First Principles Thinking

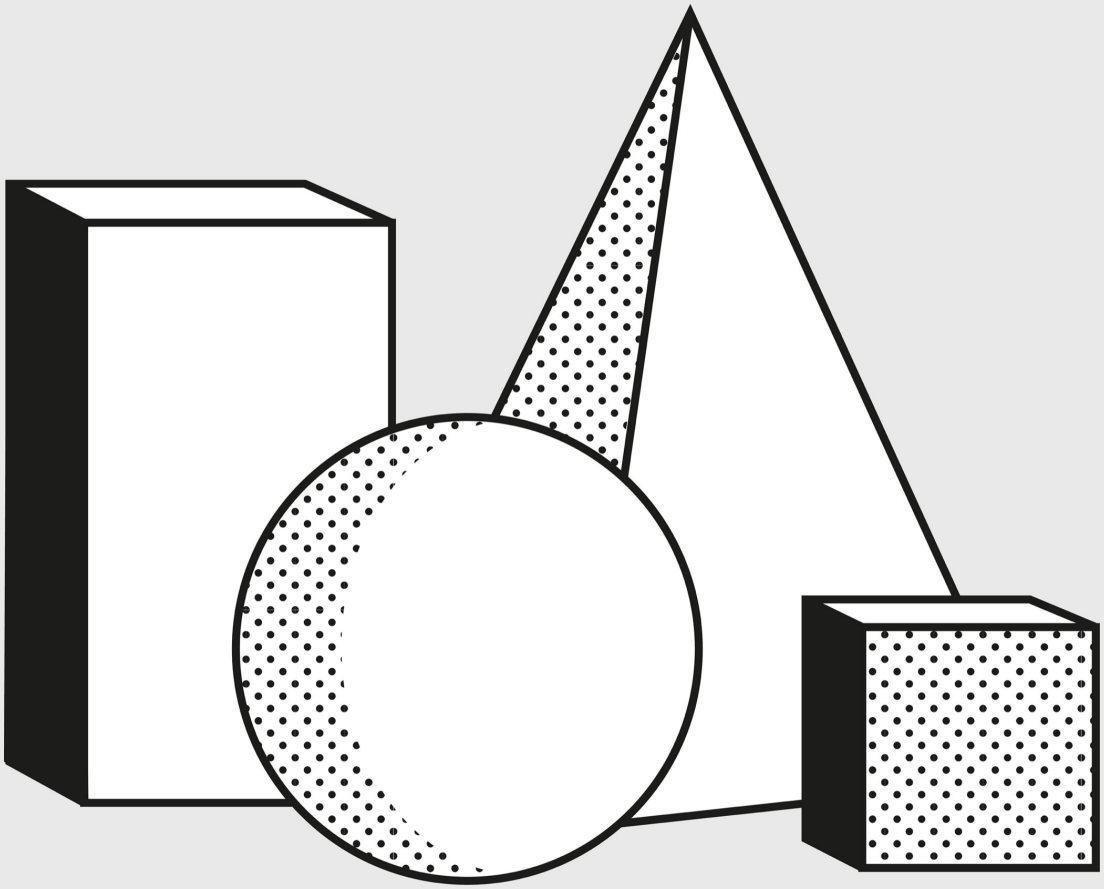

Go back to basics.

I don’t know what’s the matter with people: they don’t learn by understanding; they learn by some other way—by rote or something. Their knowledge is so fragile!

—RICHARD FEYNMAN[1]

First principles thinking is one of the best ways to discover new solutions. Sometimes called “reasoning from first principles,” it’s a tool to help break down complicated problems by separating what we know is absolutely true from anything that is an assumption. What remain are the essentials. If you know the first principles of something, you can build the rest of your knowledge around them to produce something new.

While you could take this way of thinking down to an atomic level, a lot of value is gained by simply going a level or two deeper than most people. [2] Solutions are based on what you see. Different answers reveal themselves at different levels.

If I hand you a house made from Lego blocks, you know it’s possible to make a house. Thinking at the first layer, you might move a few blocks around and, in the process, slightly improve the house. Most people stop here. They are presented with something that already exists and they endeavor to make it slightly better. Going a layer deeper and breaking the Lego house into individual pieces opens the door to possibility: not only can you build a better house, you can build something entirely different.

Everything that exists is effectively a set of Lego blocks, assembled in a certain way, that can be taken apart and reassembled. A bike is just a seat, chain, body, handlebars, etc. Breaking the bike down into its parts allows you to reassume the parts into something new. However, you can also go deeper, melting the parts into their core metals and making a shield, sword, or anything else, limited only by material and imagination.

The idea of building knowledge from first principles has a long tradition in philosophy. In the Western canon it goes back to Plato, with significant contributions from Aristotle and Descartes. Essentially, these thinkers were looking for foundational knowledge that would not change and on which

we could build everything else, from our ethical systems to our social structures.

First principles thinking doesn't have to be quite so grand. When we do it, we aren't necessarily looking for absolute truths—millennia of epistemological inquiry have shown us that these are hard to come by, and the scientific method has demonstrated that knowledge can be built only when we are actively trying to falsify it (see "Supporting Idea: Falsifiability"). Rather, first principles thinking identifies the elements that are, in the context of any given situation, irreducible.

First principles do not provide a checklist of things that will always be true; our knowledge of first principles changes as we understand more. They are the foundation on which we must build, and thus will be different in every situation—but the more we know, the more we can challenge. For example, if we are considering how to improve the energy efficiency of a refrigerator, the laws of thermodynamics can be taken as first principles. However, a theoretical chemist or physicist might want to explore entropy, and thus further break the second law of thermodynamics into its underlying principles and the assumptions that were made because of them. First principles are the boundaries that we must work within in any given situation, so when it comes to thermodynamics, an appliance maker might have different first principles than a physicist.

Techniques for Establishing First Principles

If we never learn to take something apart, test our assumptions about it, and reconstruct it, we end up bound by what other people tell us is possible. We end up trapped in the way things have always been done. When the environment changes, we just continue as if things were the same, making costly mistakes along the way.

Some of us are naturally skeptical of what we're told: Maybe it doesn't match up to our experiences. Maybe it's something that used to be true but isn't true anymore. Or maybe we just think differently about something.

When it comes down to it, everything that is not a law of nature is just a shared belief. Money is a shared belief. So is a border. So is Bitcoin. So is love. The list goes on.

There are two techniques we can use to change the level where we are looking at a situation, identify the first principles, and cut through the dogma and shared belief: Socratic questioning and the Five Whys.

Socratic questioning: Socratic questioning can be used to establish first principles through stringent analysis. This is a disciplined questioning process used to establish truths, reveal underlying assumptions, and separate knowledge from ignorance. The key distinction between Socratic questioning and ordinary discussion is that the former seeks to draw out first principles in a systematic manner. Socratic questioning generally follows this process:

1. Clarifying your thinking and explaining the origins of your ideas. (Why do I think this? What exactly do I think?)
2. Challenging assumptions. (How do I know this is true? What if I thought the opposite?)
3. Looking for evidence. (How can I back this up? What are my sources?)
4. Considering alternative perspectives. (What might others think? How do I know I am correct?)
5. Examining consequences and implications. (What if I am wrong? What are the consequences if I am?)
6. Questioning the original questions. (Why did I think that? Was I correct? What conclusions can I draw from the reasoning process?)

Socratic questioning stops you from relying on your gut and limits strong emotional responses. This process helps you build something that

lasts.

The Five Whys: The Five Whys is a method rooted in the behavior of children. Children instinctively think in first principles; just like us, they want to understand what's happening in the world. To do so, they intuitively break through the fog with a game some parents have come to dread but that is exceptionally useful for identifying first principles: repeatedly asking "why."

The goal of the Five Whys is to traverse different levels until we land on a "what" or "how." It is not about introspection, such as asking, "Why do I feel like this?" Rather, it is about systematically delving further into a statement or concept so that you can separate reliable knowledge from assumption. If your "whys" result in a statement of falsifiable fact, you have hit a first principle. If they end up with a "because I said so" or "it just *is*," you know you have landed on an assumption that may be based on popular opinion, cultural myth, or dogma. These are not first principles.

There is no doubt that both of these methods slow us down in the short term. They seem to get in the way of what we want to accomplish. We must pause, think, and research. And after we employ them a couple of times, we realize that often, after one or two questions, we are lost. We actually don't know how to answer most of the questions. But when we are confronted with our own ignorance, we can't just give up or resort to self-defense. If we do, we will never identify the first principles we have to work with and will instead make mistakes that will slow us down in the long term.

> Science is much more than a body of knowledge. It is a way of thinking.
>
> **—CARL SAGAN**[3]

Using First Principles Thinking to Blow Past Inaccurate Assumptions

The discovery that a bacterium, not stress, causes the majority of stomach ulcers is a great example of what can be accomplished when we push past assumptions to get at first principles. For centuries following the discovery of bacteria, scientists thought that bacteria could not grow in the stomach, on account of its acidity. If you had surveyed doctors and medical research scientists in the 1960s or ’70s, they likely would have postulated this as a first principle. When a patient came in complaining of stomach pain, no one ever looked for a bacterial cause.

It turned out, however, that a sterile stomach was not a first principle—it was an assumption. As Kevin Ashton writes in his book on creativity, discovery, and invention, “the dogma of the sterile stomach said that bacteria could not live in the gut.”[4] Because this dogma was taken as truth, for a long time, no one ever looked for evidence that it could be false.

That changed for good with the discovery of *Helicobacter pylori* bacterium and its role in stomach ulcers. When pathologist Robin Warren saw bacteria in samples from patients’ stomachs, he realized that stomachs were not, in fact, sterile. He started collaborating with Barry Marshall, a gastroenterologist, and together they found bacteria in loads of stomachs. If the sterile stomach wasn’t a first principle, then, when it came to stomachs, what was?

Marshall, in an interview with *Discover*, recounts that Warren gave him a list of twenty patients identified as possibly having cancer—but when Warren looked, he had found, instead, the same bacteria in all of them. He said, “Why don’t you look at their case records and see if they’ve got anything wrong with them?” Since they now knew stomachs weren’t sterile, they could question all the associated dogma about stomach disease and use some Socratic-type questioning to identify the first principles at play. They spent years challenging their related assumptions, clarifying their thinking, and looking for evidence.[5]

Their story ultimately had a happy ending: in 2005, Marshall and Warren were awarded the Nobel Prize, and now stomach ulcers are regularly treated effectively with antibiotics, improving and saving the lives of millions of people. But many practitioners and scientists rejected their

findings for decades. The dogma of the sterile stomach was so entrenched as a first principle that it was hard for many to admit that it rested on some incorrect assumptions that ultimately ended with the explanation, “because that’s just the way it is.” Even though, as Ashton notes, “*H. pylori* has now been found in medical literature dating back to 1875,”[6] it was Warren and Marshall who were able to show that “because I said so” wasn’t enough to count the sterile stomach as a first principle.

Incremental Innovation and Paradigm Shifts

Understanding how and why something works is a key step to improving it. First principles thinking helps us avoid the problem of relying on someone else’s tactics without understanding the rationale behind them.

Temple Grandin is famous for a couple of reasons. First, she is autistic, and was one of the first people to publicly disclose this fact and give insight into the inner workings of one type of autistic mind. Second, she is a scientist who has developed many techniques to improve the welfare of livestock in the agricultural industry.

One of the approaches Grandin pioneered was the curved cattle chute. Before her experiments, cattle were herded through a straight chute. Curved chutes, Grandin found, “are more efficient for handling cattle because they take advantage of the natural behavior of cattle. Cattle move through curved races more easily because they have a natural tendency to go back to where they came from.”[7] Of course, science doesn’t stop with one innovation, and animal scientists continue to study the best way to treat livestock animals.

Stockmanship Journal presented research that questioned the efficiency of Grandin’s curved chute. It demonstrated that sometimes, the much simpler straight chute would achieve the same effect in terms of cattle movement. The journal then sought out Grandin’s response, which is invaluable for teaching us the necessity of first principles thinking.

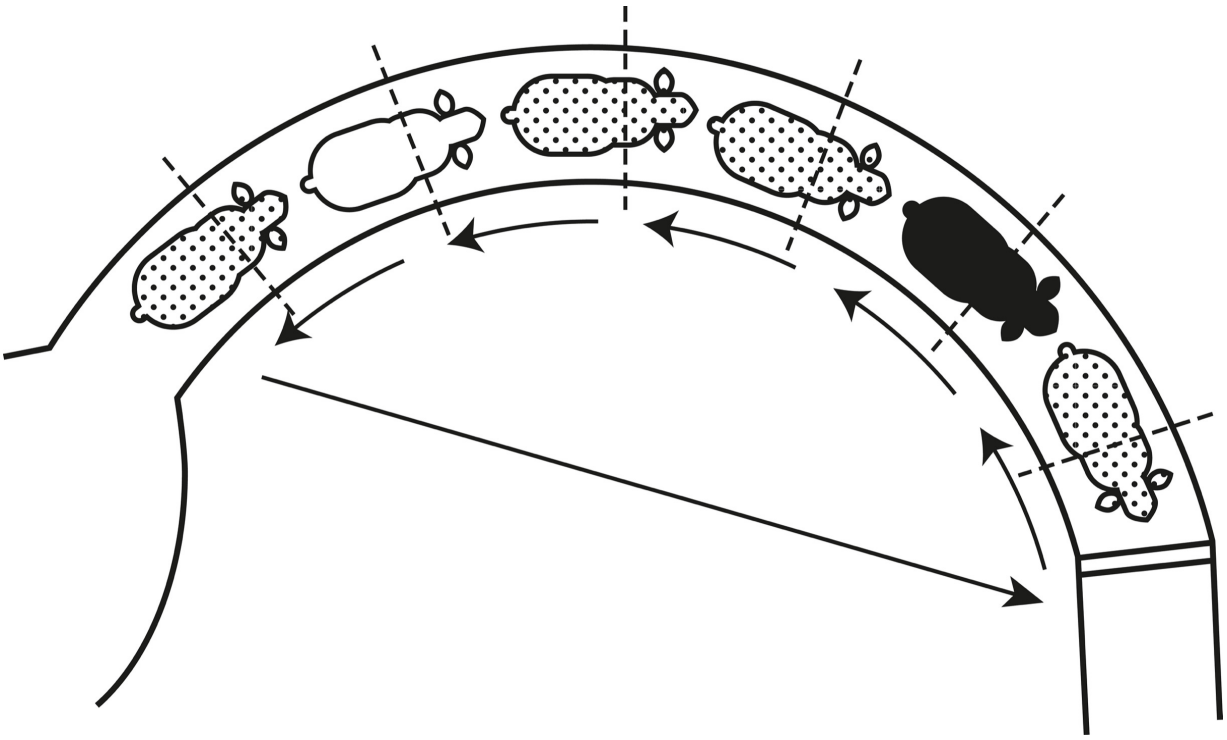

Grandin explained that curved chutes are not a first principle. She designed them as a tactic to address the first principle of animal handling that she identified in her research—essentially, that reducing stress to the animals is the single most important aspect of handling them and affects everything from their conception rates to their weight to their immune systems. When designing a livestock environment, she noted, a straight chute could work if it is part of a system that reduces stress to the animals. If you know the principles, you can change the tactics.[8]

Sometimes, we don’t want to fine-tune what is already there—we are skeptical, or curious, and are not interested in accepting what already exists as our starting point. When we start with the idea that the way things are might not be the way they have to be, we put ourselves in the right frame of mind to identify first principles. The real power of first principles thinking is moving away from random change and into choices that have a real possibility of success.

> As to methods, there may be a million and then some, but principles are few. The man who grasps principles can successfully select his own methods. The man who tries methods, ignoring principles, is sure to have trouble.
>
> **—HARRINGTON EMERSON**[9]

Conclusion

First principles thinking is the art of breaking down complex problems into their most fundamental truths. It's a way of thinking that goes beyond the surface and allows us to see things from a new perspective.

Thinking in terms of first principles allows us to identify the root causes and strip away the layers of complexity and focus on the most effective solutions. Reasoning from first principles allows us to step outside the way things have always been done and instead see what is possible.

First principles thinking is not easy. It requires a willingness to challenge the status quo. This is why it's often the domain of rebels and disrupters who believe there must be a better way. It's the thinking of those who are willing to start from scratch and build from the ground up.

In a world focused on incremental improvement, first principles thinking offers a competitive advantage because almost no one does it.

Second-Order Thinking

What happens next?

Technology is fine, but the scientists and engineers only partially think through their problems. They solve certain aspects, but not the total, and as a consequence it is slapping us back in the face very hard.

—BARBARA MCCLINTOCK[1]

Almost everyone can anticipate the immediate results of their actions. However, few people think about what happens next.

First-order thinking is easy and common. Second-order thinking is harder and requires thinking further ahead and thinking holistically. It requires us to consider not only our actions and their immediate consequences, but the subsequent effects of those actions as well. Failing to consider the second- and third-order effects of our decisions can unleash disaster.

First-order thinking is almost always about satisfying the immediate problem. Second-order thinking, on the other hand, avoids problems before they happen by asking, “And then what?”

Without second-order thinking, it can be hard to appreciate just how often what appears to solve the immediate problem takes you further away from your objective. First-order thinking tells you the chocolate bar tastes good and will satisfy your cravings. Second-order thinking tells you that when the sugar high wears off, you’ll crash.

It is often easier to find examples of when second-order thinking *didn’t* happen—when people did not consider the effects of the effects. When someone tried to do something good, or even just benign, and instead brought calamity, we can safely assume the negative outcomes weren’t factored into their original thinking. Very often, the second level of effect is not considered until it’s too late. This concept is often referred to as the “Law of Unintended Consequences” for this very reason.

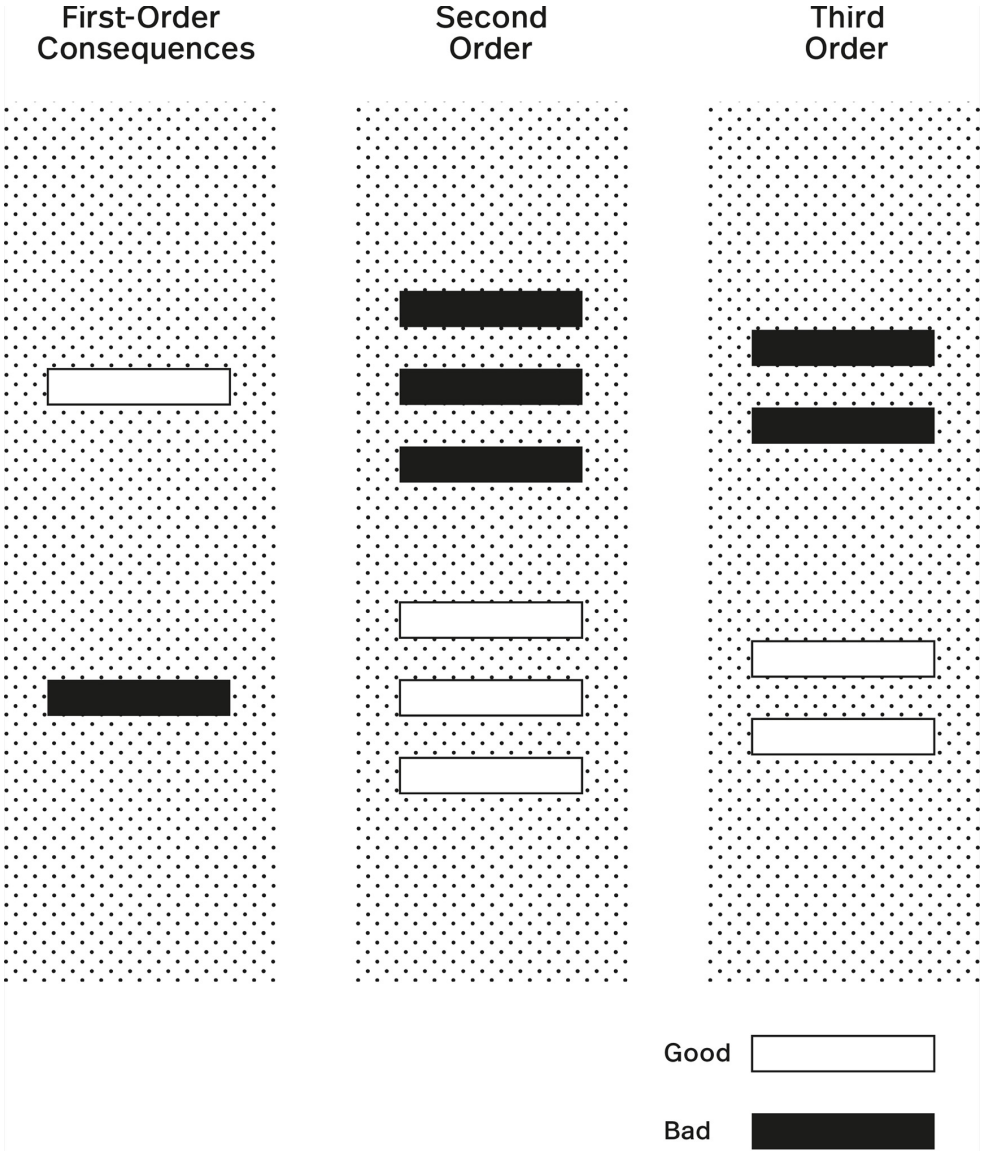

We see examples of this oversight throughout history. The British are a well-intentioned nation with an ample supply of smart politicians. However, during its colonial rule of India, the British government began to worry about the number of venomous cobras in Delhi. To reduce the population, they instituted a reward for every dead snake brought to officials. In response, Indian citizens dutifully began breeding the snakes to slaughter and bring to officials. The snake problem became worse than when the government first intervened, because the British officials didn’t think at the second level.

Second-order effects occur even with something as simple as adding traction on tires: it seems like such a great idea, because the more traction you have, the less likely you are to slide, the faster you can stop, and, thus, the safer you are. However, the second-order effects are that your engine must work harder to propel the car, you get worse gas mileage (releasing more detrimental carbon dioxide into the atmosphere), and you leave more rubber particles on the road.

This is why any comprehensive thought process considers the effects of the effects of a decision seriously. You're going to have to deal with them anyway. The genie never goes back in the bottle; you can never delete consequences to arrive back at the original starting conditions.

> Stupidity is the same as evil if you judge by the results.
>
> **—MARGARET ATWOOD**[2]

In an example of second-order-thinking deficiency, we have been feeding antibiotics to livestock for decades, to make the resulting meat safer and cheaper. Only in recent years have we begun to realize that in doing so, we have helped create bacteria that we cannot defend against.

In 1963, UC Santa Barbara ecologist Garrett Hardin proposed his First Law of Ecology: "You can never merely do one thing."[3] We operate in a world of multiple, overlapping connections, like a web, with many significant, yet obscure and unpredictable, relationships. Hardin developed second-order thinking into a tool, showing that if you don't consider "the effects of the effects," you can't really claim to be doing any thinking at all.

When it comes to the overuse of antibiotics in meat, the first-order consequence is that the animals gain more weight per pound of food consumed, and thus, there is profit for the farmer. Animals are sold by weight, so the less food you need to use to bulk them up, the more profit you make when you go to sell them. The second-order effects, however, include many serious, negative consequences. The bacteria that survive this

continued antibiotic exposure are antibiotic resistant. That means that the agricultural industry, when using these antibiotics as bulking agents, is allowing massive numbers of drug-resistant bacteria to become part of our food chain.

High degrees of connection make second-order thinking all the more critical, because denser webs of relationships make it easier for actions to have far-reaching consequences. You may be focused in one direction, not recognizing that the consequences of your decisions are rippling out all around you. Things are not produced and consumed in a vacuum.

> When we try to pick out anything by itself, we find it hitched to everything else in the Universe.
>
> **—JOHN MUIR**[4]

Second-order thinking is not a way to predict the future. You are only able to think of the likely consequences of your decisions' consequences based on the information available to you. However, this is not an excuse to power ahead and wait for post facto scientific analysis.

Could the consequences of putting antibiotics in livestock feed have been anticipated? Likely, yes, by anyone with even a limited understanding of biology. We know that organisms evolve. They adapt based on environmental pressures, and those with shorter life cycles, like bacteria, can do it quite quickly, because they have more opportunities to do so. Antibiotics, by definition, kill bacteria. Bacteria, just like all other living things, want to survive. The pressures put on them by continued exposure to antibiotics increase their pace of evolution. Over the course of many generations, eventually, mutations will occur that allow certain bacteria to resist the effects of the antibiotics. These are the bacteria that will then reproduce more rapidly, creating the situation we are now in.

Second-Order Problem

Warren Buffett used a very apt metaphor once to describe the second-order problem, likening it to a crowd at a parade: Once a few people decide to stand on their tiptoes to see better, *everyone has to stand on their tiptoes*. No one can see any better, but they're all worse off.[5]

Second-order thinking teaches us two important concepts that underline the utility of this model. If we're interested in understanding how the world works, we must think about second- and subsequent-level effects. We must understand that just because there is no immediate and visible impact from our decisions doesn't mean that we are not moving closer to or further from our objectives. How often is short-term gain worth protracted, long-term pain?

Let's look at two areas where second-order thinking can be used to great benefit:

1. Prioritizing long-term interests over immediate gains
2. Constructing effective arguments

Prioritizing Long-Term Interests

Thinking long-term eliminates a lot of poor behavior. Most people prefer to give in to instant gratification. If we want to avoid problems, however, we need to see past the immediate moment and into the future. If we forgo the immediate pleasure of candy, we improve our long-term health. The first-order effect of candy is the amazing feeling triggered by an influx of pure sugar in our system. But what are the second-order effects of regular candy consumption? Is that what I want my body or life to look like in ten years? Second-order thinking involves asking ourselves if what we are doing now is moving us closer to or further away from our objectives.

The most dangerous form of short-term thinking is one that doesn't understand that just because results are not visible doesn't mean they are not accumulating. Thinking long-term helps us see how the accumulation of tiny gains or losses moves us toward or away from our intended future.

Finding historical examples of second-order thinking can be tricky, because we don't want to evaluate based solely on the outcome: "It all turned out well, so he must have thought through the consequences of his

actions.” Even if you can glimpse the long-term gain from your short-term pain, there is no guarantee you’ll get there.

In 48 BC, Cleopatra of Egypt was in a terrible position.[6] Technically co-regent with her brother, in a family famous for murdering siblings, she was encamped in a swampy desert, ousted from the palace, with no solid plan for how to get back. She was queen, but she had made a series of unpopular decisions that left her with little political support and that gave her brother ample justification for trying to have her assassinated. What to do?

At the same time, the great Roman general Caesar arrived in Egypt, chasing down his enemy Pompey and making sure the Egyptians knew who really was in charge on the Mediterranean. Egypt was an incredibly fertile, wealthy country, and as such was of great importance to the Romans. The way they inserted themselves in Egypt, however, made them extremely unpopular there.

To survive, Cleopatra had to make some tough decisions. Should she try to work things out with her brother? Should she try to marshal some support from another country? Or should she try to align herself with Caesar?

In *Cleopatra: A Life*, Stacy Schiff explains that even in 48 BC, at the age of twenty-one, Cleopatra would have had a superb political education, based on both historical knowledge and firsthand exposure to the tumultuous events of life on the Mediterranean. She would have observed actions taken by her father, Auletes, as well as various family members, that resulted in exile, bribery, and murder from either a family member, the Romans, or the populace. She would have known that there were no easy answers. As Schiff explains, “What Auletes passed down to his daughter was a precarious balancing act. To please one constituency was to displease another. Failure to comply with Rome would lead to intervention. Failure to stand up to Rome would lead to riots.”[7]

In this situation, it was thus imperative that Cleopatra consider the second-order effects of her actions. Short-term gain might easily lead to execution (as indeed it already had for many of her relatives). If she wanted to be around for a while, she needed to balance her immediate goals of

survival and possession of the throne with the future need for support to stay on it.

In 48 BC, Cleopatra chose to align herself with Caesar. The first-order effects of this decision, it seems likely she would have known: namely, that it would anger her brother, who would increase his plotting to have her killed, and that it would anger the Egyptian people, who didn't want a Roman involved in their affairs. She probably anticipated that there would be short-term pain, and there was. Cleopatra effectively started a civil war, including a siege on the palace that left her and Caesar trapped there for months. In addition, she had to be constantly vigilant against the assassination schemes of her brother. So why did she do it?

We will never know for sure. We can only make an educated guess. But given that Cleopatra ruled Egypt quite successfully for many years after these events, her decision was probably based on seeing the effects of the effects: if she could somehow make it through the short-term pain, her leadership had a much greater chance of being successful with the support of Caesar and Rome than without it. As Schiff notes, "The Alexandrian War gave Cleopatra everything she wanted. It cost her little."[8] In winning the civil war, Caesar got rid of all major opposition to Cleopatra and firmly aligned himself with her reign.

Being aware of second-order consequences and using them to guide your decision making may mean the short term is less spectacular, but the payoffs for the long term can be enormous. By delaying gratification now, you will save time in the future. You won't have to clean up the mess you made on account of not thinking through the effects of indulging your short-term desires.

Developing Trust for Future Success

Trust and a sense of trustworthiness are the results of multiple interactions. This is why second-order thinking is so useful and valuable. Going for the immediate payoff in our interactions with people, unless the result is a win-win, almost always guarantees that interaction will be a one-off. Maximizing benefits is something that becomes possible only over time. Thus, considering the effects of the effects of our actions on others, or on our reputations, is critical to getting people to trust us and to enjoying the benefits of cooperation that come with that trust.[9]

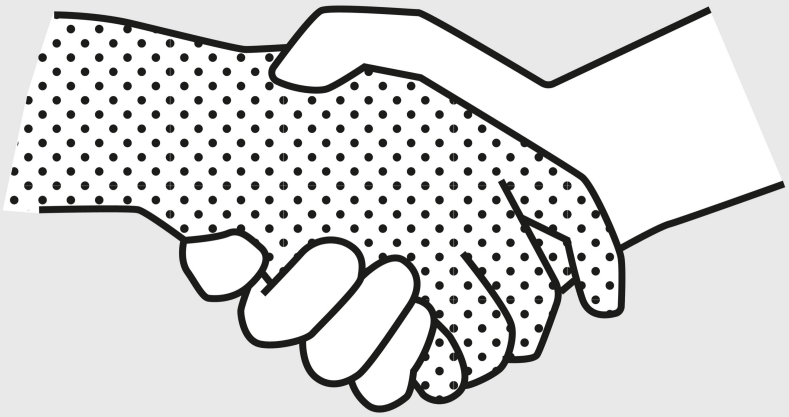

Constructing effective arguments: Second-order thinking can help you avert problems and anticipate challenges that you can then address in advance.

For example, you construct arguments every day: convincing your boss to take a chance on a new product, convincing your spouse to try a new parenting technique. Life is filled with the need to be persuasive. Arguments are more effective when we demonstrate that we have considered the second-order effects of a decision and put effort into verifying that these are desirable as well.

In late-eighteenth-century England, women had very few rights. Philosopher Mary Wollstonecraft was frustrated that this lack of rights limited a woman's ability to be independent and make choices on how to live her life. Instead of arguing, however, for why women should have rights, she recognized that she had to demonstrate the value that these rights would confer. She explained the benefits to society that would be realized because of the granting of those rights. She argued for the education of women because it would, in turn, make them better wives and mothers, more able to both support themselves and raise smart, conscientious children.

Her thoughts, from her book *A Vindication of the Rights of Woman*, are a demonstration of second-order thinking:

> Asserting the rights which women in common with men ought to contend for, I have not attempted to extenuate their faults; but to prove them to be the natural consequence of their education and station in society. If so, it is reasonable to suppose that they will change their character, and correct their vices and follies, when they are allowed to be free in a physical, moral, and civil sense.[10]

Empowering women was a first-order effect of recognizing that women should have rights. But by discussing the logical consequences this empowerment would have on society—the second-order effects—

Wollstonecraft started a conversation that eventually resulted in what we now call feminism. Not only would women eventually get freedoms they deserved, they would become better women and better members of society.

A Word of Caution

Second-order thinking must be tempered in one important way: you can't let it lead to the paralysis of the "slippery slope effect," the idea that if we start with action A, everything after is a slippery slope down to hell, with an inevitable chain of consequences including B, C, D, E, and F.

Garrett Hardin smartly addresses this danger in *Filters Against Folly*:

> Those who take the wedge (Slippery Slope) argument with the utmost seriousness act as though they think human beings are completely devoid of practical judgment. Countless examples from everyday life show the pessimists are wrong…. If we took the wedge argument seriously, we would pass a law forbidding all vehicles to travel at any speed greater than zero. That would be an easy way out of the moral problem. But we pass no such law.[11]

In practical life, everything has limits. Even if we consider secondary and subsequent effects, we can only go so far. During waves of prohibition fever in the United States and elsewhere, conservative abstainers have frequently made the case that taking even a *first* drink would be the first step toward a life of sin. They're right: it's true that drinking a beer *might* lead you to become an alcoholic. But not most of the time.

Thus, we need to avoid the slippery slope and the analysis paralysis it can lead to. Second-order thinking needs to evaluate the most likely effects and their most likely consequences, checking our understanding of what the *typical* results of our actions will be. If we worried about all possible effects of the effects of our actions, we would likely never do anything, and we'd

be wrong. How you balance the need for higher-order thinking with practical, limiting judgment must be taken on a case-by-case basis.

Conclusion

Second-order thinking is a method of thinking that goes beyond the surface level, beyond the knee-jerk reactions and short-term gains. It asks us to play the long game, to anticipate the ripple effects of our actions and to make choices that will benefit us not just today, but in the months and years to come.

Second-order thinking demands we ask: And then what?

Think of a chess master contemplating her next move. She doesn't just consider how the move will affect the next turn, but how it will shape the entire game. She's thinking many steps ahead. She's considering not just her own strategy, but her opponent's likely response. This is second-order thinking in action.

In our daily lives, we're often driven by first-order thinking. We make decisions based on what makes us happy now, what eases our current discomfort or satisfies our immediate desires.

Second-order thinking asks us to consider the long-term implications of our choices, to make decisions based not just on what feels good now, but on what will lead to the best outcomes over time.

In the end, second-order thinking is about playing the long game. It's about making choices not just for the next move, but for the entire journey.

Probabilistic Thinking

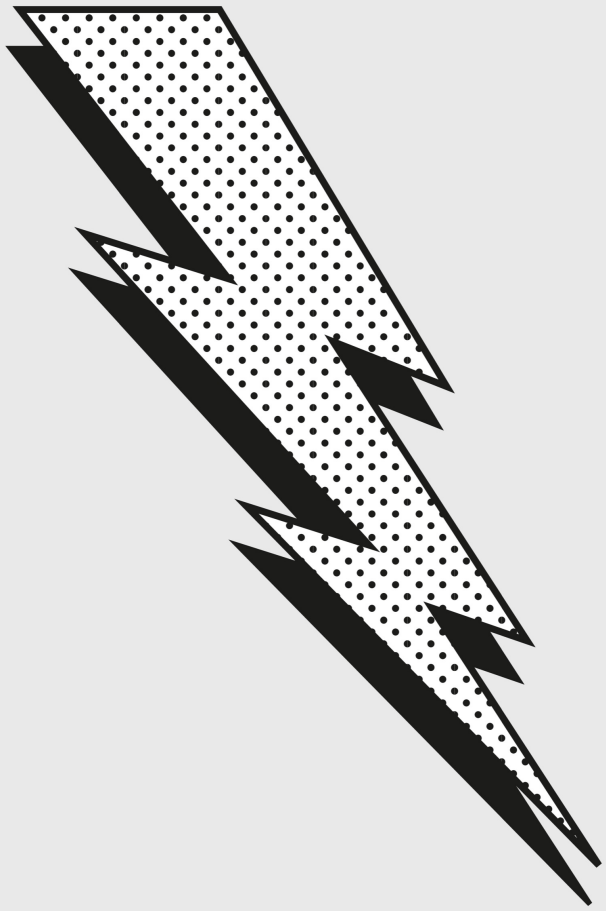

What are the chances?

The theory of probability is the only mathematical tool available to help map the unknown and the uncontrollable. It is fortunate that this tool, while tricky, is extraordinarily powerful and convenient.

—BENOIT MANDELBROT[1]

Probabilistic thinking is essentially trying to estimate, using some math and logic, the likelihood of any specific outcome occurring. It is one of the best tools we have to improve the accuracy of our decisions. In a world where each moment is determined by an infinitely complex set of factors, probabilistic thinking helps us deal with uncertainty. When we know these, our decisions can be more precise and effective.

Are You Going to Get Hit by Lightning or Not?

It's worth asking why we need to think in probabilities at all. Things either are or are not, right? Either we *will* get hit by lightning today or we *won't*. The problem is, we just don't know until we live out the day—which doesn't help us at all when we make our decisions in the morning about what to do. The future is far from predetermined, and we can better navigate it by understanding the likelihood of events that could impact us.

Very few things are 100 percent certain. Nearly everything is a probability. Our lack of perfect information about the world gives rise to all of probability theory, and to its usefulness. We know now that the future is inherently unpredictable, because not all variables can be known, and even the smallest error in our data very quickly throws off our predictions. The best we can do is estimate the future by generating realistic, useful probabilities. So how do we do that?

Probability is everywhere, down to the very bones of the world. The probabilistic machinery in our minds—the cut-to-the-quick "heuristics" made so famous by the psychologists Daniel Kahneman and Amos Tversky—was evolved by the human species in a time before computers, factories,

traffic, middle managers, and the stock market. It served us in a time when human life was about *survival* and still serves us well in that capacity.

Conditional Probability

Conditional probability is like Bayesian thinking (see below) in practice but comes at it from a different angle. When you use historical events to predict the future, you must be mindful of the conditions that surrounded that event.

Events can be independent, like tossing a coin, or dependent. A dependent event is one whose outcome is conditional on what preceded it. Let's say that the last three times I've hung out with you, we've gone for ice cream. I've picked vanilla each time. Do you conclude that vanilla is my favorite, and thus I will always choose it? You'd want to check first whether my choosing vanilla is independent or dependent. Am I the first to choose from among a hundred flavors? Or am I further down the line, when chocolate is no longer available?

My ice cream choice is independent if all the flavors are available each time someone in my group makes a choice. It is dependent if the preceding choices of my friends reduce what choices are available to me. In this case, the probability of my choosing vanilla is conditional on what is left after my friends make their choices.

Thus, using conditional probability means being very careful to observe the conditions preceding an event you'd like to understand.

But what about today—a time when, for most of us, survival is not so much the issue? Today, we want to *thrive*. We want to compete, and win. Mostly, we want to make good decisions in complex social systems that were not part of the world in which our brains evolved their (quite rational) heuristics.

To achieve these aims, we need to consciously add in a layer of probability awareness to our thinking.

What is probability awareness, and how can you use it to your advantage? There are three important aspects of probability that we need to explain so you can integrate them into your thinking, to get you into the ballpark and improve your chances of catching the ball:

1. Bayesian thinking
2. Fat-tailed curves
3. Asymmetries

Bayesian thinking: Thomas Bayes was an English minister in the first half of the eighteenth century, whose most famous work, "An Essay Toward Solving a Problem in the Doctrine of Chances," was brought to the attention of the Royal Society by his friend Richard Price in 1763—two years after his death. The essay, the key to what we now know as Bayes's Theorem, concerned how we should adjust probabilities when we encounter new data.

The core of Bayesian thinking (or Bayesian updating, as it can be called) is this: given that we have limited but useful information about the world, and are constantly encountering new information, we should consider what we already know—as much of it as possible—when we learn something new. Bayesian thinking allows us to use *all* relevant prior information in making decisions. Statisticians might call it a "base rate"—taking in outside information about past situations like the one you're in.

Consider the headline "Violent Stabbings on the Rise." Without Bayesian thinking, you might become genuinely afraid, because your

chance of being a victim of assault or murder is higher than it was a few months ago. But a Bayesian approach will have you putting this information into the context of what you already know about violent crime: You know that violent crime has declined to its lowest rates in decades. Your city is safer now than it has been since this measurement was started. Let's say your chance of being a victim of a stabbing last year was 1 in 10,000, or 0.01 percent. The article states, with accuracy, that violent crime has doubled. Your chance of being stabbed is now 2 in 10,000, or 0.02 percent. Is that worth being terribly worried about? The prior information here is key. When we factor it in, we realize that our safety has not really been compromised.

If we look at diabetes statistics in the United States, our application of prior knowledge would lead us to a different conclusion. Here, a Bayesian analysis indicates you should be concerned. In 1958, 0.93 percent of the population was diagnosed with diabetes. In 2015, it was 7.4 percent. When you look at the intervening years, the climb in diabetes diagnoses is steady, not a spike. So the prior relevant data, or "priors," indicate a trend that is worrisome.

It is important to remember that priors themselves represent probability estimates. For each bit of prior knowledge, you are not putting it in a binary structure, saying it is true or not. You're assigning it a *probability* of being true. Therefore, you can't let your priors get in the way of processing new knowledge. In Bayesian terms, this is called the "likelihood ratio" or the "Bayes factor." Any new information you encounter that challenges a prior simply means that the probability of that prior being true may be reduced. Eventually, some priors are replaced completely. Bayesian thinking is an ongoing cycle of challenging and validating what you believe you know. When making uncertain decisions, it's nearly always a mistake not to ask: What are the relevant priors? What might I already know that I can use to better understand the reality of the situation?

Fat-tailed curves: Many of us are familiar with the bell curve, that nice, symmetrical wave that captures the relative frequency of so many things from heights to exam scores. The bell curve is great because it's easy

to understand and easy to use. Its technical name is “normal distribution.” If we know we are in a bell curve situation, we can quickly identify our parameters and plan for the most likely outcomes.

Fat-tailed curves are different. Let’s take a look.

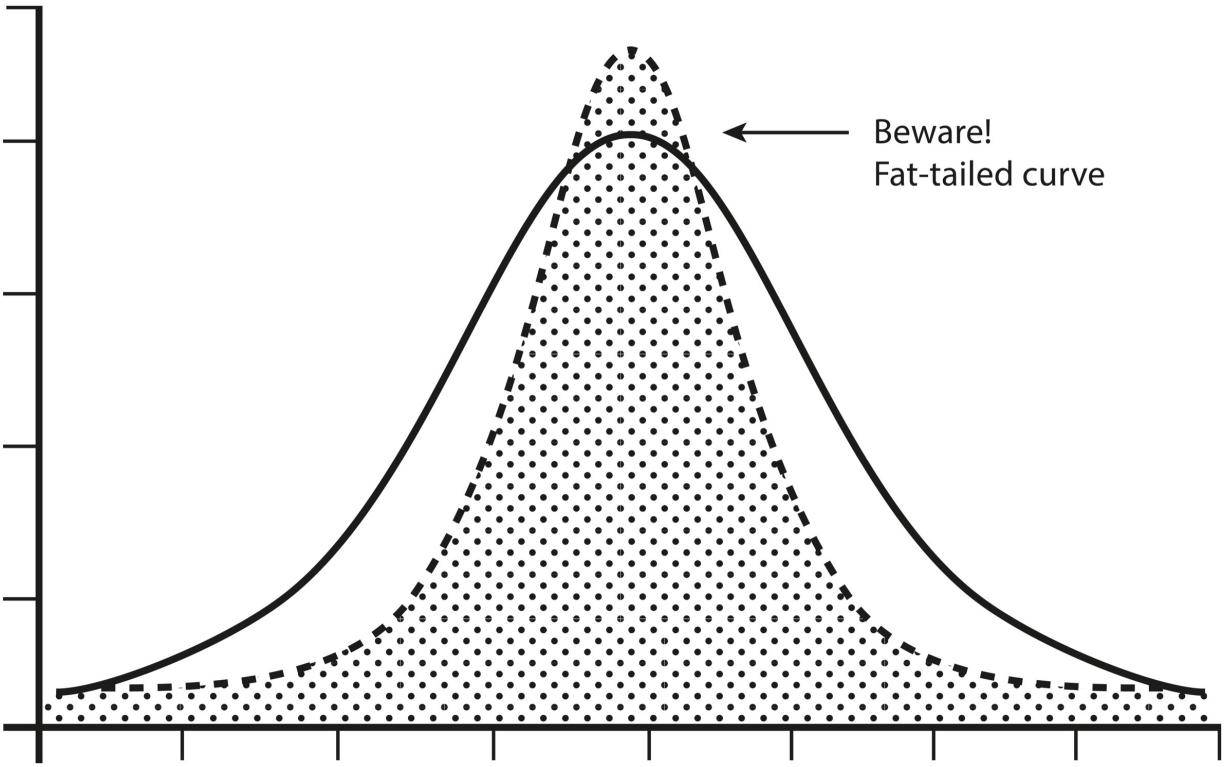

Always be extra mindful of the tails: They might mean everything.

At first glance, the two figures seem similar enough. Common outcomes cluster together, creating a wave. The difference is in the tails. In a bell curve, the extremes are predictable. There can only be so much deviation from the mean. In a fat-tailed curve, there is no real cap on extreme events.

The more extreme events that are possible, the longer the tails of the curve get. Any one extreme event is still unlikely, but the sheer number of options means that we can’t rely on the most common outcomes as representing the average. The more extreme events that are possible, the

higher the probability that one of them will occur. Crazy things are going to happen, and we have no way of identifying when.

Orders of Magnitude

Nassim Nicholas Taleb puts his finger in the right place when he points out our *naive* use of probabilities. In *The Black Swan*, he argues that any small error in measuring the risk of an extreme event can mean we're not just slightly off but *way* off—off by orders of magnitude, in fact. In other words, we're not just 10 percent wrong but ten times wrong, or a hundred times wrong, or a thousand times wrong. Something we thought could happen only once every thousand years might be likely to happen in any given year! Using false prior information results in us underestimating the probability of the future distribution being different.[2]

Think of it this way: In a bell curve situation, such as displaying the distribution of heights or weights in a human population, there are outliers on the spectrum of possibility, but the outliers have a fairly well-defined scope. You'll never meet a man who is ten times the size of an average man. But in a curve with fat tails, like wealth, the central tendency does not work the same way. You may regularly meet people who are ten, a hundred, or ten thousand times wealthier than the average person. That is a very different type of world.

Let's reapproach the example of the risk of violence we discussed in relation to Bayesian thinking. Suppose you heard that you had a greater risk of slipping on the stairs and cracking your head open than being killed by a terrorist. The statistics, the priors, seem to back it up: a thousand people slipped on the stairs and died last year in your country, and only five hundred died in terrorist attacks. Should you be more worried about stairs or terror events? Some people use examples like these to prove that terror risk is low—since the recent past shows very few deaths, why worry?[3] The problem is in the fat tails: The risk of terror violence is more like wealth, while stair-slipping deaths are more like height and weight. In the next ten years, how many events are possible? How fat is the tail?

The important thing is not to sit down and imagine every possible scenario in the tail (which, by definition, is impossible) but to deal with fat-tailed domains in the correct way: by positioning ourselves to survive or even benefit from the wildly unpredictable future, by being the only ones thinking correctly and planning for a world we don't fully understand.

Antifragility

How do we benefit from the uncertainty of a world we don't understand, one dominated by "fat tails"? Here, Nassim Nicholas Taleb's work is again instructive. In his book *Antifragile*, he explains it thus: We can think about three categories of objects—ones that are *harmed* by volatility and unpredictability, ones that are *neutral* to volatility and unpredictability, and ones that *benefit* from it.[4] The last category is antifragile—like a package that *wants* to be mishandled. Up to a point, certain things benefit from volatility, and that's how we want to be ourselves. Why? Because the world is fundamentally unpredictable and volatile, and large events—panics, crashes, wars, bubbles, and so on—tend to have a disproportionate impact on outcomes.

There are two ways to handle such a world: try to predict, or try to prepare. Prediction is tempting. For all of human history, seers and soothsayers have turned a comfortable trade. The problem is that nearly all studies of "expert" predictions in such complex real-world realms as the stock market, geopolitics, and global finance have shown again and again that, for the rare and impactful events in our world, predicting is impossible! It's more efficient to prepare.

What are some ways we can prepare—arm ourselves with antifragility—so we can benefit from the volatility of the world?

The first one is what Wall Street traders would call "upside optionality"—that is, seeking out situations that we expect to have good odds of offering us opportunities. Take the example of attending a cocktail party where a lot of people you might like to know are in attendance. While nothing is *guaranteed* to happen—you may not meet those people, and if you do, it may not go well—you give yourself the benefit of serendipity and randomness. The worst thing that can happen is...nothing. One thing you know for sure is that you'll never meet these people sitting at home. By going to the party, you improve your odds of encountering opportunity—your upside optionality.

The second thing we can do is to learn how to fail properly. Failing properly has two major components: First, never take a risk that will do you in—never get taken out of the game completely. Second, develop the personal resilience to *learn* from your failures and start again. If you follow these two rules, you can only fail temporarily.

No one likes to fail. It hurts. But failure carries with it one huge antifragile gift: learning. Those who are not afraid to fail (properly) have a huge advantage over the rest. What they learn makes them less vulnerable to the volatility of the world. They benefit from it, in true antifragile fashion.

Let's say you'd like to start a successful business, but you have no business experience. Do you attend business school, or start a business that might fail? Business school has its benefits, but business itself—the rough, jagged real-world experience of it—teaches through rapid feedback loops of success and failure. In other words, trial and error carries the precious commodity of information.

The *Antifragile* mindset is a unique one. Whenever possible, try to create scenarios where randomness and uncertainty are your friends, not your enemies.

Asymmetries: Finally, you need to think about something we might call "metaprobability"—the probability that your probability estimates themselves are any good.

This massively misunderstood concept has to do with asymmetries. If you look at nicely polished stock pitches made by professional investors, nearly every time an idea is presented, the investor looks their audience in the eye and states that they think they're going to achieve a rate of return of 20 to 40 percent per annum, if not higher. Yet *exceedingly* few of them ever attain that mark. It's not because they don't pick any winners—it's because they get so many so wrong. They are consistently overconfident in their probabilistic estimates. (For reference, the general stock market in the United States, over a long period, has returned no more than 7 percent to 8 percent per annum, before fees.)

Another common asymmetry is people's ability to estimate the effect of traffic on travel time. How often do you leave "on time" and arrive 20 percent early? Almost never? How often do you leave "on time" and arrive 20 percent late? All the time? Exactly. Your estimation errors are asymmetric, skewing in a single direction. This is often the case with probabilistic decision making.[5]

Far more probability estimates are wrong on the "over-optimistic" side than the "under-optimistic" side. You'll rarely read about an investor who aimed for 25 percent annual return rates and who subsequently earned 40 percent over a long period of time, whereas you can throw a dart at *The Wall Street Journal* and hit the names of lots of investors who aim for 25 percent per annum with each investment and end up closer to 10 percent.

The Spy World

Successful spies are very good at probabilistic thinking. High-stakes survival situations tend to make us evaluate our environment with as little bias as possible.

When Vera Atkins was second in command of the French unit of the Special Operations Executive (SOE), a British intelligence organization during World War II that reported directly to Winston Churchill,[6] she had to make hundreds of decisions by figuring out the probable accuracy of inherently unreliable information.

Atkins was responsible for the recruitment of British agents and their deployment into occupied France. She had to decide who could do the job and where the best sources of intelligence were. These were literal life-and-death decisions, and all were based in probabilistic thinking.

First, how do you choose a spy? Not everyone can go undercover in high-stress situations and make the contacts necessary to gather intelligence. The result of failure in France during the war was not getting fired; it was death. What factors of personality and experience show that a person is right for that job? Even today, with advancements in psychology, interrogation, and polygraph tests, it's still a judgment call.

For Vera Atkins, in the 1940s, it was very much a process of assigning weight to the various factors and coming up with a probabilistic assessment of who had a decent chance of success. Who spoke French? Who had the necessary confidence? Who was too tied to family? Who had the problem-solving capabilities? From recruitment to deployment, her development of each spy was a series of continually updated educated estimates.

Getting an intelligence officer ready to go is only half the battle. Where do you send them? If your information was so great that you knew exactly where to go, you probably wouldn't need an intelligence mission. Choosing a target is another exercise in probabilistic thinking. You need to evaluate the reliability of the information you have and the networks you have set up. Intelligence is not evidence. There is no chain of command or guarantee of authenticity.

The stuff coming out of German-occupied France was at the level of grainy photographs, handwritten notes that passed through many hands on the way back to headquarters, and unverifiable wireless messages sent quickly, sometimes sporadically, and with the operator under incredible

stress. When deciding what to use, Atkins had to consider the relevance, quality, and timeliness of the information she had.

She also had to make decisions based not only on what had happened but on what possibly could. Trying to prepare for every eventuality would mean that spies would never leave home, but they had to somehow prepare for a good deal of the unexpected. After all, a spy's job is often executed in highly volatile, dynamic environments. The women and men Atkins sent over to France worked in three primary occupations: organizers were responsible for recruiting locals, developing the network, and identifying sabotage targets; couriers moved information all around the country, connecting people and networks to coordinate activities; and wireless operators had to set up heavy communications equipment, disguise it, get information out of the country, and be ready to move at a moment's notice. All these jobs were dangerous. The full scope of the threats was never completely identifiable. There were so many things that could go wrong, so many possibilities for discovery or betrayal, that it was impossible to plan for them all. The average life expectancy in France for one of Atkins's wireless operators was six weeks.

Finally, the numbers suggest an asymmetry in the estimation of the probability of success of each individual agent. Of the four hundred agents that Atkins sent over to France, a hundred were captured and killed. This is not meant to pass judgment on her skills or smarts. Probabilistic thinking can only get you in the ballpark. It doesn't guarantee 100 percent success.

There is no doubt that Atkins relied heavily on probabilistic thinking to guide her decisions in the challenging quest to disrupt German operations in France during World War II. It is hard to evaluate the success of an espionage career, because it is a job that comes with a lot of loss. Atkins was extremely successful in that her network conducted valuable sabotage to support the Allied cause during the war, but the loss of life this work entailed was significant.

Conclusion

Probabilistic thinking is the art of navigating uncertainty. Successfully thinking in shades of probability means roughly identifying what matters, coming up with a sense of the odds, doing a check on our assumptions, and then deciding.

The challenge of probabilistic thinking is that it requires constant updating. As new information emerges, the probabilities change. What seemed likely yesterday may seem unlikely today. This both explains why probabilistic thinkers are always revising their beliefs with new data and why it's so uncomfortable for many people.

It's much easier to believe something false is true than deal with the fact that it might not be true. Being a probabilistic thinker means being willing to say, "I don't know for sure, but based on the evidence, I think there's a 63 percent chance of X."

The rewards of probabilistic thinking are immense. By embracing uncertainty, we can make better decisions, avoid the pitfalls of overconfidence, and navigate complex situations with greater skill and flexibility. We can be more open-minded, more receptive to new ideas, and more resilient in the face of change.

Insurance Companies

The most probability-acute businesses in the world are insurance companies—because they must be. When we think of insurance, we might think of life insurance (the probability of a policyholder dying at a certain age), or auto insurance (the probability of being in a car accident), or maybe home insurance (the probability of a tree falling on the house). With the statistics available to us, the probabilities of these things are easy to price and predict across a large enough population.

But insurance is deeply wide-ranging, and insurers will insure almost any event, for a price. Insurance policies have been taken out on Victoria's Secret models' legs, on baseball players' arms, on the Pepsi Challenge and the NCAA tournament, and even on a famous country singer's breasts!

How is this possible? Only with close attention to probability. What the great insurance companies in the world know how to do is pay attention to the important factors in a situation, even if they're not totally predictable, and price accordingly.

What is the probability of a Victoria's Secret model injuring her legs badly enough to end her career? One in 10,000? One in 100,000? Getting this calculation right would mean evaluating her lifestyle, her habits, her health, her family history—and coming up with a price and a set of conditions that are good enough, on average, to provide a profit. It's not unlike handicapping a race at the horse track. You can always say yes to insuring, but the trick is to come up with the right price. For that, we need probability.

SUPPORTING IDEA:

Causation vs. Correlation

CONFUSION BETWEEN THESE TWO TERMS often leads to a lot of inaccurate assumptions about the way the world works. We notice two things happening at the same time (correlation) and mistakenly conclude that one causes the other (causation). We then often act upon that erroneous conclusion, making decisions that can have immense influence across our lives. The problem is, without a good understanding of what is meant by these terms, these decisions fail to capitalize on real dynamics in the world and instead are rendered successful only by luck.

No Correlation

The correlation coefficient between two measures, which varies between −1 and 1, is a measure of the relative weight of the factors they share. For example, two phenomena with few shared factors, such as bottled-water consumption versus suicide rate, should have a correlation coefficient of close to 0. That is to say, if we looked at all countries in the world and plotted suicide rates in a specific year against per capita consumption of bottled water, the plot would show no pattern at all.

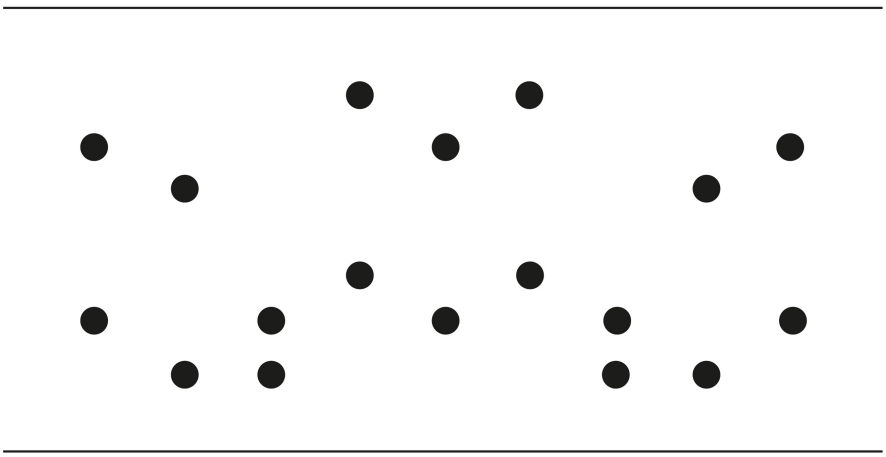

No Correlation

Perfect Correlation

By contrast, there are measures that are solely dependent on the same factor. A good example of this is temperature. The only factor governing temperature—velocity of molecules—is shared by all scales. Thus, each degree in Celsius will have exactly one corresponding value in Fahrenheit. Therefore, temperature in Celsius and Fahrenheit will have a correlation coefficient of 1, and the plot will be a straight line.

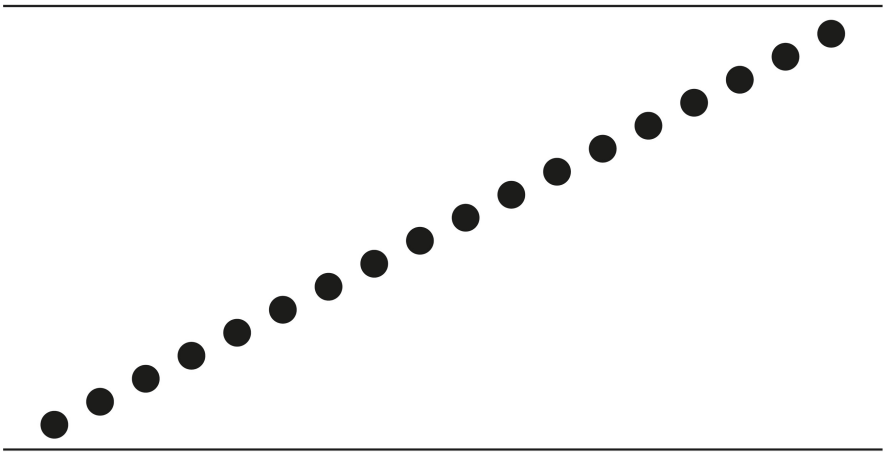

Perfect Correlation

Weak to Moderate Correlation

There are few phenomena in human sciences that have a correlation coefficient of 1. There are, however, plenty where the association is weak to moderate, and there is *some* explanatory power between the two phenomena. Consider the correlation between height and weight, which would land somewhere between 0 and 1. While virtually every three-year-old will be lighter and shorter than every grown man, not all grown men or three-year-olds of the same height will weigh the same.

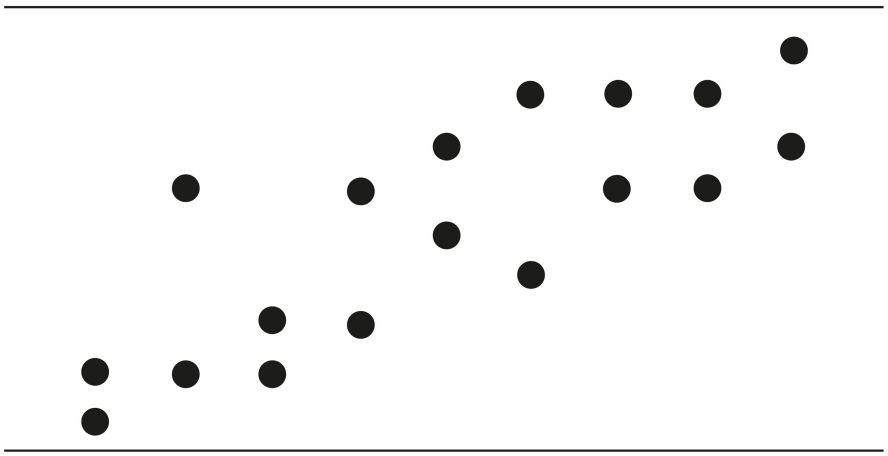

Weak to Moderate Correlation

This variation, and the corresponding lower degree of correlation, implies that, while height is a good predictor of weight, there clearly are factors other than height at play.

In addition, correlation can sometimes work in reverse. Let's say you read a study that compares alcohol consumption rates in parents and their children's corresponding academic success. The study shows a relationship between high alcohol consumption and low academic success. Is this causation or correlation? It might be tempting to conclude there's causality, such as the more parents drink, the worse their kids do in school.

However, this study has demonstrated only a *relationship*, not proved that one causes the other. The factors correlate—meaning that alcohol consumption in parents has an inverse relationship with academic success in children. It is entirely possible that having parents who consume a lot of alcohol leads to worse academic outcomes for children. It is also possible, however, that the reverse is true, or even that having kids who do poorly in school causes parents to drink more. Trying to invert the relationship can help you sort through claims to determine if you are dealing with true causation or just correlation.

Causation

Whenever correlation is imperfect, extremes will soften over time. The best will always appear to get worse, and the worst will appear to get better, regardless of any additional action. This is called regression to the mean, and it means we have to be extra careful when diagnosing causation. This is something that the general media, and sometimes even trained scientists, fail to recognize.

Consider the example Daniel Kahneman gives in *Thinking, Fast and Slow*:[7]

> Depressed children treated with an energy drink improve significantly over a three-month period. I made up this newspaper headline, but the fact it reports is true: if you treated a group of depressed children for some time with an energy drink, they would show a clinically significant improvement. It is also the case that depressed children who spend some time standing on their head or hug a cat for twenty minutes a day will also show improvement.

Whenever we come across such headlines, it is very tempting to jump to the conclusion that energy drinks, standing on the head, or hugging cats are all perfectly viable cures for depression. These cases, however, once again embody the concept of regression to the mean:

> Depressed children are an extreme group—they are more depressed than most other children—and extreme groups regress to the mean over time. The correlation between depression scores on successive occasions of testing is less than perfect, so there will be regression to the mean: depressed children will get somewhat better over time even if they hug no cats and drink no Red Bull.

We often mistakenly attribute a specific policy or treatment as the cause of an effect, when the change in the extreme groups would have happened anyway. This presents a fundamental problem: How can we know if the effect is real or simply due to variability?

Luckily, there is a way to tell between a real improvement and something that would have happened anyway. That is the introduction of the so-called “control group,” which is expected to improve by regression alone. The aim of the research is to determine whether the treated group improves more than regression can explain.

In real-life situations assessing the performance of specific individuals or teams, where the only real benchmark is past performance and no control group can be introduced, the effects of regression can be difficult, if not impossible, to disentangle. We can compare against industry average, peers in the cohort group, or historical rates of improvement, but none of these is a perfect measure.

Inversion

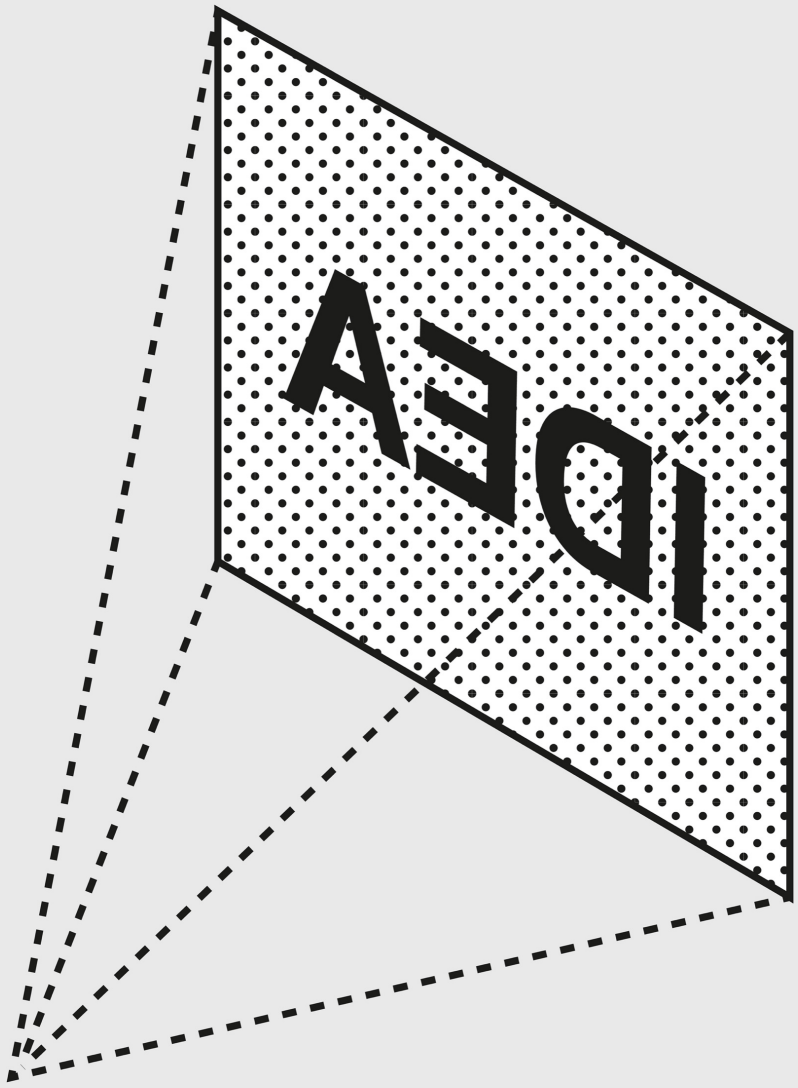

Change your perspective.

The test of a first-rate intelligence is the ability to hold two opposing ideas in mind at the same time and still retain the ability to function. One should, for example, be able to see that things are hopeless yet be determined to make them otherwise.

—F. SCOTT FITZGERALD[1]

It can be difficult to appreciate just how much avoiding the standard ways of failing dramatically increases the odds of success.

Inversion is all about identifying and removing the obstacles to success. The root of inversion is “invert,” which means to upend or turn upside down. As a thinking tool, it means approaching a situation from the opposite end of the natural starting point.

Most of us tend to think one way about a problem: forward. Inversion allows us to flip the problem around and think backward. Sometimes it’s good to start at the beginning, but it can be more useful to start at the end.

Avoiding stupidity is easier than seeking brilliance. Even when we don’t know how to achieve a particular objective, we can often identify what prevents it from happening. Perhaps you don’t know all the things that create a good night’s sleep. We can invert the problem by identifying some of the standard things that prevent us from getting a good night’s sleep, such as eating right before going to bed or consuming a lot of alcohol. Simply avoiding those two things dramatically improves the quality of our sleep.

Avoiding the Standard Ways of Failing

Warren Buffett and Charlie Munger are two of the most successful investors of all time. Their track record at Berkshire Hathaway is legendary. One underappreciated aspect of their success is how they avoided the standard ways of failing.

When asked about how otherwise smart people fail, the duo commented that it often involves drugs and leverage. “It’s insane to risk what you have and need for something you don’t really need,” Buffett says of borrowing. “You will not be way happier if you double your net worth.”

In one of his last interviews, Munger commented that one of the keys to success in life is avoiding common traps:

“My game in life was always to avoid all standard ways of failing. You teach me the wrong way to play poker and I will avoid it. You teach me the wrong way to do something else, I will avoid it. And, of course, I’ve avoided a lot, because I’m so cautious.

“Crazy is way more common than you think,” Munger said. “It’s easy to slip into crazy. Just avoid it, avoid it, avoid it.”[2]

Inversion teaches us that a great deal of wisdom can be found in knowing what to avoid.

There are two approaches to applying inversion in your life:

1. Start by assuming that what you're trying to prove is either true or false, then show what else would have to be true to make that so.
2. Instead of aiming directly for your goal, think deeply about what to *avoid* and then see what options are left over.

Set your assumptions: The nineteenth-century German mathematician Carl Jacobi became famous for several reasons—including solving some incredibly difficult problems—but is perhaps best remembered for his advice to "invert, always invert."[3]

Jacobi solved a range of problems by starting with the endpoint. When faced with proving an axiom in a difficult math problem, he might instead assume that a property of the axiom was correct and then try to determine the consequences of this assumption. From that point, he could work out surprising, and at times counterintuitive, insights.

Jacobi was not the first mathematician to use inversion. In fact, inversion is a staple of mathematical, philosophical, and scientific inquiry. We can look around today and appreciate that we can't see atoms and quarks, but we know they exist because we can make predictions about their behavior and test those predictions.

Or, we can go back 2,300 years and look at the work of the Greek mathematician Hippasus, a follower of Pythagoras.[4] (Yes, the one with the theorem.) His attempts to derive the square root of two, and his original direct approach to solving the problem (essentially, dividing larger and larger whole numbers into each other) were both fruitless and time consuming. He hit an impasse, realizing that he'd never be able to definitely solve the problem by thinking forward. In his increasing frustration, Hippasus decided to take the reverse route, thinking about what the square root of two might *imply*, and working backward from there. If he couldn't find it the way he had expected to, he'd start by proving what the number

couldn't be. His quest forever changed what we understood about mathematics and led to the discovery of the first irrational number.

Mathematics is not the only area where using inversion can produce surprising and nonintuitive results. In the 1920s, the American Tobacco Company wanted to sell more of their Lucky Strike cigarettes to women. Men were smoking, but women weren't. There were pervasive taboos against women smoking—it was seen as a man's activity. Women, therefore, presented an untapped market that had the potential to provide huge revenue. Riding on the slimness trend that had already begun, the head of the company thought that they needed to convince women that smoking would make them thinner, so he hired Edward Bernays, who came up with a truly revolutionary marketing campaign.[5],[6]

In the style of the inversion approach described above, Bernays did not ask, "How do I sell more cigarettes to women?" Instead, he wondered, if women bought and smoked cigarettes, what else would have to be true? What would have to change in the world to make smoking desirable to women and socially acceptable? Then—a step further—once he knew what needed to change, how would he achieve that?

To tackle the idea of smoking as a slimming aid, Bernays mounted a large anti-sweets campaign. After dinner, it was all about cigarettes, not dessert. Cigarettes were slimming, while desserts would ruin one's figure. But Bernays's real stroke of genius did lie solely in coming out with advertisements to convince women to stay slim by smoking cigarettes. As author Alan Axelrod puts it, "Instead, he sought nothing less than to reshape American society and culture."[7] He solicited journalists and photographers to promote the virtues of being slim. He sought testimonials from doctors about the health value of smoking after a meal. He combined this approach with "altering the very environment, striving to create a world in which the cigarette was ubiquitous." Axelrod details the full scope of Bernays's efforts:

> He mounted a campaign to persuade hotels and restaurants to add cigarettes to dessert-list menus, and he provided such magazines as

> *House and Garden* with feature articles that included menus designed to preserve readers "from the dangers of overeating."… The idea was not only to influence opinion but to remold life itself. Bernays approached designers, architects, and cabinetmakers in an effort to persuade them to design kitchen cabinets that included special compartments for cigarettes, and he spoke to the manufacturers of kitchen containers to add cigarette tins to their traditional lines of labeled containers for coffee, tea, sugar, and flour.[8]

The result was a complete shift in the consumption habits of American women. It wasn't just about selling the cigarette; it was about reorganizing society to make cigarettes an inescapable part of the American woman's daily experience.

Bernays's efforts to make women's smoking in public socially acceptable had equally startling results. He linked cigarette smoking with women's emancipation: to smoke was to be free. Cigarettes were marketed as "torches of freedom." He orchestrated public events, including an infamous parade on Easter Sunday in 1929, which featured women smoking as they walked in the parade. He left no detail unattended, so that public perception of smoking was changed almost overnight. He both normalized it and made it desirable in one swoop.

Although the Lucky Strike campaign utilized more principles than just inversion, it was the original decision to invert the approach that provided the framework from which the campaign was created and executed. Bernays didn't focus on how to sell more cigarettes to women within the existing social structure. If he had, undoubtedly sales would have been a lot more limited. Instead, he thought about what the world would look like if women smoked often and anywhere, and then set about trying to make that world a reality. Once he did that, selling cigarettes to women was comparatively easy.

This inversion approach became a staple of Bernays's work. He used the descriptor "appeals of indirection," and each time he was hired to sell a

product or service, “he instead sold whole new ways of behaving, which appeared obscure but over time reaped huge rewards for his clients and redefined the very texture of American life.”[9]

Decide what to avoid: Instead of thinking through the achievement of a positive outcome, another way to use inversion is to ask ourselves how we might achieve a *terrible* outcome, and let that guide our decision making.

Index funds are a great example of stock market inversion, promoted and brought to bear by Vanguard’s John Bogle.[10] Instead of asking how to beat the market, as so many before him had, Bogle simply recognized the difficulty of the task. Everyone is trying to beat the market. No one is doing it with any consistency, and in the process real people are losing actual money. So Bogle inverted the approach. The question then became, how can we help investors minimize losses to fees and poor money manager selection? The results were one of the greatest ideas—index funds—and one of the greatest powerhouse firms in the history of finance.

Index funds operate on the idea that accruing wealth has a lot to do with minimizing loss. Think about your personal finances: Often, we focus on positive goals, such as “I want to be rich,” and use this to guide our approach. We make investing and career choices based on our desire to accumulate wealth. We chase after magical solutions, like attempting to outsmart the stock market. These inevitably get us nowhere, and we have usually taken some terrible risks in the process that leave us worse off.

Inverting our approach, we can instead ask ourselves what the common pitfalls in investing are and how we can avoid them. For example, spending more than we make, taking on too much leverage (or paying high interest rates on debt so that we can’t tackle paying back the principal), and not starting to save as early as we can so as to take advantage of the power of compounding are all concrete financial behaviors that cost us money. We can more readily secure wealth by using inversion to make sure we are not doing the worst things that prevent the accumulation of wealth.

Guarantee a Life of Misery

In one of the more unique graduation speeches ever delivered, Johnny Carson offered the Harvard School some peculiar life advice. While he couldn't tell the graduating class how to be happy, Carson inverted and offered three guaranteed prescriptions for misery:

1. Ingesting chemicals in an effort to alter mood or perception
2. Envy
3. Resentment

Let's briefly explore each.

"Ingesting chemicals to alter mood or perception" refers to the misuse of substances like drugs or alcohol to change one's mental state. If you want to spiral your life out of control, turn to alcohol or other substances for stress relief.

Envy involves feeling discontent or resentful longing aroused by someone else's possessions, qualities, or luck. To ensure a perpetual state of dissatisfaction, constantly compare yourself to others who have more.

Resentment is a failure to let go of anger or bitterness toward someone due to a past slight or injustice. If you want to poison your personal relationships and mental health, hold on tightly to tiny slights.

In a later speech, Charlie Munger added to Carson's prescriptions to ensure a life of misery.

"First," he offered, "be unreliable. Do not faithfully do what you have engaged to do. If you will only master this one habit you will more than counterbalance the combined effect of all your virtues, howsoever great." Being unreliable is a sure way to be excluded and distrusted.

Second on the list was "not learning from others' mistakes." Mastering this habit ensures you make every mistake possible. In the process, you will be surpassed by those who master the best of what other people have figured out.

Finally, Munger offered, "go down and stay down." A surefire way to fail is to quit trying.[11]

One of the theoretical foundations for this type of thinking comes from psychologist Kurt Lewin.[12] In the 1930s, he came up with the idea of "force field analysis," which essentially recognizes that in any situation where change is desired, successful management of that change requires applied inversion. Here is a brief explanation of his process:

1. Identify the problem.
2. Define your objective.
3. Identify the forces that support change toward your objective.
4. Identify the forces that impede change toward the objective.
5. Strategize a solution! This may involve both augmenting, or adding to, the forces in step 3 and reducing or eliminating the forces in step 4.

Even if we are quite logical, most of us stop after step 3. Once we figure out our objective, we focus on the things we need to put in place to make it happen—the new training or education, the messaging and marketing. But Lewin theorized that it can be just as powerful to remove obstacles to change.

The inversion happens between steps 3 and 4. Whatever angle you choose to approach your problem from, you need to then follow up with consideration of the opposite angle. Think about not only what you could do to solve a problem but what you could do to make it worse—and then avoid doing that, or eliminate the conditions that perpetuate it.

This inversion approach was used by Florence Nightingale to help significantly reduce the mortality rate of British soldiers in military hospitals in the mid-nineteenth century. Nightingale is often remembered as the founder of modern nursing, but she was also an excellent statistician and was the first woman elected to the Royal Statistical Society, in 1858.

During the first winter of the Crimean War, 1854–55, the British Army endured a death rate of 23 percent. The next winter that rate had dropped to

2.5 percent.[13] The main reason for the change was a much better understanding of what was actually killing the soldiers, an understanding that rested on the detailed statistics that Florence Nightingale collected. She demonstrated that the leading cause of death by far was poor sanitation. In her famous polar-area chart—a completely new way of presenting data at the time—she captured a visual representation of the statistics that made them easy to understand. Improve the sanitary conditions in the hospitals, she explained, and many soldiers' lives would be saved.

Nightingale's use of statistics helped to identify the real problem of army-hospital deaths. She was able to demonstrate not only what the army could do to improve outcomes but, just as important, what they had to avoid doing to stop making things worse. She reflected on the knowledge that could be derived from statistics and, in another instance of inversion thinking, she advocated for their use as a means of prevention.[14] The question became not so much "how do we fix this problem?" but "how do we stop it from happening in the first place?" Nightingale took the knowledge and experience she gained in Crimea and began gathering statistics not just for British Army field hospitals but for domestic ones as well. She demonstrated that unsanitary conditions in military hospitals were a real problem causing many preventable deaths.[15]

Nightingale's advocacy for statistics ultimately went much further than British military hospitals. But her use of statistics to improve sanitary conditions can be seen as an example of applied inversion. She used them to advocate for both solving problems and the invert, preventing them.

> Hence to fight and conquer in all your battles is not supreme excellence; supreme excellence consists in breaking the enemy's resistance without fighting.
>
> **—SUN TZU**[16]

Conclusion

A lot of advantage is gained simply by avoiding the standard paths to failure.

Inversion is not the way we are taught to think. We are taught to identify what we want and explore things that will move us closer to our objective. However, by spending time identifying things that will ensure we *don't* get what we want, we dramatically increase our odds of success.

Often, we get so fixated on solving a problem in a particular way that we miss simpler, more elegant solutions. Inversion forces us to consider the opposite side of the equation.

Instead of asking, "How do I solve this problem?" inversion asks, "What would guarantee failure?" Instead of asking, "How can I achieve this goal?" it asks, "What is preventing me from achieving it?" By inverting the question, we can gain insights that our normal thought patterns might miss.

The next time you're grappling with a difficult problem or striving toward an ambitious goal, try inverting your thinking. Ask yourself how you could guarantee failure. The answers may surprise you and open up new avenues for possible solutions.