



At War with Luck

WRITTEN IN BLOOD

In the words of the nineteenth-century German philosopher Friedrich Nietzsche, thus spoken by his ancient Persian prophet Zarathustra, “Of all that is written, I love only what a person hath written with his blood.”

If so, Nietzsche would have loved this book.

It was written with the blood of war against luck, fought over the last more than a quarter century in my life as a trader. It grew organically out of an investing and risk-mitigation practice as a hedge fund manager and professional *safe haven investor*. The message of this book has been and will always be lived by me and my hedge fund firm, Universa Investments. (It is our manifesto.)

Talk is cheap. Ideas and commentary are just that. Significance only comes from the doing, from action within the arena. It is not my business, like Sherlock Holmes, “to know what other people do not know.” It is my business to *do* what other people do not and cannot *do* (as well as, just as importantly, to know what *I* do not know). Doing and demonstrating effective

safe haven investing is far, far more important than arguing about what it should be. And even among most of those who claim to do it, they neglect those pithy words from Hemingway to “never confuse movement with action.”

This book tells of the foundation and methodology behind how, as of this writing, Universa risk-mitigated portfolios have, over their decade-plus life to date, outperformed the S&P 500 by over 3% on an annualized, net basis. More to the point, this performance is a *direct consequence* of having *far less risk*. This level of outperformance is rare in the hedge fund industry and among risk-mitigation strategies in general, which have pretty much all underperformed the S&P 500 during this and most periods. The markets have been good to us because we haven’t tried to cheat them; we haven’t tried to predict or outsmart them. We have only aligned our investing, in a focused way, with our beliefs about the way they work.

People think of risk mitigation as a liability, as a trade-off against wealth creation, because it usually is. Universa is, if nothing else, a real-life case study and out-of-sample test that unequivocally proves the point that risk mitigation doesn’t have to be viewed that way. Risk mitigation can and should be thought of as being additive to portfolios over time—with the right risk mitigation, that is. This is the mark that I want Universa to make in the markets.

Writing this book has been a labor of love, though it has had a difficult time competing for my attention, which is consumed by Universa. But the book has provided very important opportunities for introspection. It has also made me think more deeply about questions that I am always asked by people about what small “lay” investors can do to protect their portfolios.

As this book is, in part, a response to those questions, I do want to ensure that expectations are set appropriately at the start. This is not a “how-to” book, but it is a “why-to” as well as a “why-not-to” book. Let’s be clear: What I do specifically as a safe haven investor is not to be attempted by nonprofessionals

(nor—perhaps even especially—by most professionals). Nothing that I could tell you in a book will change that.

So, I will not be holding your hand and teaching you how to do it; I will not be revealing much in the way of trade secrets, and I have no interest in selling you anything as an investment manager. This book is not about the workings of a specific safe haven strategy, *per se*; nor is it an encyclopedic survey of all the major safe haven investments. Moreover, it has little if any current market commentary—as this would be entirely unnecessary to the book's point.

That said, there are big investing problems to be solved, and I do explicitly employ at Universa their conclusive solution that will be found over the course of this book. Surely, that's a good thing. After all, if it's such a great solution, wouldn't it be a big red flag if I didn't put it into practice?

My intention in this book is to present the basic concepts underlying my approach in a straightforward manner. One-eighth of the iceberg is above water, and that's all we need. (While I am forced into the mathematical weeds at times, believe me, I tried my best to stay out of them.) I offer a logical and practical analytical framework to demystify safe havens, for viewing and thinking about their value in mitigating systematic risk, and what that even means in the first place. That is to say, I offer a framework for rigorously testing hypotheses about safe havens and their very existence.

If the only things readers get out of this book are a more realistic and rational premise of safe haven investing and a foundation from which to assess and tackle it—and thus avoid its traps—then the book will have achieved its purpose. It should help the reader and pay back their expense of buying this book manyfold. You will make money investing and you will lose money investing, but, when you look back at it all, what really will have mattered is getting that foundation right.

Investing is a deceptive, uneven playing field for many reasons, and I aim to help level it out a little. (To show you the

honor of my intentions, I pledge all of my proceeds from the sales of this book to charity.)

PROBING BETS

This book is the result of a discovery process and problem-solving effort that began many years ago. I grew up a scrappy, poor kid in the Chicago commodities trading pits, back when they were the center of the financial universe. As a mere teenager, I learned from my mentor, Everett Klipp (“the Babe Ruth of the Chicago Board of Trade”), who told me over and over that “a small loss is a good loss,” and that risk mitigation and survival are everything in trading and investing. His words still ring true today: Take care of the losses; the profits will then take care of themselves. Profits matter only relative to the losses; stay in the game by protecting your capital base, your means of playing the game. Don’t predict.

Pretty obvious stuff, except people don’t really focus on losses, especially the potential for big ones. In all my experience, most investors don’t think about the impact of the downside the way they need to; they just don’t.

For me, it began with hands-on, trial-and-error experiments in my twenties as a “local” (or independent) trader in the bond pit in Chicago. (I earned much of my stake in college by writing and then selling the first portfolio management handheld computer program on the trading floor.) It continued in my thirties as a bank proprietary and hedge fund derivatives trader in New York, and even a bit in the hallowed halls of New York University’s Courant Institute of Mathematical Sciences. I hashed it out in endless discussion-walks through the streets of Manhattan with Nassim Nicholas Taleb, in the days when we started the very first formal tail hedging program, and he would later become Scientific Advisor at Universa. These ideas were the point of every conversation during countless,

reckless longboard skateboard cruises up and down the treacherous hills and dales of Central Park with Brandon Yarckin as we plotted the launch of Universa, where he would become COO. (I still blame Brandon for my separated shoulder.)

And it goes without saying that I have not made these observations from scratch. I owe an enormous gratitude to the explorers who plumbed the depths (“the triad”).

These were the formative years of probing bets, of bold conjectures and refutations, and finding answers to how things worked through success and failure—of a bloody war against luck. The conclusion was my decision to start Universa in 2007, where we have taken our discoveries and solutions and transformed them into a formal, rational, and practical investment approach that solves the big problems and *moves the needle* for my investment partners.

By opening this book and, I would hope, opening your mind to the ideas it contains, you will have the opportunity to change your perceptions about investing and risk mitigation. These pages will present a fundamentally different, unconventional perspective that I firmly believe is the most effective one for successful investing. It opposes so much of what is hailed as gospel within the collective that is the investment profession. So, we need to think objectively about the orthodoxy and dead-end dogma of modern finance. We need to be the uncarved block.

To be sure, I have no interest in changing the existing investing paradigm by making it obsolete. I’m no pied piper, and I harbor no delusions of grandeur that the message of this book will ever become “the new gospel according to Mark.” I do not expect for one moment that my approach will ever become another cookie-cutter, rubber-stamped strategy of the investment consultant herd. And that’s important. Becoming conventional is self-defeating in this business. It’s the kiss of death. We take the road less traveled by, and that has made all the difference.

I have often exclaimed to my team at Universa, plundering the words of Steve Jobs, “We are pirates! Not the Navy!” (Mine are the smartest, savviest, and most experienced Great Pirates in the derivatives world.)

WHAT’S A SAFE HAVEN?

Everyone has some intuitive understanding about what a safe haven is and why we would invest in one. Most likely, it would go something like “a place of refuge for when things go bad” or, more specifically, “an asset that provides safety from risk.” These are spot on. The term *risk* would likely mean scary things like stock market crashes, financial and banking crises, pandemics, monsters hiding under the bed, and so forth. Risk, then, is circularly defined as that equivocal thing we need safety from.

Plain and simple, risk is exposure to bad contingencies. Most of these bad contingencies will likely *never* happen, but they *can* happen. In investing, the bad contingencies have financial consequences of economic loss within a portfolio. Investment risk isn’t just some theoretical and spurious numerical value, like *volatility* or *correlation* or whatever. It is the potential for loss, and the scope of that loss. Nothing else.

We have many diverging roads ahead, many potential paths forward with so many twists and turns, more than we can ever count. Some of those potential paths will be very pleasant, and some of those potential paths won’t be. Of all the potential paths, we don’t know which is the one and only path that we will actually traverse. Now, that’s risk!

A safe haven is an investment that mitigates risk, or the bad potential economic contingencies in your investment portfolio. That is a necessary condition for safe haven status. It protects against consequential loss that happens to everyone, everywhere, all at the same time, because that loss is tied to the cycles of broad macroeconomic growth and contraction. Because of

its ubiquitous and systematic nature, you can't just diversify that risk away with groupings of things that supposedly won't experience such loss simultaneously.

And remember this: A safe haven isn't so much a thing or an asset. It is a *payoff*, one that can take many different forms. It might be a chunk of metal, a stock selection criterion, a cryptocurrency, or even a derivatives portfolio. **Whatever forms they may take, it is their function that makes safe havens what they are: They preserve and protect your capital. They are a shelter from financial storms.**

So *safe haven investing* is *risk mitigation*. To me, these two terms are synonymous, and I will be using them interchangeably throughout this book. (The former made for a catchier title.)

What's more, risk mitigation *is* investing *itself*. Treat this as a fundamental premise. Even the most renowned proponent of the bottom-up approach to investing, Benjamin Graham, the "father of value investing," declared: "The essence of investment management is the management of risks, not the management of returns. Well-managed portfolios start with this precept." Moreover, "Confronted with a challenge to distill the secret of sound investment into three words, we venture the motto, Margin of Safety." Truer words have never been written on the subject. For Graham, "safety of principal" is what separates investing from speculating. *It makes investing—investing!* (He learned the hard way, in the 1929 stock market crash, that all great ideas can be dashed, simultaneously and systematically.)

But this is the pretty obvious part. It still misses what's so special about what a safe haven should be. Any punter can devise something that does well in a financial crash. Safe haven investing is so much more. And this is where they start to look really different from each other—so different that our classification that would put them all in the same category starts to lose its meaning. Often, they are more different than they are alike. We run into issues similar to what biologists encountered when trying to define what a species is. We will need a specific *safe haven*

concept—like their *species concept*—in order to classify them and evaluate if they are even what they claim to be. **This will be a major line of inquiry in this book: Are safe havens a classifiable thing? And can they really add economic value?**

THE GREAT DILEMMA

There is a monumental problem facing investors, *the great dilemma of risk*. If you take too much risk, it will likely cost you wealth over time. And at the same time, if you don't take enough risk, it will also likely cost you wealth over time. You are trapped in a "Catch-22": damned if you do and damned if you don't. Pick your poison. You can try calibrating between these two bad choices in hopes of finding a happy medium, but this still leaves you with a bad choice—it is still poison. Modern finance is really all about the quest for this theoretical happy medium, the supposed "Holy Grail of investing." Despite this valiant quest and lofty name, the results have shown that this Holy Grail is a myth; the happy medium is not so happy, and it can even provide the *worst* of both worlds. And so, with this thinking, you are left with only one real choice: to make bold predictions and then roll the dice on them.

This great dilemma is the most important problem in all of investing; it is one that desperately needs solving. It is also the rationale for what I do as a hedge fund manager, and the reason for this book. Thanks to the sheer scale and scope of the problem, the stakes have never been higher. In particular, the broader problem is of liabilities exceeding assets, and this applies to mammoth pools of capital and even to individuals with small investment accounts. Just think of the massively underfunded public and private pension funds today, which must generate specific, high-target rates of return over many years or else face insolvency as their liabilities consume their capital. They can't just hide away, idling in less risky assets in an impaired portfolio, or try to diversify away their risks; and yet investing in riskier

assets brings, by definition, acute risks of unrecoverable loss. The standard approach to risk mitigation has really failed them—just as it has failed everyone. And the problem is only going to get worse. It is a looming, ticking time bomb.

The consequences of failing to solve the great dilemma of risk won't just appear as some abstract figures in the newspaper. They are all too real: people's savings wiped out, governments that must tax or inflate their economies to death—human tragedy with real economic consequences. This is not my opinion. It is just simple math.

This monumental problem is further complicated today by the massive distortions built up in global financial markets from years of hubristic monetary interventions by global central banks, enabling the reckless accumulation of debt and leverage. Though these distortions are on an unprecedented scale and are intricately related to the underfunding problem itself, they are nonetheless beside the point. They are both beyond the scope of this book (I have already written plenty about them elsewhere) and, most importantly, completely unnecessary to the book's message. I don't need to convince you of any ideological, Cassandra-like premise that markets are risky so that you will accept my conclusions about safe havens. It will not matter to our methodology. We can and will remain agnostic, not roll the dice, and, most important, *not predict*.

To find a solution to this monumental problem, we need to reduce the costliness of risk—specifically the costliness of losses—and do so in a way that does not end up costing us even more. In other words, **we need a cure that is not worse than the disease. Risk mitigation must be cost-effective.**

In this book, we will learn how finding that solution comes from recognizing that not all risks are created equal—because not all losses are created equal. They don't all add up cleanly in an accounting ledger. Therefore, we need to think about losses and our investment returns differently, through a different lens and a different framing.

As in all things, “the good God is in the details.” And in our case, these details, while not terribly complicated, often appear counterintuitive and paradoxical. As we will see, there are *emergent* dynamics at play here that make cost-effective risk mitigation extremely challenging, perhaps more so than anything else in the realm of investing. We need to proceed cautiously.

The problem is that investing is approached by most professionals and academics (and even the reigning PhD quants of modern finance) in a highly *reductionist* way. But, as we will see throughout this book, in safe haven investing the whole is, indeed, not the same as the sum of its parts—and it is often much greater.

Cost-effective safe haven investing will turn out to be an awesome variation on the theme from my first book, *The Dao of Capital*. It’s an idea that has perhaps become my shibboleth: *roundabout* investing. That is, the indirect approach, seeming to go backwards in order to go forwards, as in Sun Tzu’s and von Clausewitz’s approaches to “lose the battle to win the war.” What will look like a bad idea for one roll of the dice strangely becomes the best idea over many.

But how is successful investing possible without predictions? It sounds too good to be true. That investing is about forecasting returns is a tenet of the industry. As such, most people think they need to look very far ahead in investing and risk mitigation—with a magic crystal ball; they think that they need to see around corners. Not only is that pretty much impossible, it is actually a misconception about investing. **Investing really needn’t be about making grandiose forecasts**, any more than it is in, say, sports or other games like poker or backgammon—though one could easily make that mistaken assumption from the outside looking in. It isn’t even necessarily about getting the probabilities right. You can get the probabilities right all day but still do very poorly. It’s really about *getting the payoffs right*. Playing good defense that leads to good offense. So there’s more room for error, more room for being

right, more room to get it right after getting it wrong. This is cost-effective risk mitigation. You look and feel like you can see around corners, even though, in actuality, you can't.

As an archer, you don't try to forecast or pinpoint exactly where your arrow will hit once it leaves your bow. That would be an unproductive way to approach it—leading to *target panic*. Once you shoot the arrow (and even *as* you shoot the arrow), it is out of your control and susceptible to endless perturbations. So, instead, as in Herrigel's *Zen in the Art of Archery*, you aim by deliberately *not* taking aim—you hone your process and structure (focusing “behind the line” rather than down range) with the intent to specifically tighten your shot grouping around your target. There is this ancient Stoic notion of a *dichotomy of control* that applies here to investing, as it does to archery: **We need to control what we can control in a way that gets us closer to our target (of higher wealth)—and certainly not further from it.**

This is cost-effective risk mitigation. You look and feel like you can see around corners and can always hit the bullseye, even though, in actuality, of course, you can't.

You see, a cost-effective safe haven doesn't just slash risk. It actually lets you simultaneously take more risk. If that twist gave you pause for a moment, then good. It should!

FIRST PRINCIPLES

We will start here at square one with a few basic foundational truths about investing. As the great fourth-century BCE Greek philosopher (and the world's first *real scientist*) Aristotle wrote, “The naturally proper direction of our road is from things better known and clearer to us,” from what he called “first principles” or “the first basis from which a thing is known.” They are what comes first, our *a priori* propositions or universal premises.

We start with these first principles and use them as conceptual building blocks to eventually form deductive,

testable hypotheses. While these first principles may sound obvious and sort of trivial at first, they are not. In fact, they are going to become really important as we explore some counter-intuitive things about a pretty technical subject. They will allow us to change and even exploit the commonly held, though false, heuristics at the core of modern finance. They will even afford us a certain existential authenticity, allowing us to make our investing consistent with what we believe—to “bet our beliefs” as they say, despite all the external pressures to conform.

Principle number one is that investing is a process that happens sequentially through time. Investing is not static. It does not occur in just one interval of time, nor in many intervals of time aggregated together as one. (Albert Einstein purportedly noted that “the only reason for time is so that everything doesn’t happen at once.”) Time is the medium through which life takes place, and so it is the medium through which investing takes place. We are stretched across time. Investing and risk are a multi-period problem; and returns are an iterative, multiplicative process. *They compound*: In each period, we generally invest what we are left with from the last period. Like the geometric growth of offspring across generations, *we parlay our capital*. This principle fundamentally determines the nature of investing and the way we need to think about and interpret returns.

But as obvious as it is, just try telling that to a hedge fund manager whose incentive fee is based on annual performance: *investing happens through time*. Or, deliver that message to the pension fund that is harshly judged on its ability to meet an annual benchmark over the short term rather than over a timeline consistent with its beneficiaries. And try telling that to the economics and behavioral finance communities who label behavior “irrational” when it doesn’t appear optimal within their single-period, timeless framework. Do so and you should expect funny looks. But they criticize what they can’t understand; they’re no Einsteins. As for us, we need to ask ourselves

every day: What is the meaning of time in investing? The answer, as you'll see, changes everything.

Principle number two is that there is only one explicit purpose or goal of investing, and that is to maximize our wealth over time. Period. This is exactly what we are trying to do with every additional, incremental decision that we make as investors. It is the target we shoot at. I'm not talking about the elusive *mathematical expectation* of wealth, nor our wealth relative to some arbitrary benchmark (though there are plenty of managers incentivized to care only about that). Rather, I'm talking about our *actual* realized ending wealth—meaning the outcome we're actually left with. (They are not the same.) This is equivalent to maximizing the rate that we grow or compound wealth over time—our compound annual growth rate, or CAGR. All investors are absolute return, compounding investors. I don't believe any thinking person, and certainly not any practitioner, should disagree with this principle. It is common sense.

And yet, some would still protest that the goal of investing should be to further humanity's progress, ease its burdens, do good to consumers and the world—and the profits will follow. But this is only another way of restating our principle. The consumer is the sovereign king whom capital must serve in order to be profitable; it is because of this that capital investment and entrepreneurship have objectively done more good for the world than any government or charity ever could.

Others might still protest that the goal of investing should be to maximize our wealth, given a certain level of theoretical risk taken. (This is where things would start to get a little muddled were it not for our next principle.) Since we know that risk mitigation is investing, we can deduce that the explicit purpose of risk mitigation is the very same as the explicit purpose of investing: to maximize the rate at which we compound our wealth over time—in other words, to *cost-effectively lower risk*.

This leads to principle number three: If a risk-mitigation strategy achieves its purpose by cost-effectively

lowering a portfolio's risk, then adding that strategy raises the portfolio's CAGR over time.

It makes sense. If we mitigate risk effectively by constraining it deliberately, shouldn't the point be to experience less loss as a result, such that *over time* we end up making more? And if we don't end up making more, will we still be glad we did it? What would the point have been? Is there any other reason to mitigate risk?

Risk mitigation, when done well, should provide a tangible, positive economic effect relative to its cost. That is, it should be cost-effective, and thus a good value proposition. Accordingly, a risk-mitigation strategy should be evaluated based on its degree of cost-effectiveness at lowering risk—not just on its effectiveness at lowering risk.

Of course, we may have mitigated the risk of a remote, extreme loss that never happened. And yet, such a remote loss could happen suddenly at any time, and our risk mitigation could thus raise our CAGR (relative to our position had we not used that mitigation). This is the problem of induction, where the arrival of just one black swan falsifies any claim that all swans are white. But when I say that cost-effective risk mitigation raises a portfolio's CAGR over time, this also means over a sufficiently broad range of observable outcomes, which largely resolves these epistemological problems. (We will address this with what is known as a bootstrap in later chapters.)

Amusingly, this principle that risk mitigation, done well, should raise our CAGR over time is actually quite controversial. In fact, most practitioners and academics would probably call it a crackpot idea. Lower risk comes at a cost, they believe, so risk mitigation is incompatible with higher returns.

To wit, one of the historic commodities trading houses had the German motto: "*Besser gut schlafen, als gut essen.*" ("It is better to sleep well than to eat well.") And academics claim that lower returns actually accompany lower risk or volatility *as a*

consequence, ipso facto. Their story goes, as you hedge and diversify away all of your correlated and uncorrelated risks, respectively, your returns will approach the lowly *risk-free rate*. Or going the other way, the academics argue that in order to induce an investor to hold an asset that is relatively volatile, its price drops until its expected return becomes high enough to justify the additional risk. In their world, it sounds reasonable, and it's a comfortable story, asserted but not proved: You've got to take more risk to make higher returns; sleeping well comes at a cost. No guts, no glory.

To make matters worse, academics furthered this idea by positing that investing and risk mitigation are about lowering or calibrating a portfolio's volatility relative to its average return—the *risk-adjusted return* or dreaded *Sharpe ratio*—unwittingly at the expense of the growth rate of wealth. They thus claim an intellectually dishonest victory based on their own theoretical scoreboard. It is a solution in search of a problem, and a bad idea. (It's even a big reason for our great dilemma.) I don't really believe that most investors even have this bad idea. Rather, to paraphrase Carl Jung, the idea has them.

We need to measure our success as investors by the practical scoreboard that counts, rather than the theoretical ones that don't. And there is just one scoreboard that counts, just one bullseye.

But we are often lured away from such practical objectives by gratuitous mathematical formulas. Modern quantitative finance suffers from a certain science or physics envy. After all, according to the American physicist Richard Feynman, "Physics is like sex: sure, it may give some practical results, but that's not why we do it."

Well, practical results are precisely why we do what we do in investing *and* risk mitigation: to maximize the growth rate of wealth by lowering risk. And the best practices of the scientific method can actually help us with this.

MODUS TOLLENS

Aristotle is generally considered to be the earliest and foremost developer of the idea of *deductive* reasoning, or *sullogismos*. Deduction is “top-down” logic, whereby general rules or premises are applied to particular cases or conclusions. This contrasts with induction, which is “bottom-up” logic and goes in the opposite direction, whereby particular cases or premises are applied to reach a general rule or conclusion. **Examining the geometry of a die to estimate the frequency that any side will come up over repeated rolls is deductive reasoning. Reasoning in the other direction, by repeatedly rolling a die and using those results to estimate the geometry of the die, is inductive reasoning.** (We will be rolling the dice in both directions in this book.)

A syllogism applies deductive reasoning to draw a valid conclusion from assumed premises. One example is the syllogism called *modus tollens* or “denying the consequent.” It is the main logical method to avoid mistakes of reason in science—what Feynman has described as “what we do to keep from lying to ourselves.” It is an ideal BS filter (so don’t be too surprised if you haven’t encountered it in the context of investing).

A *modus tollens* takes the form of “If *H*, then *O*. Not *O*. Therefore, not *H*” (with *H* for hypothesis and *O* for observable). There are two premises—an explanatory hypothesis, made up of an antecedent and consequent, paired with an observable; brought together, they yield a conclusion, which follows logically from the premises. The logic goes, if a statement is true, then so is its contrapositive.

Think of this example of *modus tollens* involving my dog Nana:

- If Nana is good at catching groundhogs, then I won’t have a groundhog problem.
- I have a groundhog problem.
- Therefore, Nana isn’t good at catching groundhogs.

We can see that a *modus tollens* serves the specific role of *falsifying* or eliminating a hypothesis. But neither it, nor anything else for that matter, can ever be used to *verify* a hypothesis as *true*. When we pair our proposed hypothesis with a minor premise that is an observable fact, we have a well-constructed test of that hypothesis. *Modus tollens*, then, is the logical principle of the empirical sciences, and the scientific method itself; it allows us to clarify our ideas and move them away from mere *metaphysics*. Scientific rigor demands that we be able to pose, experimentally test, and ultimately falsify theories or conjectures in this way. When, like sleuths, we disqualify false theories whenever we can, then step by step we approach the truth. It's all very Sherlock Holmesian: "When you have eliminated the impossible, whatever remains, however improbable, must be the truth."

Most significantly, the twentieth-century Austrian philosopher of science Karl Popper constructed his whole *falsification principle* around it—as the fundamental demarcation between science and pseudoscience. "Universal statements are never derivable from singular statements, but can be contradicted by singular statements," as Popper wrote in *The Logic of Scientific Discovery*. "Consequently, it is possible by means of purely deductive inferences (with the help of the *modus tollens* of classical logic) to argue from the truth of singular statements to the falsity of universal statements. Such an argument to the falsity of universal statements is the only strictly deductive kind of inference that proceeds, as it were, in the 'inductive direction'; that is, from singular to universal statements."

So far, we have been discussing cost-effective risk mitigation as if it were something worthy of our attention—something that *exists*. Even uttering the phrase presupposes it. (And this is why no one ever really utters the phrase "cost-effective" in the context of risk mitigation in investing. Have you noticed?) Taking such a conclusion for granted is begging the question.

It might even appear that we did this in our principle number three. However, the principle only claimed what risk

mitigation should be, not necessarily what it always is. Cost-effective risk mitigation could still be only theoretical, and not actually possible.

So, instead, we need to treat this principle as a conditional premise. It is an explanatory hypothesis, and this conveniently suggests our own *modus tollens* syllogism for safe havens, which we will be testing and investigating over and over:

- *If a strategy cost-effectively mitigates a portfolio's risk, then adding that strategy raises the portfolio's CAGR over time.*
- *Adding that strategy doesn't raise the portfolio's CAGR over time.*
- *Therefore, it does not cost-effectively mitigate the portfolio's risk.*

What we have here is a natural, testable conjecture about safe haven investing. And it's important to understand what this hypothesis testing can and cannot do. It can only refute or falsify the hypothesis. If a safe haven strategy does *not* raise a portfolio's CAGR over time, then the null hypothesis—that the strategy cost-effectively mitigates the portfolio's risk—does not hold. If it disagrees with experiment, it is wrong—it is not a cost-effective safe haven strategy. What we cannot do, however, is prove that something *is* a cost-effective safe haven strategy. Such is the scientific method.

To illustrate why you can't prove things in reverse, it's important to note that I could not have posed this syllogism instead, as the inverse of our premises: *If a strategy does not cost-effectively mitigate a portfolio's risk, then adding that strategy lowers the portfolio's CAGR over time.* That would be deductively invalid; it mistakes a sufficient condition for a necessary condition. **Observing that adding the strategy raises the portfolio's CAGR over time actually proves nothing about cost-effective risk mitigation. This is because there are other ways that the strategy could have raised the portfolio's**

CAGR; the strategy needn't have even mitigated risk at all, and it may have even added risk. We would need to delve deeper into the source of that outperformance. (As Hemingway said, "Being against evil doesn't make you good.")

This is equivalent to the fallacy of *affirming the consequent*: If "O is true," can we then conclude that "H is true"? Thinking that we can is a mistake often committed even in the physical sciences: "If my theory is correct, then we will observe this data. We observe this data; therefore, my theory is correct."

In my previous example with Nana, what if I *don't* have a groundhog problem? Can I proudly claim that my Nana is good at catching groundhogs? After all, there could be myriad other possible reasons that I don't have a groundhog problem. Maybe they were scared off by our resident fox. Or perhaps my son has been playing in our woods with his bow and arrow and ghillie suit.

We would similarly mistake sufficiency for necessity with the fallacy of *denying the antecedent*. In this case, we might conclude that I have a groundhog problem as a logical consequence of knowing that Nana is not good at hunting groundhogs.

All knowledge is a hypothesis; it is all conjectural and provisional, and it can only ever be falsified, never confirmed.

Now, a critical part of this scientific method is the way we go about choosing the very hypothesis that we are putting to the test. We specifically need to avoid *ad hoc* hypotheses that simply fit our observations. They are always there, ripe for the picking. We need logical explanations that are independent of and formed prior to those observations.

Scientific knowledge is not just to know *that* something is so; it is just as important to know *why* something is so.

The deductive thinking behind my hypothesis "Nana is good at catching groundhogs" was so important because it's

going to hang around as sort of our working hypothesis until I manage to falsify it. Did I have a sound deductive reason to think that her skill, if it exists, would really result in the disappearance of groundhogs? Skilled or not, does she prefer to spend her summer days sleeping indoors? (As a Bernese Mountain Dog, she much prefers snow and air conditioning.) Does she more often chase or get chased by the groundhogs?

DEDUCTIVE DICE

We need a sound, deductive framework for understanding why and how risk mitigation can lead to higher compound growth rates. Is it even possible? How can we expect risk mitigation to be cost-effective? What's the mechanism behind it? And how might we recognize this ability if or when we see it?

There are so many forces swirling around in investing and markets; any attempts at explaining risk mitigation based on that swirling data will likely get us into trouble. So we will need to figure out, deductively, the mechanism that gives rise to our hypothesis. It won't be enough just to show that a strategy raises wealth and, *poof!*, it is cost-effective risk mitigation. We need to understand the forces at work to make that happen.

The best deductive tool at our disposal will be the same deductive tool used from prehistory right up to today to discover and comprehend the general science of probability and the formal idea of risk.

Archaeological excavations have uncovered collections of talus bones, or ankle bones of goats and sheep, dating back to 5,000 BCE. These four-sided bones were the oldest of all gambling devices. Our more familiar six-sided "bones" started showing up by 3,000 BCE. Visibly, dice have been embedded in civilization's history as the generators of fate, then chance, and then even skill (in the early precursors of modern backgammon). They were ubiquitous throughout both the depth

and breadth of that history and are a part of our collective unconscious.

But it took ages, and the need to better understand wagers on games, for dice to become the intuitive pedagogical tools of deductive inference that would set in motion a theory of probability. As early as the fourth-century BCE, Aristotle casually pointed out that, while it is easy to make a couple of lucky rolls of a die, with 10,000 repeated trials, divine intervention be damned, the luck of the die evens out: “The probable is that which for the most part happens.” Imagine, this was revolutionary stuff! But the ancient Greeks and Romans never really got it—they never even bothered to make sure their dice faces were symmetric. If everything was fate anyway, what difference did it make? Julius Caesar’s famous line “the die has been cast” was not a statement about probability. (For all the wisdom that we ascribe to the ancients, if you had the ability to go back in time, you would totally clean up gambling against them.)

It wasn’t until much later, by the seventeenth century, that Galileo and then Blaise Pascal and Pierre de Fermat became gambling advisor-mercenaries to various noblemen. For instance, the Chevalier de Méré needed advice on his costly observation that betting with even odds on the appearance of a 6 in four rolls of a single die was profitable in the long run, whereas betting with the same odds on a double-6 in 24 rolls of two dice was not. (In 1952, the famed New York City gambler “Fat the Butch” rediscovered this same deductive fact in his own rather costly hypothesis test.)

At that point, probability was all about deductive reasoning, starting with the known properties of the generator (a die) and then reasoning forward with expectations about its particular outcomes. Repeatability was implicitly a necessary condition to probabilistic inference. This was the *frequentist* perspective, where the very meaning of probability was the frequency of occurrences over many trials. It is the logic of the gambler with an edge, the logic of the casino. Probability was truly coming

of age as a neat trick to allow mathematicians to pick off degenerate gamblers. These were the original *quants*—money on the line has a way of sparking innovation. (Heck, it took having an options position for me to ever start really thinking about math.)

Of course, we have always understood risk and its mitigation in our bones; that's how we made it this far, after all. But along with advancements in our understanding of probability grew a gradual formalization and sophistication of risk mitigation. And we can think of the growth of that formality first and foremost as the growth of innovations in insurance—which itself would facilitate an explosion in risk taking and innovations. Insurance is an ancient idea, and a key part of the very progress of our civilization. It began as solidarity, as risks were shared—spreading throughout small villages, for instance, as commitments to mutually self-insure and share the replacement costs of homes within the community. **This aggregation of individual risks created a frequentist's perspective where there otherwise was none, effectively expanding an individual's sample size from 1 to the size of their community.**

Fast forward to the twentieth century, and an epic nerd skirmish would erupt over this perspective. The newly founded *Bayesians* and Popper's *propensity theory*—where *probability* instead meant “degrees of belief” or “tendencies,” respectively—went head-to-head with the simpleton frequentists. And it doesn't really matter who was right. All that matters is which perspective is being used. When your sample size is small, and worse yet unique and unrepeatable, no matter your subjective probabilities, there is so much noise in your sample you can hardly know anything anyway. Your N equals 1. You are a punter, hoping for good luck or good fate. But if your success or failure relies on many outcomes, over many rolls of the dice, then you naturally care about the properties over many rolls of the dice. Your N is large. You are the house exploiting the “house edge” through repetition, to quash the randomness.

And the house doesn't gamble. You are, as the poker theorist David Sklansky wrote, "at war with luck."

Most people would say (or at least their actions imply) that they are the house in their investing. And 93% of people also say that they are above average drivers. They aren't, in either case.

"At war with luck"—using our "skills to minimize luck as much as possible"—does indeed apply to risk mitigation. That is precisely how this war is fought in investing. In my case, it is no less than everything I do as an investor. So it is fitting that we begin to understand risk mitigation by deductively stripping it down to its bare bones, literally.

This is one of the most valuable things I have learned from Nassim Taleb, his valid warnings against the *ludic fallacy* notwithstanding (where "the narrow world of games and dice" have so little in common with the untamed risk of the real world): that playing around with and meditating on simplified Monte Carlo simulations, or "alternative histories," is the best way to figure things out.

After all, according to Popper, science is "the art of systematic over-simplification."

Beyond epistemological rigor, the biggest advantage I gain from building my safe haven hypothesis deductively and piecemeal using games of dice is *transparency*. You are going to see some things about safe haven investing that will seem to defy common sense. Couple that with the fact that there are frighteningly many ways that the investment industry regularly smokes people with complicated and unfalsifiable (and thus pseudoscientific) theories and cherry-picked market data, and you can see my concern about a healthy skepticism. "Trust, but verify": When something doesn't smell right, go back to the beginning, to our simple and transparent deductive dice illustrations. Roll your own and play along at home. There's no place to hide there.

As Feynman famously said:

In general, we look for a new law by the following process. First, we guess it. Then, we compute the consequences of the guess to see what, if this is right, it would imply. And then we compare the computation results directly with observations to see if it works. If it disagrees with experiment, it's wrong. In that simple statement is the key to science. It doesn't make any difference how beautiful your guess is, it doesn't make any difference how smart you are who made the guess, or what his name is—if it disagrees with experiment, it's wrong. That's all there is to it.

This will be my analytical framework for this book.

In Part One, we start with “what comes first” (the *a priori*), with an intuitive construction and examination of those fundamental safe haven mechanisms, with the help of our deductive dice. “First, we guess it.” Then in Part Two, with “what comes after” (the *a posteriori*), we start to formulate testable safe haven hypotheses based on those mechanisms—hypotheses about how we might expect them to work. We will conduct clinical trials or experiments on different idealized safe havens (what I call cartoons)—to “compute the consequences of the guess.” Then, we can compare those results “directly with observations to see if it works,” by running those same experiments on the diverse range of real-world safe havens themselves. Our aim is to try to falsify, in a meaningful and rigorous way, the hypothesis that safe havens as a group—and various safe havens in particular—can raise wealth by lowering risk. This is not a foregone conclusion; after all, it is considered a crackpot idea.

Through this methodology, you will hopefully understand what works in risk mitigation, what does not, and why. And this understanding will protect you more than any individual safe haven ever could. It will guide you toward our goal as investors.

If you think about it, cost-effective risk mitigation—or raising compound growth rates and thus wealth through lower risk—is really our comprehensive

goal as investors. It is the true essence of investment management. In and of itself, it is the specific meta-purpose or meaning we pursue when we deploy capital—what we hunt for relentlessly, our buried treasure.

Yes, there really is a buried treasure for investors, one that solves our monumental problem by showing that the great dilemma of risk—the ostensible tradeoff between higher returns and lower risk—is actually a false choice. But this treasure wasn't so much hidden away by pirates of lore as it was cloaked behind the flawed apparatus of modern finance, shrouded behind its veil of rigor. As a result, it appears as a myth, an idealized and elusive goal. But that's only because investors have been looking too narrowly and in the wrong places for it. We need a more holistic approach; we also need a treasure map to know where to dig.

But just because that buried treasure exists doesn't mean we will ever find it. The greatest value—more than in the treasure itself—will be in what we gain from the hunt.

