplified. Like the flat tire in the rain, we are capable of treating things that will have little effect on our long-term happiness as having significant impact. Our decision-making becomes reactive, focused on off-loading negative emotions or sustaining positive emotions from the latest change in the status quo. We can see how this can result in self-serving bias: fielding outcomes to off-load the negative emotions we feel in the moment from a bad outcome by blaming them on luck and sustaining the positive emotions from good outcomes by taking credit for them. The decisions driven by the emotions of the moment can become a self-fulfilling prophecy, degrading the quality of the bets we make, increasing the chances of bad outcomes, and making things worse.

“Yeah, but what have you done for me lately?”

Watching the ticker doesn't just magnify what has happened in the very recent past. It distorts our view of it as well. To understand the additional element of distortion, the casino is a great place to look.

Imagine that you go to a casino for an evening of blackjack with your friends. In the first half hour, you go on a winning streak and are ahead \$1,000. You keep playing because you and your friends are having such a good time. For the next hour and a half, it seems like you never win a hand. You lose back the \$1,000 and break even for the night. How are you feeling about that?

Now imagine that you lose \$1,000 in the first half hour and stick around playing with your friends because they are having a great time. In the next hour and a half you go on a winning streak that erases the early loss, and you end up breaking even for the night. How are you feeling about that?

I'm guessing you are pretty sad and morose about starting off with the big win, only to break even. In the second example, you're probably so happy that the drinks are on you. While you took a different path to get there, in both cases you didn't win or lose a dime at the end of the two hours. But in one case you are really sad about the result and the other really happy.

As they say in the infomercial world, “But, wait! There's more!”

Imagine you go up that same \$1,000 in the first half hour but now, over the next hour and a half, you can't seem to win a hand and lose \$900 back, ending the night with a \$100 win. How does that feel? Now imagine that you lost that same \$1,000 in the first half hour but then went on a winning streak to end the night down only \$100. How does that feel? Most likely, you're pretty glum about the \$100 win but still buying drinks for everyone after recovering from that terrible start to only lose \$100. So you're sad that you won \$100 and happy that you lost \$100.

The way we field outcomes is path dependent. It doesn't so much matter where we end up as how we got there. What has happened in the recent past drives our emotional response much more than how we are doing overall. That's how we can win \$100 and be sad, and lose \$100 and be happy. The zoom lens doesn't just magnify, it distorts. This is true whether we are in a casino, making investment decisions, in a relationship, or on the side of the road with a flat tire. If we got a big promotion last week and have a flat tire right now, we are cursing our lives, complaining about how unlucky we are. Our feelings are not a reaction to the average of how things are going. We feel sad if we are breaking even (or winning) on an investment that used to be valued much higher. In relationships, even small disagreements seem big in the midst of the disagreement. The problem in all these situations (and countless others) is that our in-the-moment emotions affect the quality of the decisions we make in those moments, and we are very willing to make decisions when we are not emotionally fit to do so.

Now imagine if you had gone for that night of blackjack a year ago. When you think about the outcomes as having happened in the distant past, it is likely your preference for the results reverses, landing in a more rational place. You are now happier about the \$100 win than about the \$100 loss. Once we pull ourselves out of the moment through time-traveling exercises, we can see these things in proportion to their size, free of the distortion caused by whether the ticker just moved up or down.

This is a constant challenge in poker. While the moving scoreboard has the upside of reminding players that all their decisions have consequences, there is also a downside. The scoreboard, like a stock ticker, reflects the most recent changes, creating a risk that players get caught up in ticker watching, responding emotionally and disproportionately to momentary fluctuations. Poker players think about this problem a lot.

Tilt

Surfers have more than twenty terms to describe different kinds of waves. The reason is that the type of wave, the way it breaks, the direction it's coming from, the bottom depth, etc., create differing challenges for surfers. There are *closeouts* (waves that break all at once) and *double-ups* (a type of wave created when two waves meet to form one wave) and *reforms* (a wave that will break, then die down, then break again). Non-surfers just call all of these “waves.” On rare occasions when we non-surfers need to be more specific, we just add a lot of extra words. Those extra words don't cost us much because it doesn't come up very often—maybe never. But for people involved in specialized activities, it's worth it to be able to communicate a complex concept in a single word that laypeople would need lengthy phrases to convey. Having a nuanced, precise vocabulary is what jargon is all about. It's why carpenters have at least a dozen names for different kinds of nails, and in the field of neuro-oncology, there are more than 120 types of brain and central nervous system tumors.

Because poker players are in a constant struggle to keep in-the-moment fluctuations in perspective, their jargon has a variety of terms for the concept that “bad outcomes can have an impact on your emotions that compromise your decision-making going forward so that you make emotionally charged, irrational decisions that are likely to result in more bad outcomes that will then negatively impact your decision-making going forward and so on.” The most common is *tilt*. Tilt is the poker player's worst enemy, and the word instantly communicates to other poker players that you were emotionally unhinged in your decision-making because of the way things turned out.* If you blow some recent event out of proportion and react in a drastic way, you're on tilt.

The concept of tilt comes from traditional pinball machines. To keep players from damaging the machines by lifting them to alter the course of the ball, the manufacturers placed sensors inside that disabled the machine if it was violently jostled. The flippers stopped working, the lights went off, and the word “tilt” flashed at numerous places on the layout. The origin of tilt in pinball is apt because what's going on in our brain in moments of tilt is like a shaken pinball machine. When the emotional center of the brain starts

pinging, the limbic system (specifically the amygdala) shuts down the prefrontal cortex. We light up . . . then we shut down our cognitive control center.

There are emotional and physiological signs of tilt. In poker, you can hear a poker player on tilt from several tables away. Every several hands, you hear a raised voice in an incredulous tone: “Seriously? Again?” or “I don’t know why I bother playing. I should just hand over all my money.” (Imagine the inflection of exasperation and a lot of swearing.) Along with these verbal cues, there are physiological signs of tilt. We can feel our cheeks flush and our heart race. Our respiration speeds up.

Tilt, of course, is not just limited to poker. Any kind of outcome has the potential for causing an emotional reaction. We can be tempted to make a reactive, emotional decision in a disagreement with a relationship partner, or because of bad service in a restaurant, or a comment in the workplace, or making a sale only to have it canceled, or having an idea dismissed. We’ve all had this experience in our personal and professional lives: blowing out of proportion a momentary event because of an in-the-moment emotional reaction.

By recognizing in advance these verbal and physiological signs that ticker watching is making us tilt, we can commit to develop certain habit routines at those moments. We can precommit to walk away from the situation when we feel the signs of tilt, whether it’s a fight with a spouse or child, aggravation in a work situation, or losing at a poker table. We can take some space till we calm down and get some perspective, recognizing that when we are on tilt we aren’t *decision fit*. Aphorisms like “take ten deep breaths” and “why don’t you sleep on it?” capture this desire to avoid decisions while on tilt. We can commit to asking ourselves the 10-10-10 questions or things like, “What’s happened to me in the past when I’ve felt this way?” or “Do I think it’s going to help me to be in this state while I’m making decisions?” Or we can gain perspective by asking how or whether this will have a real effect on our long-term happiness.

If you are part of a truthseeking pod, that pod can incorporate questions designed to sniff out tilt and reduce the number of decisions we execute while on tilt. We can incorporate vigilance around ticker watching when evaluating each other’s decisions, including the most obvious question: “Do you think maybe you are/were on tilt?” We can follow that with time-traveling

questions like, “Do you think this will really matter in the long run?” If we make the concept of tilt and its negative impact on decision quality part of the discussion, it creates accountability around tilt to the group. Ignoring the signals of emotional decision-making raises the prospect of having to answer for it. That, in turn, will get us positive reinforcement from the group for recognizing the signs of tilt and avoiding decision-making in that state. It also trains good habits of mind so we can run these processes on our own, acting as our own decision buddy.

At the very beginning of my poker career, I heard an aphorism from some of the legends of the profession: “It’s all just one long poker game.” That aphorism is a reminder to take the long view, especially when something big happened in the last half hour, or the previous hand—or when we get a flat tire. Once we learn specific ways to recruit past and future versions of us to remind ourselves of this, we can keep the most recent upticks and downticks in their proper perspective. When we take the long view, we’re going to think in a more rational way.

Ulysses contracts: time traveling to precommit

The most famous traveler of antiquity, the Homeric hero Odysseus, was also a mental time traveler. One of the legendary trials on his journey home involved the island of the Sirens. Sailors passing the island became so entranced by the Sirens’ song that they would steer toward the shore, crashing to their deaths on the rocky shoal around the island. Aware of the fate that befell any sailor who heard the song, Odysseus told his crew to tie his hands to the mast and fill their ears with beeswax as they approached the island. They could then steer safely, unaffected by the song they could not hear, while he would get to hear the Sirens’ song without imperiling the ship.

The plan worked perfectly. This action—past-us preventing present-us from doing something stupid—has become known as a Ulysses contract. (Most translations of Homer use the hero’s ancient Greek name, Odysseus. The time-travel strategy uses the hero’s ancient Roman name, Ulysses.)

It’s the perfect interaction between past-you, present-you, and future-you. Ulysses recognized that his future-self (along with his crew) would become

entranced by the Sirens and steer toward the rocks. So he had his crew fill their ears with wax and tie his hands to the mast, literally binding his future-self to better behavior. One of the simplest examples of this kind of contract is using a ride-sharing service when you go to a bar. A past version of you, who anticipated that you might decide irrationally about whether you are okay to drive, has bound your hands by taking the car keys out of them.

Most illustrations of Ulysses contracts, like the original, involve *raising* a barrier against irrationality. But these kinds of precommitment contracts can also be designed to lower barriers that interfere with rational action. For example, if we are trying to eat healthier, we might identify that an irrational decision point occurs when we go to the mall with someone, agree to meet them in a couple hours, and spend idle time in the food court. A barrier-*inducing* Ulysses contract could involve us not going to the mall at all or budgeting our time tightly so we have just enough time to accomplish our intended purpose. A barrier-*reducing* contract would be to precommit to carry healthy snacks in our bag, so we can increase the probability, if we're doing any idle eating, that we can make a better choice since we have drastically reduced the effort it takes to grab a healthier snack.

Ulysses contracts can come in varying levels of how much your hands are bound, ranging from physically preventing acting on a decision to just committing in advance to certain actions without any barriers save the commitment itself. Regardless of the level of binding, precommitment contracts trigger a decision-interrupt. At the moment when we consider breaking the contract, when we want to cut the binding, we are much more likely to stop and think.

When you are physically prohibited from deciding, you are interrupted in the sense that you are prevented from acting on an irrational impulse; the option simply isn't there. That's the brute-force way to do this kind of time traveling. Past-Ulysses interrupted present-Ulysses's decision by taking the decision, literally, out of his hands.

In most situations, you can't make a precommitment that's 100% tamper-proof. The hurdles aren't necessarily high, but they nevertheless create a decision-interrupt that may prompt us to do the bit of time travel necessary to reduce emotion and encourage perspective and rationality in the decision. A lawyer attending a settlement negotiation can make a precommitment, with the client or other lawyers on their team, as to the lowest amount they would

accept in a settlement (or the highest amount they would agree to pay to settle). Home buyers, understanding that in the moment they might get emotionally attached to a home, can commit in advance to their budget. Once they decide on a house they want to buy, they can decide in advance what the maximum amount they'd be willing to pay for it is so that they don't get caught up in the moment of the bidding.

Throwing out all the junk food in our house makes it impossible for midnight-us to easily, mindlessly, down a pint of ice cream. But as long as we have a car or food delivery services, that kind of food is still available somewhere. It just takes a lot more effort to get it. The same is true if we ask the waiter not to put the bread basket on the table at the restaurant. We can still, obviously, get bread, but now we have to ask the waiter to bring it. In fact, even Ulysses had to rely on his crew to ignore him if, upon hearing the Sirens' song, he signaled them to free him.

Ulysses contracts can help us in several ways to be more rational investors. When we set up an automatic allocation from our pay into a retirement account, that's a Ulysses contract. We could go through the trouble of changing the allocation, but setting it up initially gives our goal-setting, System 2–self a chance to precommit to what we know is best for our long-term future. And if we want to change the allocation, we have to take some specific steps to do so, creating a decision-interrupt.

Investment advisors do this with clients, determining in advance, as they are discussing the client's goals, the conditions under which they would buy, sell, hold, or press their positions on particular stocks. If the client later wants to make an emotional decision in the moment (involving, for example, a sudden rise or drop in the value of an investment), the advisor can remind the client of the discussion and the agreement.

In all these instances, the precommitment or predecision doesn't completely bind our hands to the mast. An emotional, reactive, irrational decision is still physically possible (though, to various degrees, more difficult). The precommitments, however, provide a stop-and-think moment before acting, triggering the potential for deliberative thought. Will that prevent an emotional, irrational decision every time? No. Will we sometimes still decide in a reflexive or mindless way? Of course. But it will happen less often.

Decision swear jar

We all know about the concept of a swear jar: if someone swears, they put a dollar in the jar. The idea behind it is that it will make people mindful about swearing and reduce how much they do it. A “decision swear jar” is a simple kind of precommitment contract that we can apply to many of the key concepts of this book. For the decision swear jar, we identify the language and thinking patterns that signal we are veering from our goal of truthseeking. When we find ourselves using certain words or succumbing to the thinking patterns we are trying to avoid because we know they are signs of irrationality, a stop-and-think moment can be created. You can think about this as a way to implement accountability.

We have discussed several patterns of irrationality in the way we lodge beliefs and field outcomes. From these, we can commit to vigilance around words, phrases, and thoughts that signal that we might be not be our most rational selves. Your list of those warning signs will be specific to you (or your family, friends, or enterprise), but here is a sample of the kinds of things that might trigger a decision-interrupt.

- Signs of the illusion of certainty: “I know,” “I’m sure,” “I knew it,” “It always happens this way,” “I’m certain of it,” “you’re 100% wrong,” “You have no idea what you’re talking about,” “There’s no way that’s true,” “0%” or “100%” or their equivalents, and other terms signaling that we’re presuming things are more certain than we know they are. This also includes stating things as absolutes, like “best” or “worst” and “always” or “never.”
- Overconfidence: similar terms to the illusion of certainty.
- Irrational outcome fielding: “I can’t believe how unlucky I got,” or the reverse, if we have some default phrase for credit taking, like “I’m at the absolute top of my game” or “I planned it perfectly.” This includes conclusions of luck, skill, blame, or credit. It includes equivalent terms for irrationally fielding the outcomes of others, like, “They totally had that coming,” “They brought it on themselves,” and “Why do they always get so lucky?”

- Any kind of moaning or complaining about bad luck just to off-load it, with no real point to the story other than to get sympathy. (An exception would be when we're in a truthseeking group and we make explicit that we're taking a momentary break to vent.)
- Generalized characterizations of people meant to dismiss their ideas: insulting, pejorative characterizations of others, like "idiot" or, in poker, "donkey." Or any phrase that starts by characterizing someone as "another typical _____." (Like David Letterman said to Lauren Conrad, he dismissed everyone around him as an idiot, until he pulled himself into deliberative mind one day and asked, "What are the odds that *everyone* is an idiot?")
- Other violations of the Mertonian norm of universalism, shooting the message because we don't think much of the messenger. Any sweeping term about someone, particularly when we equate our assessment of *an idea* with a sweeping personality or intellectual assessment of the person delivering the idea, such as "gun nut," "bleeding heart," "East Coast," "Bible belter," "California values"—political or social issues. Also be on guard for the reverse: accepting a message because of the messenger or praising a source immediately after finding out it confirms your thinking.
- Signals that we have zoomed in on a moment, out of proportion with the scope of time: "worst day ever," "the day from hell."
- Expressions that explicitly signal motivated reasoning, accepting or rejecting information without much evidence, like "conventional wisdom" or "if you ask anybody" or "Can you prove that it's not true?" Similarly, look for expressions that you're participating in an echo chamber, like "everyone agrees with me."
- The word "wrong," which deserves its own swear jar. The Mertonian norm of organized skepticism allows little place in exploratory discussion for the word "wrong." "Wrong" is a conclusion, not a rationale. And it's not a particularly accurate conclusion since, as we know, nearly nothing is 100% or 0%. Any words or thoughts denying the existence of uncertainty should be a signal that we are heading toward a poorly calibrated decision.
- Lack of self-compassion: if we're going to be self-critical, the focus should be on the lesson and how to calibrate future decisions. "I have

the worst judgment on relationships” or “I should have known” or “How could I be so stupid?”

- Signals we’re being overly generous editors when we share a story. Especially in our truthseeking group, are we straying from sharing the facts to emphasize our version? Even outside our group, unless we’re sharing a story purely for entertainment value, are we assuring that our listener will agree with us? In general, are we violating the Mertonian norm of communism?
- Infecting our listeners with a conflict of interest, including our own conclusion or belief when asking for advice or informing the listener of the outcome before getting their input.
- Terms that discourage engagement of others and their opinions, including expressions of certainty and also initial phrasing inconsistent with that great lesson from improvisation—“yes, and . . .” That includes getting opinions or information from others and starting with “no” or “but . . .”

This is by no means a complete list, but it provides a flavor of the kinds of statements and thinking that should trigger vigilance on our part.

Once we recognize that we should watch out for particular words, phrases, and thoughts, when we find ourselves saying or thinking those things, we are breaking a contract, a commitment to truthseeking. These terms are signals that we’re succumbing to bias. Because if we are cutting the binding when we catch ourselves saying or thinking these things, it can trigger a moment of reflection, interrupting us in the moment. Popping out of the moment like that can remind us of why we took the trouble to list terms that signal potential decision traps.

The swear jar is a simple example of a Ulysses contract in action: we think ahead to a hazard in our decision-making future and devise a plan of action around that, or at least commit that we will take a moment to recognize we are veering away from truthseeking. Better precommitment contracts result from better anticipation of what the future might look like, what kinds of decisions we want to avoid, and which ones we want to promote. That takes thoughtful reconnaissance.

Reconnaissance: mapping the future

Operation Overlord, the Allied forces operation to retake German-occupied France starting in Normandy, was the largest seaborne invasion in military history. It involved planning and logistics on an unprecedented scale. What if the forces were delayed at the start by bad weather? What if the airborne landing force had trouble communicating by radio due to the terrain? What if significant numbers of paratroopers were blown off course? What if currents interfered with the beach landings? What if the forces on the different beaches remained separated? Countless things could go wrong, with tens of thousands of lives at stake and, potentially, the outcome of the war.

All those things *did* go wrong, along with many other challenges encountered on D-Day and immediately thereafter. The Normandy landings still succeeded, though, because they prepared for as many potential scenarios as possible. Reconnaissance has been part of advance military planning for as long as horses have been used in battle. The modern military, of course, has evolved from sending scouts by horse and reporting back to the main forces to planes, drones, satellites, and other high-tech equipment gathering information about what to expect in battle.

The Navy SEAL team that caught and killed Osama bin Laden wouldn't have entered his compound without knowing what they would find beyond the walls. What buildings were there? What was their layout and purpose? What differences might it make if they conducted the raid in different kinds of weather or different times of day? What other people would be present and what risks would they pose? What would they do if bin Laden wasn't there? What was the team trying to commit to, given what they knew about each of those things (and, of course, numerous others)? Just as they relied on reconnaissance, we shouldn't plan our future without doing advance work on the range of futures that could result from any given decision and the probabilities of those futures occurring.

For us to make better decisions, we need to perform reconnaissance on the future. If a decision is a bet on a particular future based on our beliefs, then before we place a bet we should consider in detail what those possible futures might look like. Any decision can result in a set of possible outcomes.

BETTING ON A FUTURE

BELIEF → BET → (SET OF OUTCOMES)

Thinking about what futures are contained in that set (which we do by putting memories together in a novel way to imagine how things might turn out) helps us figure out which decisions to make.




Figure out the possibilities, then take a stab at the probabilities. To start, we imagine the range of potential futures. This is also known as scenario planning. Nate Silver, who compiles and interprets data from the perspective of getting the best strategic use of it, frequently takes a scenario-planning approach. Instead of using data to make a particular conclusion, he sometimes takes the approach of discussing all the scenarios the data *could* support. In early February 2017, he described the merits of scenario planning: “When faced with highly uncertain conditions, military units and major corporations sometimes use an exercise called scenario planning. The idea is to consider a broad range of possibilities for how the future might unfold to help guide long-term planning and preparation.”

After identifying as many of the possible outcomes as we can, we want to make our best guess at the probability of each of those futures occurring. When I consult with enterprises on building decision trees and determining probabilities of different futures, people frequently resist having to make a guess at the probability of future events mainly because they feel like they can’t be certain of what the likelihood of any scenario is. But that’s the point.

The reason why we do reconnaissance is *because* we are uncertain. We don't (and likely can't) know how often things will turn out a certain way with exact precision. It's not about approaching our future predictions from a point of perfection. It's about acknowledging that we're already making a prediction about the future every time we make a decision, so we're better off if we make that explicit. If we're worried about guessing, we're *already* guessing. We are already guessing that the decision we execute will result in the highest likelihood of a good outcome given the options we have available to us. By at least trying to assign probabilities, we will naturally move away from the default of 0% or 100%, away from being sure it will turn out one way and not another. Anything that moves us off those extremes is going to be a more reasonable assessment than not trying at all. Even if our assessment results in a wide range, like the chances of a particular scenario occurring being between 20% and 80%, that is still better than not guessing at all.

This kind of reconnaissance of the future is something that experienced poker players are very familiar with. Before making a bet, poker players consider each of their opponents' possible responses (fold, call, raise), and the likelihood and desirability of each. They also think about what they will do in response (if some or all of the opponents don't fold). Even if you don't know much about poker, it should make sense that a player is better off considering these things *before* they bet. The more expert the player, the further into the future they plan. Before making that decision to bet, the expert player is anticipating what they'll do following each response, as well as how the action they take now affects their future decisions on the hand. The best players think beyond the current hand into subsequent hands: how do the actions of this hand affect how they and their opponents make decisions on future hands? Poker players really live in this probabilistic world of, "What are the possible futures? What are the probabilities of those possible futures?" And they get very comfortable with the fact that they don't know exactly because they can't see their opponent's cards.

This is true of most strategic thinking. Whether it involves sales strategies, business strategies, or courtroom strategies, the best strategists are considering a fuller range of possible scenarios, anticipating and considering the strategic responses to each, and so on deep into the decision tree.

This kind of scenario planning is a form of mental time travel we can do on our own. It works even better when we do it as part of a scenario-planning group, particularly one that is open-minded to dissent and diverse points of view. Diverse viewpoints allow for the identification of a wider variety of scenarios deeper into the tree, and for better estimates of their probability. In fact, if two people in the group are really far off on an estimate of the likelihood of an outcome, that is a great time to have them switch sides and argue the other's position. Generally, the answer is somewhere in the middle and both people will end up moderating their positions. But sometimes one person has thought of a key influencing factor the other hasn't and that is revealed only because the dissent was tolerated.

In addition to increasing decision quality, scouting various futures has numerous additional benefits. First, scenario planning reminds us that the future is inherently uncertain. By making that explicit in our decision-making process, we have a more realistic view of the world. Second, we are better prepared for how we are going to respond to different outcomes that might result from our initial decision. We can anticipate positive or negative developments and plan our strategy, rather than being reactive. Being able to respond to the changing future is a good thing; being surprised by the changing future is not. Scenario planning makes us nimbler because we've considered and are prepared for a wider variety of possible futures. And if our reconnaissance has identified situations where we are susceptible to irrationality, we can try to bind our hands with a Ulysses contract. Third, anticipating the range of outcomes also keeps us from unproductive regret (or undeserved euphoria) when a particular future happens. Finally, by mapping out the potential futures and probabilities, we are less likely to fall prey to resulting or hindsight bias, in which we gloss over the futures that did not occur and behave as if the one that did occur must have been inevitable, because we have memorialized all the possible futures that could have happened.

Scenario planning in practice

A few years ago, I consulted with a national nonprofit organization, After-School All-Stars (ASAS), to work with them on incorporating scenario planning into their budgeting.* ASAS, founded in 1992 by Arnold Schwarzenegger, provides three hours of structured after-school programming for over 70,000 underserved youth in eighteen cities across the United States. They depend heavily on grants for funding and were struggling with budget planning, given the uncertainty in the grant award process. To help them with planning, I asked for a list of all their grant applications and how much each grant was worth. They provided me with a list of all their outstanding grant applications and the award amounts applied for. I told them that I didn't see how much each grant was worth in the information they provided. They pointed to the column of the award amounts sought. At that point, I realized we were working from different ideas about how to determine worth. The misunderstanding came from the disconnect between the *expected value* of each grant and the amount they would be awarded *if they got the grant*.*

Coming up with the expected value of each grant involves a simple form of scenario planning: imagining the two possible futures that could result from the application (awarded or declined) and the likelihood of each future. For example, if they applied for a \$100,000 grant that they would win 25% of the time, that grant would have an expected value of \$25,000 ($\$100,000 \times .25$). If they expected to get the grant a quarter of the time, then it wasn't worth \$100,000; it was worth a quarter of \$100,000. A \$200,000 application with a 10% chance of success would have an expected value of \$20,000. A \$50,000 grant with a 70% chance of success would be worth \$35,000. Without thinking probabilistically in this way, determining a grant's worth isn't possible—it leads to the mistaken belief that the \$200,000 grant is worth the most when, in fact, the \$50,000 grant is. ASAS recognized that uncertainty was causing them problems (to the point where they felt enslaved by it in their budgeting), but they hadn't wrapped uncertainty into their planning or resource-allocation process. They were flying by the seat of their pants.

After I worked with the national office, ASAS made estimating the likelihood of getting each grant part of their planning. The benefits they got from scenario planning were immediate and substantial:

- They created a more efficient and productive work stack. Before doing this exercise, they naturally put a higher priority on applications seeking larger dollar amounts, executing those first, putting more senior staff on them, and being more likely to hire outside grant writers to get those applications completed. By shifting to thinking about the probability of getting each grant, they were now able to prioritize by how much the grant was actually worth to the organization in making these decisions. Thereafter, the higher *value* grants went to the top of the stack rather than just the grants with the higher potential awards.
- They could budget more realistically. They had greater confidence in advance estimates of the amount of money they could expect to receive.
- Because coming up with an expected value required estimating the likelihood of getting the grant, they increasingly focused on improving the accuracy of their estimates. This prompted them to close the loop by going back to the grantors. They had previously, after rejections, followed up with grantors. Because they were now focusing on checking and calibrating their probabilities, they expanded this to the grants they won. Overall, their post-outcome reviews focused on understanding what worked, what didn't work, what was luck, and how to do better, improving both their probability estimates and the quality of their grant applications.
- They could think about ways to increase the probability of getting grants and commit to those actions.
- They were less likely to fall prey to hindsight bias because they had considered in advance the probability of getting or not getting the grant.
- They were less likely to fall prey to resulting because they had evaluated the decision process in advance of getting or not getting the grant.

- Finally, because of all the benefits ASAS received from incorporating scenario planning into budgeting and grant applications, they expanded the implementation of this type of scenario planning across departments, making it part of their decision-making culture.

Grant prospecting is similar to sales prospecting, and this process can be implemented for any sales team. Assign probabilities for closing or not closing sales, and the company can do better at establishing sales priorities, planning budgets and allocating resources, evaluating and fine-tuning the accuracy of its predictions, and protecting itself against resulting and hindsight bias.

A more complex version of scenario planning occurs when the number of possible futures multiplies and/or we go deeper into the tree, considering what we will do next in response to how things turn out and what the set of outcomes of that next decision is and so on.

Consider the scenario planning involved in Seahawks coach Pete Carroll's much-criticized Super Bowl decision, trailing by four, twenty seconds remaining, one time-out, second-and-goal at the Patriots' one-yard line. Carroll has two general choices, run or pass, and they in turn lead to multiple scenarios.

If Carroll calls for a run, these are the possible futures: (a) touchdown (immediate win); (b) turnover-fumble (immediate loss); (c) tackle short of the goal line; (d) offensive penalty; and (e) defensive penalty. Futures (c)–(e) branch into additional scenarios. The most likely failure scenario, by far, is that the runner is tackled before reaching the end zone. Seattle could stop the clock with its final time-out, but if they run the ball again and do not score, time will expire.

If Carroll calls for a pass, the possible futures are (a) touchdown (immediate success); (b) turnover-interception (immediate failure); (c) incomplete pass; (d) sack; (e) offensive penalty; and (f) defensive penalty. Again, the first two futures essentially end the game and the others branch off into additional play calling and additional outcomes.

The main difference between passing and running is that calling a pass likely gives Seattle a total of three plays to score, instead of two if Carroll calls a running play. An unsuccessful run would require that Seattle use their final time-out to stop the clock so they could run a second play. An incomplete pass would stop the clock and leave Seattle with a time-out and

the chance to call those same two running plays. An interception, which negates the possibility of a second or third offensive play, is only a 2%–3% probability, a small price to pay for three chances to score rather than two. (A turnover caused by a fumble on a running play is 1%–2%.)*

Notice that the option of the extra play doesn't reveal itself without this kind of scouting of the future. Even after the fact, with lots of time to analyze the decision, very few commentators saw this advantage.

The important thing is that we do better when we scout all these futures and make decisions based on the probabilities and desirability of the different futures. ASAS couldn't guarantee it would get awarded every grant it applied for, but it could, through good process, make better decisions about which grants to prioritize and how much revenue to expect from the basket of grant proposals they submitted. Despite the outcry, I'd like to think that Pete Carroll didn't lose much sleep over his decision to call for Wilson to pass.

Reconnaissance of the future dramatically improves decision quality and reduces reactivity to outcomes. So far, we have talked about thinking ahead to what the future might look like. But it turns out that better decision trees, more effective scenario planning, results from working backward rather than forward.

Backcasting: working backward from a positive future

All methods of imagining the future are not created equal. You know that Chinese proverb, “A journey of a thousand miles starts with a single step”? Turns out, if we were contemplating a thousand-mile walk, we'd be better off imagining ourselves looking back from the destination and figuring how we got there. When it comes to advance thinking, standing at the end and looking backward is much more effective than looking forward from the beginning.

The distorted view we get when we look into the future from the present is similar to the stereotypical view of the world by Manhattan residents, poked fun at by the famous *New Yorker* cover. The cover was a drawing of a

map from the perspective of a person from New York. In that view, half the map covers a few blocks of the city. While you can see all the buildings on Ninth Avenue and even vehicles and people, the Hudson River and New Jersey are just horizontal strips. The entire United States occupies the same amount of space as the distances between Ninth and Tenth avenues. Beyond a strip of the Pacific Ocean are three tiny lumps, labeled “China,” “Japan,” and “Russia.”

When we forecast the future, we run the risk of a similar distortion. From where we stand, the present and the immediate future loom large. Anything beyond that loses focus.

Imagining the future recruits the same brain pathways as remembering the past. And it turns out that remembering the future is a better way to plan for it. From the vantage point of the present, it’s hard to see past the next step. We end up over-planning for addressing problems we have right now. Implicit in that approach is the assumption that conditions will remain the same, facts won’t change, and the paradigm will remain stable. The world changes too fast to assume that approach is generally valid. Samuel Arbesman’s *The Half-Life of Facts* makes a book-length case for the hazards of assuming the future is going to be like the present.

Just as great poker players and chess players (and experts in any field) excel by planning further into the future than others, our decision-making improves when we can more vividly imagine the future, free of the distortions of the present. By working backward from the goal, we plan our decision tree in more depth, because we start at the end.

When we identify the goal and work backward from there to “remember” how we got there, the research shows that we do better. In a *Harvard Business Review* article, decision scientist Gary Klein summarized the results of a 1989 experiment by Deborah Mitchell, J. Edward Russo, and Nancy Pennington. They “found that prospective hindsight—imagining that an event has already occurred—increases the ability to correctly identify reasons for future outcomes by 30%.”

A huge urban planning project requires enormous amounts of money, materials, commitment—and the vision to work backward from a distant future goal. When Frederick Law Olmsted designed Central Park, for example, he recognized that a lot of the charm of the park, and a lot of enjoyment that people would get from it, would take decades, as the

landscape changed and matured. People walking through Central Park when it opened to the public in 1858 would have seen a lot of barren land. Even in 1873, when construction was substantially completed, there was a lot of undergrown foliage. The trees, shrubs, and plants were clearly recent transplants. None of those visitors would recognize Central Park today. But Olmstead would because he was starting from what it would develop into.

The most common form of working backward from our goal to map out the future is known as backcasting. In backcasting, we imagine we've already achieved a positive outcome, holding up a newspaper with the headline "We Achieved Our Goal!" Then we think about how we got there.

Let's say an enterprise wants to develop a three-year strategic plan to double market share, from 5% to 10%. Each person engaged in the planning imagines holding up a newspaper whose headline reads "Company X Has Doubled Its Market Share over the Past Three Years." The team leader now asks them to identify the reasons they got there, what events occurred, what decisions were made, what went their way to get the enterprise to capture that market share. This enables the company to better identify strategies, tactics, and actions that need to be implemented to get to the goal. It also allows it to identify when the goal needs to be tweaked. Backcasting makes it possible to identify when there are low-probability events that must occur to reach the goal. That could lead to developing strategies to increase the chances those events occur or to recognizing the goal is too ambitious. The company can also make precommitments to a plan developed through backcasting, including responses to developments that can interfere with reaching the goal and identifying inflection points for re-evaluating the plan as the future unfolds.

In planning a trial strategy after taking a new case, a trial lawyer can imagine the headline of winning at trial. What favorable rulings did the lawyer get along the way? How did the most favorable testimony play out? What kind of evidence did the judge allow or throw out? What points did the jury respond to?

If our goal is to lose twenty pounds in six months, we can plan how to achieve that by imagining it's six months from now and we've lost the weight. What are the things we did to lose the weight? How did we avoid junk food? How did we increase the amount of exercise we were doing? How did we stick to the regimen?

Imagining a successful future and backcasting from there is a useful time-travel exercise for identifying necessary steps for reaching our goals. Working backward helps even more when we give ourselves the freedom to imagine an *unfavorable* future.

Premortems: working backward from a negative future

If you know medical terms or watch forensic-crime dramas on TV, you are familiar with a postmortem: that's where a medical examiner determines the cause of death. A premortem is an investigation into something awful, but *before* it happens. We all like to bask in an optimistic view of the future. We generally are biased to overestimate the probability of good things happening. Looking at the world through rose-colored glasses is natural and feels good, but a little naysaying goes a long way. A premortem is where we check our positive attitude at the door and imagine not achieving our goals.

Backcasting and premortems complement each other. Backcasting imagines a positive future; a premortem imagines a negative future. We can't create a complete picture without representing both the positive space and the negative space. Backcasting reveals the positive space. Premortems reveal the negative space. Backcasting is the cheerleader; a premortem is the heckler in the audience.

Imagining a headline that reads "We Failed to Reach Our Goal" challenges us to think about ways in which things could go wrong that we otherwise wouldn't if left to our own (optimistic, team-player) devices. For the company with the three-year plan to double market share, the premortem headline is "Company Fails to Reach Market Share Goal; Growth Again Stalls." Members of the planning team now imagine delays in new products, loss of key executives or sales or marketing or technical personnel, new products by competitors, adverse economic developments, paradigm shifts that could lead customers to do without the product or rely on alternatives not on the market or in use, etc.

A lawyer trying a case now considers favorable evidence being disallowed, undiscovered evidence undermining the case, drawing an unsympathetic judge, and the jury disliking or distrusting the main witnesses.

When we set a weight-loss goal and put a plan to reach that goal in place, a premortem will reveal how we felt obligated to eat cake when it was somebody's birthday, how hard it was to resist the bagels and cookies in the conference room, and how hard it was to find time for the gym or how easy it was to find excuses to procrastinate going. There has been a massive amount written about visualizing success as a way to achieve our goals. Because that's such a common element in self-help strategies, conducting a premortem (with its negative visualization) may seem like a counterproductive way to succeed.

Despite the popular wisdom that we achieve success through positive visualization, it turns out that incorporating *negative visualization* makes us *more* likely to achieve our goals. Gabriele Oettingen, professor of psychology at NYU and author of *Rethinking Positive Thinking: Inside the New Science of Motivation*, has conducted over twenty years of research, consistently finding that people who imagine obstacles in the way of reaching their goals are more likely to achieve success, a process she has called "mental contrasting." Her first study, of women enrolled in a weight-loss program, found that subjects "who had strong positive fantasies about slimming down . . . lost twenty-four pounds *less* than those who pictured themselves more negatively. Dreaming about achieving a goal apparently didn't help that goal come to fruition. It impeded it from happening. The starry-eyed dreamers in the study were less energized to behave in ways that helped them lose weight."

She repeated these results in different contexts. She recruited college students who claimed to have an unrequited crush. She then prompted one group to imagine positive scenarios of initiating a relationship, and another group to imagine negative scenarios of how that would play out. The results were similar to those in the weight-loss study: five months later, the subjects indulging in the positive scenario planning were less likely to initiate the relationship. She found the same results when studying job seekers, students before a midterm exam, and patients undergoing hip replacement surgery.

Oettingen recognized that we need to have positive goals, but we are more likely to execute on those goals if we think about the negative futures. We start a premortem by imagining why we failed to reach our goal: our company hasn't increased its market share; we didn't lose weight; the jury

verdict came back for the other side; we didn't hit our sales target. Then we imagine why. All those reasons why we didn't achieve our goal help us anticipate potential obstacles and improve our likelihood of succeeding.

A premortem is an implementation of the Mertonian norm of organized skepticism, changing the rules of the game to give permission for dissent. Being a team player in a premortem isn't about being the most enthusiastic cheerleader; it's about being the most productive heckler. "Winning" isn't about the group feeling good because everyone confirms their (and the organization's) narrative that things are going to turn out great. The premortem starts with working backward from an unfavorable future, or failure to achieve a goal, so competing for favor, or feeling good about contributing to the process, is about coming up with the most creative, relevant, and actionable reasons for why things didn't work out.

The key to a successful premortem is that everyone feels free to look for those reasons, and they are motivated to scour everything—personal experience, company experience, historical precedent, episodes of *The Hills*, sports analogies, etc.—to come up with ways a decision or plan can go bad, so the team can anticipate and account for them.

Conducting a premortem creates a path to act as our own red team. Once we frame the exercise as "Okay, we failed. Why did we fail?" that frees everyone to identify potential points of failure they otherwise might not see or might not bring up for fear of being viewed as a naysayer. People can express their reservations without it sounding like they're saying the planned course of action is wrong. Because of that, a planning process that includes a premortem creates a much healthier organization because it means that the people who do have dissenting opinions are represented in the planning. They don't feel like they're shut out or not being heard. Everyone's voice now has more value. The organization is less likely to discourage dissent and thereby lose the value of diverse opinions. Those who have reservations are less likely to have resentment or regret build if things don't work out; their voices were represented in the strategic plan.

Incorporating this type of imagining of the negative space into a truthseeking group reinforces a new habit routine of visualizing and anticipating future obstacles. As always, when a group we are part of reinforces this kind of thinking, we are more likely in our own thinking to consider the downside of our decisions.

Imagining both positive and negative futures helps us build a more realistic vision of the future, allowing us to plan and prepare for a wider variety of challenges, than backcasting alone. Once we recognize the things that can go wrong, we can protect against the bad outcomes, prepare plans of action, enable nimble responses to a wider range of future developments, and assimilate a negative reaction in advance so we aren't so surprised by it or reactive to it. In doing so, we are more likely to achieve our goals.

Of course, when we backcast and imagine the things that went right, we reveal the problems if those things *didn't* go right. Backcasting doesn't, therefore, ignore the negative space so much as it overrepresents the positive space. It's in our optimistic nature (and natural in backcasting) to imagine a successful future. Without a premortem, we don't see as many paths to the future in which we don't reach our goals. A premortem forces us to build out that side of the tree where things don't work out. In the process, we are likely to realize that's a pretty robust part of the tree.

Remember, the likelihood of positive and negative futures must add up to 100%. The positive space of backcasting and the negative space of a premortem still have to fit in a finite amount of space. When we see how much negative space there really is, we shrink down the positive space to a size that more accurately reflects reality and less reflects our naturally optimistic nature.

We make better decisions, and we feel better about those decisions, once we get our past-, present-, and future-selves to hang out together. This not only allows us to adjust how optimistic we are, it allows us to adjust our goals accordingly and to actively put plans in place to reduce the likelihood of bad outcomes and increase the likelihood of good ones. We are less likely to be surprised by a bad outcome and can better prepare contingency plans.

It may not feel so good during the planning process to include this focus on the negative space. Over the long run, however, seeing the world more objectively and making better decisions will feel better than turning a blind eye to negative scenarios. In a way, backcasting without premortems is a form of temporal discounting: if we imagine a positive future, we feel better *now*, but we'll more than compensate for giving up that immediate gratification through the benefits of seeing the world more accurately, making better initial decisions, and being nimbler about what the world throws our way.

Once we make a decision and one of those possible futures actually happens, we can't discard all that work, even—or especially—if it included work on futures that did not occur. Forgetting about an unrealized future can be dangerous to good decision-making.

Dendrology and hindsight bias (or, Give the chainsaw a rest)

One of the goals of mental time travel is keeping events in perspective. To understand an overriding risk to that perspective, think about time as a tree. The tree has a trunk, branches at the top, and the place where the trunk meets the branches. The trunk is the past. A tree has only one, growing trunk, just as we have only one, accumulating past. The branches are the potential futures. Thicker branches are the equivalent of more probable futures, thinner branches are less probable ones. The place where the top of the trunk meets the branches is the present. There are many futures, many branches of the tree, but only one past, one trunk.

As the future becomes the past, what happens to all those branches? The ever-advancing present acts like a chainsaw. When one of those many branches happens to be the way things turn out, when that branch transitions into the past, present-us cuts off all those other branches that didn't materialize and obliterates them. When we look into the past and see only the thing that happened, it seems to have been inevitable. Why wouldn't it seem inevitable from that vantage point? Even the smallest of twigs, the most improbable of futures—like the 2%–3% chance Russell Wilson would throw that interception—expands when it becomes part of the mighty trunk. That 2%–3%, in hindsight, becomes 100%, and all the other branches, no matter how thick they were, disappear from view.

That's hindsight bias, an enemy of probabilistic thinking.

Judge Frank Easterbrook, a leading jurist and member of the U.S. Court of Appeals for the Seventh Circuit, warned about the danger in the legal system of assessing probabilities *after* one of the potential futures has occurred. *Jentz v. ConAgra Foods* involved a “hot” grain bin owned by

ConAgra and the company it hired (named West Side) to investigate and deal with the cause of the burning smell, smoke, and raised temperatures from the bin. After failing to solve the problem, West Side's foreman told ConAgra to call the fire department. He then told some of his employees to remove tools from a tunnel to the bin that could impair firefighters' access.

While the workers were in the tunnel, the bin exploded, severely injuring the West Side employees. Those injured sued ConAgra and West Side. The jury awarded \$180 million in compensatory and punitive damages. In addressing the punitive damages against West Side, Judge Easterbrook, writing for the court, pointed out that Illinois law required a "gross deviation" from the standard of care to award punitive damages. Finding no evidence in the record of any foreseeable likelihood of explosion at the time the foreman ordered the workers to remove tools from the tunnel, he concluded, "The verdict appears to be a consequence of hindsight bias—the human tendency to believe that whatever happened was bound to happen, and that everyone must have known it. If [the foreman] believed that an explosion was imminent, then he is a monster; but of that there is no evidence. Hindsight bias is not enough to support a verdict."

Once we know there was an explosion, it's difficult to imagine the actions of the parties when the explosion was only one of several possible futures. The members of the jury had a conflict of interest. When they heard the story of the men entering the tunnel to retrieve the tools, they knew the outcome. The jury lopped off the other branches of the tree, all the other ways things could have turned out given the situation in the grain-storage bin and, in hindsight, all they could see was a confluence of circumstances with an inevitable tragedy at the end.

Imagine how things would be different for the protagonists throughout this book if more of us acted like Judge Easterbrook.

Steve Bartman could have used a stadium full of Cubs fans with Judge Easterbrook's perspective. They would have recognized that the Cubs losing the game was only one of many ways it could have turned out. At the moment Bartman and others around him reached for the foul ball, the branch of the future where Steve Bartman touches the ball was a small one, and the chance of the Cubs losing after that was the tiniest of twigs, requiring a number of unlikely on-field developments (like the pitching ace giving up hit after hit, and the great-fielding shortstop committing a rare error on an inning-ending

double play). Just as the foreman ordering workers into the tunnel didn't cause the explosion, Bartman touching the ball didn't cause the loss. He just got an unlucky result, based on all the things that happened over which he had no control after the foul ball touched his hands.

Pete Carroll and the world of Monday Morning Quarterbacks could have used the judge's reminder that we tend to assume that, once something happens, it was bound to happen. If we don't try to hold all the potential futures in mind *before* one of them happens, it becomes almost impossible to realistically evaluate decisions or probabilities *after*.

This turned out to be a big part of the problem for the CEO in the wake of firing his president. Although he had initially characterized the decision as one of his worst, when we reconstructed the tree, essentially picking the branches up off the ground and reattaching them, it was clear that he and his company had made a series of careful, deliberative decisions. Because they led to a negative result, however, the CEO had been consumed with regret. When he looked back at his decisions, he couldn't see all those branches and their likelihoods. He could only see the trunk. All he saw was a bad result.

The CEO removed all those other branches with a chainsaw and ran them through a wood-chipper. They disappeared and he acted as if they never existed. That's what hindsight bias is, and we're all running amok through the forest with a chainsaw once we get an outcome. Once something occurs, we no longer think of it as probabilistic—or as *ever* having been probabilistic. This is how we get into the frame of mind where we say, "I should have known" or "I told you so." This is where unproductive regret comes from.

By keeping an accurate representation of what could have happened (and not a version edited by hindsight), memorializing the scenario plans and decision trees we create through good planning process, we can be better calibrators going forward. We can also be happier by recognizing and getting comfortable with the uncertainty of the world. Instead of living at extremes, we can find contentment with doing our best under uncertain circumstances, and being committed to improving from our experience.

Reaction to the 2016 election provides another strong demonstration of what happens when we lop branches off the tree. Hillary Clinton had been favored going into the election, and her probability of winning, based on an accumulation of the polls, was somewhere between 60% and 70%, according to FiveThirtyEight.com. When Donald Trump won, pollsters got the Pete

Carroll treatment, maybe no one more than Nate Silver, founder of FiveThirtyEight.com and a thoughtful analyzer of polling data. (“Nate Silver was wrong.” “The pollsters missed it.” “Just like Brexit, the bookies blew it.” Etc.) The press spun this as a certain win for Clinton, despite the Trump branch of the tree being no mere twig at 30%–40%. By the day after the election, the Clinton branch had been severed, only the Trump branch remained, and how could pollsters and the polling process have been so blind?

One of the things poker teaches is that we have to take satisfaction in assessing the probabilities of different outcomes given the decisions under consideration and in executing the bet we think is best. With the constant stream of decisions and outcomes under uncertain conditions, you get used to losing a lot. To some degree, we’re all outcome junkies, but the more we wean ourselves from that addiction, the happier we’ll be. None of us is guaranteed a favorable outcome, and we’re all going to experience plenty of unfavorable ones. We can always, however, make a good bet. And even when we make a bad bet, we usually get a second chance because we can learn from the experience and make a better bet the next time.

Life, like poker, is one long game, and there are going to be a lot of losses, even after making the best possible bets. We are going to do better, and be happier, if we start by recognizing that we’ll never be sure of the future. That changes our task from trying to be right every time, an impossible job, to navigating our way through the uncertainty by calibrating our beliefs to move toward, little by little, a more accurate and objective representation of the world. With strategic foresight and perspective, that’s manageable work. If we keep learning and calibrating, we might even get good at it.