

ALPHA ASSETS

Competitive Advantage in
Business and in Life

KARL T. ULRICH

THE WHARTON SCHOOL



ALPHA ASSETS


*Competitive Advantage
in Business and in Life*

KARL T. ULRICH

THE WHARTON SCHOOL



Philadelphia



Copyright © 2026 by Karl T. Ulrich. All rights reserved.

Published by Pontifica Press, Philadelphia. pontifica.com · alphaassets.com

This book is maintained under version control. This copy was built from the manuscript repository and reflects its state on October 9, 2026. Edition: working proof, not for distribution.

ISBN 979-8-0000000-0-0 (placeholder)

Set in Source Serif 4 and Source Sans 3. Trim 5 × 7½ inches.

Contents

1	What Explains Success?	1
2	What Is an Alpha Asset?	13
3	Lucky or Good?	25
4	How Does Luck Become an Asset?	40
5	How Does Advantage Accumulate?	51
6	Where Should You Apply Your Strengths?	66
7	How Do You Build an Advantage Deliberately?	77
8	When Does Advantage Stop Working?	89
9	What Should Success Mean?	98
A	The Alpha Asset Worksheet	111
	Notes	120
	About the Author	133
	Colophon	134
	Proof pages	135

What Explains Success?

In January 2016, at the cyclocross world championships in Zolder, Belgium, officials carried a scanner through the pit area and passed it over more than a hundred bicycles. Cyclocross is the winter discipline of bicycle racing: an hour of laps on a short course of grass, mud and sand, with barriers the riders dismount for, on bikes that look like road bikes fitted with knobbed tires. One of the bicycles had a motor in it. The bike stood in the pit of Femke Van den Driessche, nineteen years old and the reigning Belgian and European champion in the women's under-23 category, and the motor sat inside its seat tube with a battery, wired to a switch under the handlebar tape. Whether it had ever been switched on did not matter. The presence of one in a bike at a race is the offense. In April the Union Cycliste Internationale banned her for six years and annulled every result she had posted since the previous October, both titles included.

Nobody needed a result to know what the motor did. It added power to the pedals, and more power is more speed. The motor was an advantage anyone could understand before the race, from the mechanism alone, with no finish to explain. Most explanations of success run the other way, offered after the result to account for it.

The motor was also unfair, and helping was not the reason. Everything a rider brings to the line helps. The competitors had agreed to a contest of bodies and skill, and a motor is neither. Suppose instead that the rules allowed it and every rider had the same one. The motors would still add power, and having one would distinguish nobody. It would be like the skinsuit every rider already wears: a tight suit is faster than a flapping jacket, and because every rider has one, it explains nothing about who wins. An advantage the rivals can match is not an advantage.

The rules forbid very little else. The rider with the biggest engine at Zolder was born with most of it, in heart and lungs and muscle, and chose their parents no more than their rivals chose to ride without a motor. An engine in the seat tube is cheating and an engine in the chest is a gift, and the difference between them is a decision the sport has made rather than a fact about engines. Sports that want closer contests do equalize. A handicap horse race loads the better horse with more weight, and a golf handicap gives the weaker player strokes, so that a duffer and a club champion can play for money. Cycling does neither, because the race exists to find out who is fastest given what each rider was born with and what they built. The people who run a contest choose where to draw the line between the advantages it allows and the ones it forbids.

In business the line is the law, and inside the line an unfair advantage is the object of the exercise. Investors ask founders what their unfair advantage is, and they ask

it as encouragement. What a company wants is the equivalent of the motor, legitimately acquired: something that improves its results and that its rivals cannot pick up at the bike shop. It might be a production process, a reputation, a base of customers, or the scale to buy inputs for less than anyone else pays. A resource with both properties, one that improves performance and that rivals cannot easily get, I will call an alpha asset, and chapter 2 gives the definition its edges.

The questions the motor leaves open are the ones this book is about: how much of an advantage is given and how much built, what a result can tell us about whether anyone has one, and how a company or a person who wants one goes about getting it. The place to begin is the starting line at Zolder, with the motor removed.

The Starting Line

One rider can sustain more power than the others because of the body they were born with. Another has spent eight winters learning to carry speed through a muddy corner, a skill that is built. A third rides better equipment or has a team car and a mechanic in the pit. All of them train. They still arrive at the line with different prospects.

The body is a good place to see how an advantage can be real and still have mixed origins. A cyclist's sustainable power depends on maximal oxygen uptake, VO_2max , measured in milliliters of oxygen per kilogram of body weight per minute, together with the fraction of that

maximum they can hold and the efficiency with which they turn oxygen into work at the pedals. An untrained young woman typically measures in the high thirties; an untrained young man in the mid forties. Elite endurance athletes measure in the seventies and eighties. Greg LeMond, the American who won the Tour de France three times, is usually credited with 92.5, and Bjørn Dæhlie, the Norwegian cross-country skier who won eight Olympic gold medals in the 1990s, with 96, though the published figures for famous athletes come from different labs on different protocols and are better treated as a range than as a ranking. Training raises the number, by an amount that differs greatly from one body to the next. In the HERITAGE Family Study, 481 sedentary adults from 98 families followed the same twenty-week program. The average gain in maximal oxygen uptake was about 400 milliliters per minute. Some participants gained almost nothing and others gained more than a liter, and the response ran in families, with roughly half of its variation attributable to inheritance. The rider's engine is partly a gift and partly a construction, in proportions that differ from rider to rider, and nobody can read the proportions off a result.

Then there is the day. A cyclocross course is a sequence of corners, mud, sand, and barriers, and a race is decided by a few seconds out of an hour. A puncture, a crash in front of them, a line through a corner that held on one lap and gave way on the next: these decide places. A rival's puncture is a favorable event that leaves the rider no faster than they were. Companies have weather of their own. A

construction boom lifts every supplier of building materials, a customer's bankruptcy empties a supplier's order book, and a fire at a rival's plant fills it, with no change in how well anyone operates.

So the riders differ in three ways. Some differences are advantages: a capability or a position that would, in expectation, produce better results than the rivals in this contest. Some advantages were inherited and some were built, and most, like the engine, are both. And some differences in results are not advantages at all but the variation of the day, which chapter 3 will call noise. Three separate questions follow. Whether an advantage exists is a question about mechanism: does the thing add power? What explains a particular result is a question about a mixture that no single result reveals. And what credit the winner is due is a third question, which depends on the first two and on where the advantage came from. Most arguments about success run the three together. This book keeps them apart.

The same three questions apply to a company that earns more than its rivals for a decade, a fund manager who beats the market five years running, a founder whose company sells for a billion dollars, and a professor whose early papers are cited more than their classmates' papers. In each case something improved the odds, and we want to know what it was, where it came from, whether anyone else could get it, and whether it will still be there next year.

What a Result Can Tell

Michael Phelps won eight gold medals at the 2008 Olympics. The profiles written that summer gave his body as the reason: size 14 feet, a wingspan reported as 80 inches against a height of 76, and ankles said to flex farther than other swimmers' ankles. The wingspan came from a newspaper, the feet and the ankles from a blog, and the profiles repeated the figures from one another. When Scientific American put the ankle claim to an exercise physiologist, he doubted that Phelps was much more flexible than his rivals and suggested that the main difference was fitness. Phelps may have the ankles. Nobody had measured them before the explanation went to print.

Explanations of success drive decisions. A board decides whether to keep or fire a chief executive on three years of results. A limited partner commits a hundred million dollars to a venture fund on the record of its last two funds. A hiring manager chooses the candidate whose last project succeeded. A young engineer decides whether to copy the career of a founder they admire. Each of these decisions assumes that the record reveals the ability. When inheritance and the day account for most of the record, the decision rests on a record that a person with no such ability might have produced just as easily.

The problem is easier to see at a small scale. Two people sell equipment for the same company. One books two million dollars in a year, the other one million. The first may be the better salesperson. They may also have inher-

ited a territory full of customers whose machines were due for replacement, while the second spent the year winning customers who had never bought from the company at all. Asked to explain their year, the first will describe how they qualify a prospect and how they follow up, and they will mean it. The replacement schedule of the equipment their predecessor sold will not be in the account. A company that responds by distributing their presentation slides has copied the easiest thing to copy. Whoever succeeded was a participant in an experiment run once. Knowing what they did is useful. Knowing that it caused the result is something more.

Writers on investing, on randomness, and on the stories business books tell about successful companies have made the case that a record is a poor guide to skill. The question here is different. Whether the salesperson earned their territory, or Van den Driessche her 2015 titles, is a question about origins, and it is often unanswerable. Whether either held a real advantage going into the next year is a question about the present, and it is often answerable: a reputation, a customer base, a body, a position can be observed directly, and its effects can be reasoned about, whatever produced it. A lucky early result can create a real later advantage. A board, an investor, or a person planning a career needs the second answer more than the first, and most of this book is about how to get it.

My Mother's Question

I had been working on this book, in its earlier form, for several years when my mother asked me what it was about. She is a historian. I explained the basic idea: that companies and people who sustain an advantage over rivals possess something that improves their results and that rivals cannot easily get, and that most of these advantages are built over time by mechanisms I could describe. She listened and then asked what evidence I had that any of this was true.

Ouch. My education is in science and technology. I am usually the one asking for the evidence. I did not have a good answer that afternoon. I have a better one now. The claims in this book come in four kinds. Some are definitions, true by construction, which prove nothing about the world but turn a vague goal into a question a manager can act on. Some are mechanisms argued from logic, like the loop in which a growing customer base lowers cost and lower cost attracts customers; I can give no experiment that establishes such a loop, but I can state each link so the reader can judge whether any fails. Some are findings from populations, such as the persistence of returns across thousands of firms or the effect of a randomized management intervention across dozens of plants; these are the strongest evidence the book has, and they say what is true on average, never what is true of this firm or this person. And some are stories, including Phelps's ankles and including my own, which can illustrate a mechanism

and cannot establish it. When I tell a story, I will say what mechanism it illustrates, what observation supports the mechanism apart from the result, and what else might explain the result. Where I cannot, the story is decoration and I will say so.

My own career is one of the running cases in this book, alongside two companies I co-founded, Xootr and Maker-Stock, because I know what was decided and what was merely survived, which is more than I can know about Amazon. My own effort is vivid to me. What someone else would have made of my opportunities is not.

Unequal and Unjust

The second paragraph of the United States Declaration of Independence holds it to be self-evident that all men are created equal. In the ordinary sense of the word, nothing of the kind is evident. Boulders, snowflakes, dogs, companies, and countries differ from one another along every dimension anyone has measured, and so do people. Van den Driessche's rivals did not start the race with her engine, her years in the mud, or her team. Most of them did not start with a motor either.

A large part of who we are was settled before we could choose anything. The genes, the family, the country, the decade, the labor market in the year of the first job search. Economists following Canadian university graduates found that those who entered the labor market in a recession earned less for about ten years than those

who graduated a few years earlier or later. They started at lower-paying employers and recovered only partly, by moving to better ones. The graduates did nothing to earn or to deserve the difference. Phelps did nothing to acquire his wingspan. LeMond did nothing to acquire his lungs.

Unequal is a description. Unjust is a judgment. An inherited advantage is not a wrong done to anyone. A motor in the seat tube is. In between lie the advantages that are perfectly legal and that still raise the question of what a society owes to those without them. I hold two beliefs on this. Every person and every organization should have the opportunity to prove themselves and thereby earn access to further opportunities. That is equality of opportunity, as distinct from equality of outcomes, and no society fully provides it. Norway does not give every teenager a place at its Olympic training center. It does give nearly every child a pair of cross-country skis and a local race, which can lead to the next race, and the next. The second belief is that because so much of the inequality that shapes well-being is unchosen, people acting together have an obligation to mitigate it, especially where fundamental needs are concerned.

Organizations are, in one respect, more fairly placed than people. A person cannot change their VO2 max by much. A company can change almost everything about itself, given time and money, and most of the advantages this book describes are of the kind that a company builds rather than inherits. Even so, companies are born unequal. Some start with a founder's reputation, a par-

ent's customers, or a patent. A company can be a month old while its most useful resources are decades old, and the date of incorporation says little about how far back to look for the origins of its advantage. The accident of being first into a market that later turned out to matter is as unearned as a wingspan.

The rest of this book asks four questions, for a company and for a person. What improves your prospects? What stops others from acquiring the same thing? What does an observed result reveal about either? And given the answers, what should you build next? The thing that improves prospects and resists acquisition is the alpha asset, and chapter 2 defines it. A base of two hundred university workshops that a small materials company acquired one at a time over two years is one. So is the purchasing scale that lets the largest restaurant chain buy beef for less than any rival, and so is a name on a basketball shoe. Chapter 3 takes up what results reveal. Chapter 4 shows how an early result, lucky or not, becomes an asset by changing the opportunities that follow. The chapters after that are about building, choosing where to compete, knowing when an advantage has stopped working, and what success entitles anyone to claim.

A rider who knows which of their advantages were given and which were built still has to decide how to train before the next season. The rest of us face the same decision with less convenient finish lines, and the decision

has the same form: given the advantages you have and the ones you lack, what should you build next?

The motor in the seat tube settled Van den Driessche's case. The rest of us are not scanned. We are judged on results, by boards and investors and employers and by ourselves, and the judgment assumes that results reveal what produced them. The chapters that follow are about how often that assumption holds, and what to do on the many occasions when it does not.

What Is an Alpha Asset?

In 2018 I pulled from a financial database the records of every publicly traded restaurant company in the United States. For each one I computed its average return on invested capital over the ten years from 2006 to 2015, less its cost of capital. The list has familiar names on it: McDonald's, Starbucks, Yum Brands (which owns KFC and Taco Bell), Chipotle, Domino's, and several dozen smaller chains. As a group they earned, by my calculation, about six percentage points a year more than their capital cost them. But the spread was not evenly shared. The companies at the top of the list earned several times the group average, while those at the bottom destroyed value for a decade and survived anyway.

I am an engineer, not an economist, and my perhaps naive view of economic theory is that competition should erase such differences. A business that earns too much attracts imitators and entrants, supply rises, prices fall, and the excess profit disappears. On a trip to Istanbul I watched this view hold exactly. Fish bread is a street food sold from a small charcoal grill set on a couple of crates at the edge of a market. The vendor grills two pieces of fish, toasts a split loaf, and assembles a sandwich with olive oil, herbs and salt, all for the equivalent of about four dol-

lars. I ate a lot of fish bread, and to my palate it was identical from one vendor to the next, uniformly excellent, and priced the same everywhere. If a vendor raised prices by ten percent, a new grill would appear on the next set of crates within the week. You can be pretty sure that no fish bread vendor in Istanbul earns much more than any other. (I concede that a tourist's reading of the fish bread economy may be incomplete. As a young man I spent a year in Naples watching teenagers sell cigarettes from folding tables. It took me most of that year to learn about the supply chain and the management hierarchy behind what looked like a free market.)

Very few markets in a modern economy look like the fish bread market. Across thousands of public companies, the spread between the return on capital and the cost of capital persists for decades. When the spread is decomposed statistically, a large part of it belongs to the individual firm rather than to its industry. Why does it persist? The marginal competitor, the firm on the sidelines that would enter if the business looked attractive enough, cannot get what the incumbent has. Yum Brands has something that a smaller chain like Potbelly cannot easily acquire.

Two Tests

By an *advantage* I mean a capability or a position that improves expected performance relative to the relevant rivals, in a particular setting, over a particular

period. Expected, because a lucky result is not an advantage (chapter 1 was about the difference). Relative to rivals, because a capability that every competitor shares improves nobody's standing. In a setting, because the same capability can be decisive in one market and ordinary in another; a strong hand at bridge is nothing at poker. Over a period, because nearly every advantage erodes.

The thing that confers an advantage I call an *alpha asset*: a resource that an entity owns or controls that is (a) performance enhancing for the job the entity does and (b) hard for its rivals to acquire. I borrowed the term from finance. In the capital asset pricing model, a security's expected return is decomposed into beta, the part explained by exposure to the market as a whole, and alpha, the part specific to the security. Alpha assets are the underlying explanation for financial alpha, the reason a firm earns more than its exposure to the economy would predict.

The concept is not mine. It is a reframing of the dominant idea in the academic field of competitive strategy, the *resource-based view of the firm*, which Jay Barney set out in 1991. Jay, now at the University of Utah, is an affable and enthusiastic man with whom I once worked on a consulting project in healthcare. The resource-based view (RBV, in the alphabet soup of acronyms that specialized fields love) holds that a firm sustains an advantage when it controls VRIN resources. (Got it?) VRIN stands for valuable, rare, inimitable and non-substitutable. (Got it now?) I made a good-faith effort to use VRIN in an

earlier book, but it is just too wonky for anyone outside a doctoral seminar. The resource-based view is one of the most powerful ideas in social science and one of the most poorly marketed. It needs a makeover, so let's relabel VRIN resources as alpha assets and simplify the four conditions to two: valuable and non-substitutable collapse into performance enhancing; rare and inimitable collapse into hard to acquire.

A resource is performance enhancing when its presence causes a significant improvement in results, with no good substitute, for the competitors who do the job. Three phrases in that definition need defending.

No good substitute. In a solo time trial, a race against the clock, there is no substitute for a light, stiff bicycle; a rider who turns up on a Dutch utility bike finishes last. There is no substitute for electricity in a factory. But there is a substitute for the Advil brand of anti-inflammatory, since Motrin, Nurofen and the generic bottle contain the same ibuprofen, and so the molecule is closer to being uniquely performance enhancing than the brand is.

Significant. The question is whether the resource accounts for much of the difference between winning and losing, not merely whether it helps. Water is critical to every rider; a dehydrated cyclist does not finish. An earpiece with a radio link to the team car is a small improvement, decisive on the right day, but a rider can win without it. Ibuprofen explains none of the variation among the top finishers.

For the competitors who do the job. In my 2022 draft of this definition I wrote "for most competitors," meaning to exclude idiosyncratic needs such as insulin, which is performance enhancing for a diabetic cyclist and irrelevant to the rest of the field. That exclusion was too strong, especially for people, because a resource can be performance enhancing only in combination with something else the holder has. A tough general counsel may be essential to a company run by a mercurial founder and surplus to a company run by a cautious one. In such cases the asset is the combination, and the combination can be rare even when each part is ordinary. So a better statement of the test is that the resource improves results for anyone doing this job who holds the complements it needs, with the complements counted as part of the asset.

Performance enhancement is necessary but not sufficient. A restaurant needs reliable electricity, accurate books, and compliance with labor law. None of these is a source of advantage; every competitor has them.

A resource is hard to acquire when a rival would need a large amount of time or money to obtain it, or could not obtain it at all. I don't have a precise threshold, but my rule of thumb is that an investment above ten percent of the rival's revenue, or a wait of five years or more, makes a resource hard to acquire.

McDonald's scale is hard to acquire; a smaller chain like In-N-Out could not match it with any plausible investment in any plausible time. Because scale is also performance

enhancing, through purchasing, advertising and the cost of running a kitchen at volume, McDonald's scale is an alpha asset. The Nike name is hard to acquire, since no rival can use it, and put on a basketball shoe it would lift the sales of almost any shoe. It is an alpha asset. A cyclist's cardiovascular system, as chapter 1 showed, is partly inherited and only partly trainable, which makes it hard to acquire. Since it is also performance enhancing, it is an alpha asset too.

Some resources pass the second test and fail the first. Starbucks is headquartered in Seattle; it would take Luckin Coffee, a chain that opened in China, years and a great deal of money to move its headquarters to Seattle, and the move would not improve a single cup of coffee it sells. Or consider the name of Bjørn Dæhlie, the Norwegian cross-country skier. Because no other skier can use it, the name is hard to acquire; it did nothing for his speed on the trail. After he retired, he put the name on a line of sportswear, and in the Scandinavian market the name became an alpha asset for a company whose job is selling clothes.

Who Pays, and Who Gets Paid

The two tests leave out a third consideration: when a scarce, performance-enhancing asset is for sale at a price that fully reflects its value, its buyer earns nothing above the normal return. If everyone can see that a brand, a location or a patent is worth a hundred million dollars,

the seller charges a hundred million, and the buyer's return on the purchase is ordinary. (An asset that can be bought at a fair price is, almost by definition, not hard to acquire.) Above-normal returns from buying assets go to the buyer who expects more from the asset than the market does, who already holds the complements that make it worth more in their hands, or who is lucky. Barney made this argument in 1986, five years before his paper on the resource-based view, under the heading of strategic factor markets. For that reason, most of what follows is about building alpha assets rather than buying them.

A related question is who collects the return. When a firm's advantage rests on a resource that a person controls, the person can bargain for the value. Consider a software company whose results depend on one engineer: the company has an alpha asset only in the sense that the engineer has one; it pays for the asset in compensation and in the risk that they leave. Russell Coff made this argument in 1999: value created by hard-to-imitate resources often accrues to employees and other stakeholders rather than appearing as unusual returns to shareholders. If you work for someone else, you are in the engineer's position. Your scarce capability is your alpha asset, and who captures its value depends on where you work and on what terms.

Important Is Not Alpha

Many things are important without conferring advantage. An enterprise software company cannot succeed without an effective sales process, a restaurant without competent operations, or a bank without timely and accurate accounting. None of these is a source of sustained advantage. They are the price of admission: every serious competitor does them competently, and the differences among competent performers are small.

Product falls into this category more often than product people like to hear. Although a good product is important for every company, it is an alpha asset only in a few settings. The first is the short window after a disequilibrium, when the new product is dramatically better than anything else available. The second is an industry where a patent or a regulatory approval makes the product itself exclusive, as in pharmaceuticals, and the third is the small number of businesses where making the product is so hard that the capability to make it is the asset, as in leading-edge semiconductor fabrication. Elsewhere, every competitor must make a good product and none makes a decisively better one. Chapter 7 is about using the window while it is open.

With the two tests in hand, a manager or a founder can audit a business. The appendix contains the worksheet I use in class. It lists the categories of resource that most often turn out to be alpha assets: the solution itself

and the know-how behind it, the processes that deliver it, the brand, scale, the customer network, the organization's relationships upstream and downstream, its culture, and a residual category for location, balance sheet and the founder. For each resource you assign two scores from one to ten, rarity (how hard would it be for a rival to acquire?) and necessity (how much does the resource improve results, with how few substitutes?), and multiply them. In a worked example for McKinsey, the consulting firm, I give the firm's frameworks a rarity of three, since they are published and walk out the door with every departing consultant; its client relationships get a rarity of eight.

The scores are prompts for an argument, not measurements. A number on the worksheet means that someone has asserted a position about a resource that someone else can dispute with evidence. (That is more than I can say for most strategy discussions.) The worksheet also has a failure mode, the subject of chapter 3: if you fill it in after a company has succeeded, you will find alpha assets everywhere, because you already know the result. Two habits reduce the risk. First, score the resource by what a rival would have to spend to get it, a question about the rival rather than about the firm's history. Second, ask whether a venture investor who knew the firm's resources and not its results would fund a plan that depended on them.

Endowments

Chapter 1 said that organizations are born unequal; the two tests say how. A new company can start with an alpha asset it did nothing to build. In 2016 Microsoft decided to build its own business messaging product, Teams, rather than buy Slack. Because Microsoft had attached Teams to the installed base of Office 365, Teams reached twenty million daily users within three years, despite wide agreement (which I share) that Slack was the better product. The customer network was an endowment from the parent. When the singer John Legend co-founded a wine label, LVE, with Raymond Vineyards in Napa, Raymond made the wine; the alpha asset was Legend's name, which carried the label into mainstream distribution that a tiny fraction of new wineries ever reach. If you have that kind of name, congratulations. The rest of us start with speed, a better product for a while, and not much else.

A person's endowments are the same in kind and harder to change. By the same two tests, a capability or a position is an alpha asset for you if it improves your results in the work you want to do and if the people competing for that work cannot easily acquire it. For most of us the candidates are a skill that took a decade to build, a reputation, a set of relationships, and a position inside an organization that gives access to problems worth solving. You can score each one on the same worksheet. The difference from the corporate case is the objective. Whereas a firm's objective is, to a first approximation, return on invested capital, a

person's objective is not obvious. A capability that would win a contest the person does not want to enter does them no good. (I am a competent administrator, and I dislike meetings.) Chapter 9 returns to the question of what an advantage is for.

Assets Stop Being Alpha

The 1989 Tour de France came down to a final time trial of 24.5 kilometers from Versailles into Paris. The race leader, Laurent Fignon, started the day fifty seconds ahead of the American rider Greg LeMond. LeMond rode with a teardrop aerodynamic helmet and a pair of clip-on handle-bar extensions, recently introduced by Scott, that let him tuck his arms into a narrow position in front of the bike. Fignon rode bareheaded with his ponytail in the wind. LeMond beat him by fifty-eight seconds on the day and won the Tour by eight. Engineers who have since analyzed the ride agree that the equipment accounts for more than eight seconds. On that afternoon the helmet and the bars were alpha assets, but by the mid-1990s every professional rode a time trial with both.

How long does an asset remain hard to acquire? Almost never forever. Luckin Coffee opened in China with a mobile-first ordering system that had no counterpart at Starbucks. I once stood in a Luckin shop in Shanghai while a barista tried and failed to get the app onto my phone and finally gave me the coffee for free. For several years the system was an alpha asset; today every coffee chain in

China takes orders by phone. (Luckin's ordering system outlasted its accounting; in 2020 the company admitted that a large part of its reported sales had been fabricated.)

The central tenet of microeconomics is that competition drives markets toward equilibrium. An equally serious discipline could have been built on the opposite observation: at the time scales on which people and companies live, the world is never in equilibrium, because friction of every kind keeps it out. Customers are slow to switch, know-how is slow to spread even inside one company, and rivals are slow to notice, slower to copy, and slowest to copy well. Advantage exists in the gap between a new asset and its imitation, and friction sets the length of the gap. Chapter 8 asks what happens when the gap closes.

Lucky or Good?

A professional basketball player retires having shot something like ten thousand free throws. Each attempt is the same fifteen feet, the same rim, the same ball, and the result is unambiguous. Their career percentage is known to a decimal place. If they shot 90 percent and a teammate shot 75, nobody doubts that the difference is skill.

A partner in a venture capital firm leads perhaps thirty investments over a career, each unlike the others. The returns are drawn from a distribution so skewed that one or two investments typically account for most of a fund's gains, and the average time from the decision to a resolved outcome is more than seven years. If their investments returned 15 percent a year and a colleague's returned 8, no reasonable observer can say from the outcomes alone that they are the better investor.

A founder's record is usually one number: what the company sold for, or that it did not sell at all. That number contains the state of the market in the year of the exit, the competitors who did or did not appear, a decade of economic ups and downs, and a few hundred decisions the founder made. If the company sold for a billion dollars, nothing in the figure says what share of it was the founder.

Chapter 1 named three sources of a record: ability, circumstance and chance. When does a record tell us anything about the first? The inference can fail in two distinct ways. The first is statistical: the record may contain too little signal, relative to its noise, to support any conclusion about the person. The second is narrative: the explanation we attach to a record may be shaped by the record itself, so that what looks like independent evidence of ability is just the result restated in flattering language. The two failures have different remedies; for the first, I draw on work I did with my colleague Linda Zhao.

Two Dials

How much a record can reveal about ability depends on two quantities. The first is how noisy each outcome is: how much a single result varies, for a given level of true ability, because of everything that is not ability. The second is how many independent outcomes the record contains. A free-throw shooter sits at one extreme on both. Each attempt is nearly noise-free, and there are thousands of them. The founder sits at the other extreme, with a single outcome drawn from a distribution so wide that the same founder, replayed in a different year, might have produced almost any result.

The arithmetic is short. Let's suppose each actor has a persistent ability that does not change over the period we observe, and that each outcome equals that ability plus noise. Take the spread of ability across actors and divide

it by the spread of the noise in a single outcome; call the result the signal-to-noise ratio, SNR. Averaging N independent outcomes shrinks the noise by the square root of N . So the ratio of typical differences in ability to the noise left in an actor's average record is SNR times the square root of N . The statistical half of this chapter follows from that product, which I call the detection budget. The budget rises in proportion to SNR but only with the square root of N . Doubling the signal-to-noise ratio of each observation doubles the budget. Doubling the length of a career multiplies it by only the square root of two.

How big does the budget have to be? To declare with ordinary statistical confidence that an actor is better than average by some amount d , measured in units of the ability spread, the budget must exceed roughly 2.5 divided by d . The 2.5 is not my choice; it is the sum of the two numbers a statistician uses to detect an effect with 95 percent confidence and an 80 percent chance of finding it when it is there. (Anyone who prefers a stricter standard can raise it.) For a medium-sized difference, d of about one half, that means a budget of five. For a large difference, d of about 0.8, it means a budget of about three. Below those levels the record does not support the claim. The actor may well have the ability; the trouble is that an average actor with ordinary luck could have produced the same record.

Consider first a shooter with five thousand career free throws. The spread of true shooting ability across players is about 0.10 and the noise in a single attempt is about

0.43, so the signal-to-noise ratio is about 0.23; five thousand attempts give a square root of about 71, and the budget is about 16, well past the threshold of five. For a chief executive, the measured spread of true effect on annual firm returns is about 0.04 against year-to-year noise of about 0.18, a ratio of about 0.22; a tenure yields perhaps seven independent strategic decisions, with a square root of about 2.6, so the budget is about 0.6. Because the two ratios are nearly the same, the difference between a knowable domain and an unknowable one is almost entirely the number of outcomes a career can produce.

The first of two refinements is that detection and prediction are different questions. A record too weak to establish that an actor is good still shifts a forecast of their next result, and the shift has an exact size. A rational forecaster should keep a share of their observed deviation from the average equal to the budget squared, divided by one plus the budget squared. A budget of one keeps half; a budget of two keeps 80 percent; a free-throw shooter's budget of about sixteen keeps more than 99 percent. Nothing discontinuous happens at any threshold, so when I call a record "unknowable" in what follows, I mean that it barely moves a sensible starting estimate, not that it moves it by nothing.

The second refinement is that outcomes are rarely independent. A venture partner's thirty investments share the years in which they were made and the sectors they were in; a chief executive's forty quarters share one strategy and one economy. If any two outcomes have a correlation

of ρ , the usual correction says that N such outcomes are worth only N divided by one plus $(N \text{ minus one})$ times ρ independent ones, a number called the effective N . The formula has a ceiling: as N grows, the effective number of outcomes approaches one over ρ . At a correlation of 0.1, no career, however long, yields more than ten independent observations.

The Map

Linda Zhao and I placed ten domains on a map whose axes are the signal-to-noise ratio and the effective number of outcomes a career can produce. We used published estimates of how much of the variation in outcomes comes from ability where they exist, and stated ranges where they do not. Three domains anchor the knowable corner: career free throws, with a budget around sixteen; tournament chess, where a serious player accumulates thousands of rated games, the Elo rating system is built to measure skill from those results, and the budget is above thirty; and high-frequency trading, where the effective N runs to millions and the budget is in the hundreds. Geopolitical forecasting sits on the boundary, and only because of the forecasting tournaments that Philip Tetlock's group ran. They imposed a structure (hundreds of resolvable questions scored against the truth) that pundits in their natural state never face.

The three domains in which the largest sums of money and the most prestige are allocated on the strength of

records all sit in the unknowable region. The first is long-only active equity management, where per-year noise is far larger than the spread of true ability across managers, and a career contains perhaps thirty annual returns, giving a budget near two before correlation is counted and near one after. The second is venture capital, with a budget around 2.4 at twenty investments; the heavy right tail makes even that optimistic. The third is the non-founder chief executive of a large public company. A tenure of five to ten years contains on the order of five to fifteen independent strategic decisions. The measured spread in executives' effect on firm performance is small relative to the year-to-year noise in benchmark-adjusted returns, and the budget is around 0.6. Even an executive whose true effect is a full standard deviation and a half above the median cannot be reliably distinguished, from the record, from an ordinary executive who happened to run a firm in a good industry in a good decade. The economists Marianne Bertrand and Sendhil Mullainathan showed in 2001 that chief executives are paid for exactly such industry-wide luck.

The founder is the limiting case: with one outcome per actor, the variance observed across founders is the sum of the variance of ability and the variance of noise, and no data the domain generates can split the sum. So the record cannot pick out a skilled founder, or even say how much of the spread among founders is skill. In a study of venture-backed founders, Paul Gompers and his colleagues found that founders with a prior success were more likely to suc-

ceed again, roughly 30 percent against 18 percent for first-timers. That tells us something about serial founders as a group and nothing about which individual's success was skill. Chapter 4 returns to that finding, since skill is not the only mechanism behind it.

But surely the top of the distribution is different? The median fund manager may be indistinguishable from an index, but the manager who beat the market for fifteen years must be good. Two results say the opposite: the top is where inference from outcomes is least trustworthy.

The first is selection: the budgets on the map are computed for a single actor named in advance. A manager who comes to an evaluator's attention because of an extreme record has been selected from a population of thousands, and the probability that the best record among thousands was produced by luck is not small. The finance economists Eugene Fama and Kenneth French applied this logic to the universe of US equity mutual funds. They found that the spread of returns across funds looked very much like what chance alone would produce, with evidence of skill sufficient to cover costs confined to a thin tail. Laurent Barras and his colleagues, using methods for controlling false discoveries, reached a similar conclusion. The record that got the manager into the room is the record you must discount most.

The second result is stranger. Jerker Denrell and Chengwei Liu showed that in domains where outcomes are noisy and heavy-tailed (a few extreme results account for much

of the total), the most extreme performances are disproportionately produced by high-variance, unreliable processes rather than by high ability. A process that produces a consistent result rarely produces the very best result in any given year. A process that swings wildly does so more often and also produces the very worst. In such domains the top of the distribution should be actively discounted, not merely treated as uninformative. The hedge fund that returned 60 percent last year is more likely to be the one that will lose 40 this year than the one with the best process.

The next chapter develops a further complication: even where ability is real and detectable, the apparent effect of an exceptional actor includes the advantages that their prior outcomes generated. Morten Sorensen, using a structural model of how venture capital firms and startups match, estimated that about two thirds of the persistence in top-quartile returns reflects the selection of better deals by better-known firms rather than better decisions once the deals are in hand. The better deals come because of the earlier results, and the earlier results may or may not have been skill.

The Halo

The second way the inference fails has nothing to do with sample size. It would happen with perfect measurement and infinite data, because the problem is in the explanations rather than in the outcomes.

Phil Rosenzweig made the case in *The Halo Effect* in 2007. His central observation is that our descriptions of a company's leadership, culture, strategy and customer focus are not independent of its results. When results are good, observers describe the culture as cohesive and the leader as visionary. When the same company's results turn, the same culture becomes insular and the same leader arrogant. Rosenzweig's examples include companies whose practices were praised in the business press in one year and blamed the next, with nothing about the practices having changed. The problem is not that the descriptions are false but that they are not evidence. A variable that was measured by looking at the outcome cannot explain the outcome. Yet most of what passes for research on why companies succeed, including surveys of managers who already know how their firm is doing, measures its explanatory variables that way.

Of the eight further delusions Rosenzweig names, three bear directly on this book. Connecting the winning dots: selecting a sample of successful companies, finding what they have in common, and calling it a cause, without asking how many failed companies had the same traits. The delusion of lasting success: treating a decade of outperformance as proof of a formula, when the base rate of regression is high. And the delusion of absolute performance: judging a company by its results rather than by its results relative to what rivals did in the same conditions. (A company can do everything right and lose to a rival who did it better.)

Rosenzweig's principal example is uncomfortable for me. It is *Good to Great*, by the business writer Jim Collins, whose best-known image is the flywheel; chapter 5 of this book is built around flywheels. Collins selected eleven companies whose stock had beaten the market for fifteen years after a transition, compared them with companies that had not made the leap, and reported the traits the eleven shared. Two of the eleven, Fannie Mae and Circuit City, failed within a decade of the book's publication. Collins's method could not have caught that. He identified the traits by looking at companies already known to have succeeded and chose the comparison firms to be similar in everything except the result. Whatever the eleven shared, his sample cannot say whether it caused the results, how many firms with the same traits failed, or whether the traits were themselves descriptions of success.

Strategy scholars raised the same objection against the resource-based view six years before Rosenzweig raised it against the business press. Richard Priem and John Butler argued in 2001 that if a valuable resource is defined as one that improves performance, then the claim that valuable resources improve performance is empty. In practice, they argued, the resources of successful firms get identified by their success. The alpha asset worksheet in the appendix is exposed to exactly this. Fill it in for a company you already know has done well, and every resource looks rare and necessary.

Three habits protect the worksheet and this book from the halo. State the mechanism before the result: say how the resource lowers cost or raises willingness to pay, in terms that could be checked without knowing whether the firm succeeded. Score rarity by what a rival would have to spend to get the resource. And look for the comparison the halo hides: the firm with the same resource that did not succeed, or the firm without it that did. The same habits apply to every story in this book, including the ones about my own companies. Where I cannot name the mechanism, point to an observation that does not depend on the result, and say what else might explain the outcome, the story is decoration.

None of this means that management has no effect. Two findings are the kind of evidence the halo cannot contaminate. When Nicholas Bloom and his colleagues randomly assigned management consulting to Indian textile plants, the treated plants became about a sixth more productive in the first year. The plants were chosen at random, so nobody's view of which plants would succeed could color the comparison. When Morten Bennesen and his colleagues studied what happened to firms whose chief executive was hospitalized, the firms' profitability and investment moved. So the person at the top has an effect, though the study says nothing about which executives are better. Management makes a difference; which manager made the difference, in any particular case, is the question the record usually cannot answer.

Researchers fall for a welcome result too. In late 2024 a working paper from a doctoral student at MIT reported a study of artificial intelligence in a materials-science laboratory. The tool nearly doubled the output of the best researchers while the weakest third gained almost nothing, a result that supported, precisely, the thesis that AI rewards judgment at the top. It was widely circulated and admired until May 16, 2025, when MIT's economics department announced that it had no confidence in the provenance, reliability or validity of the data, or in the veracity of the research, and asked that the paper be withdrawn. I return to the question of what AI does to advantage in chapter 6, with the evidence that did survive, which is weaker and points partly the other way.

A population is different: across many actors in the same domain, the total number of outcomes is far larger than any one career contains, and the population can support inferences that no individual record can. The spread between return on capital and its cost persists across thousands of firms; management practice moves productivity across dozens of plants; serial founders succeed more often than first-timers across thousands of ventures. Each of these is knowable at the level of the population and unknowable at the level of the case.

Medicine has used that asymmetry for seventy years. The Framingham Heart Study followed thousands of residents of one Massachusetts town from 1948, recording blood pressure, cholesterol, smoking and the rest at regu-

lar visits, and waiting for the sparse, unambiguous cardiovascular events. Doctors then apply the relationship estimated across the population to a new patient, using their observable features alone. Nobody waits for the patient's own heart attacks to accumulate before estimating their risk. The same approach works in the unknowable region of the map. A board that cannot learn from three years of results whether its chief executive is good can still ask whether they have adopted the practices that, across hundreds of executives and firms, predict better results. A limited partner who cannot tell skill from deal flow in two fund records can still compare the firm's process with the process features that predict returns across the population of funds. The map sets the weight to put on the record versus the population. Where the budget is high, trust the record; where it is low, trust the population and treat the record as a small adjustment.

The medical comparison also shows the limits of the method. Features validated across a population are, by construction, generic. A practice can be validated across hundreds of firms only if it means roughly the same thing in each of them. Such studies therefore concentrate on monitoring, targets and incentives rather than on the idiosyncratic strategic choices that strategy scholars care most about. An idiosyncratic choice is an observation without a population. The method rescues evaluation of the transferable part of ability and leaves the idiosyncratic part where the map put it. And population features can be gamed, as any metric can. The method does not escape

that problem; it moves the target from an outcome that is equally gameable and far less informative.

The domain can also be changed. Although forecasting pundits are unknowable, forecasters in Tetlock's tournaments were detectable: the ones selected into the top two percent after one year stayed there in the following two, defying the regression that luck predicts. The forecasters were the same; the structure, many resolvable questions scored honestly, is what moved the domain across the boundary. Any organization that wants to know who is good can often do better by asking what structure would make the question answerable than by asking it of the records it already has.

An N of One

I have founded several companies, and I am a professor. Both records are the kind this chapter says are unknowable from the outside. Two companies I co-founded survived and grew. Others did not. My early academic papers were cited more than most of my classmates' papers. Put any of this through the map and the budget is well below three. Anyone judging me by it should shrink the record most of the way toward the average. In the chapters that follow, where my account of a decision cannot be separated from my knowledge of how it turned out, I say so.

That leaves the question chapter 1 raised with the salesperson's territory. Suppose the record cannot tell us whether the venture partner, the executive, or the

founder earned their early success. The early success still happened. It brought them deals, recruits, capital and invitations. Are those advantages now, and do they belong to the person or to the position? Because those questions concern the present rather than the past, the record can help with them. The next chapter takes them up.

How Does Luck Become an Asset?

In 2004 three researchers at Columbia built a website on which visitors could listen to forty-eight songs by unknown bands, rate them, and download the ones they liked. More than fourteen thousand people took part, each assigned at random to one of nine worlds. In one world, each visitor saw only the songs; in the other eight, visitors also saw how many times each song had already been downloaded by the people before them in that world. The eight social worlds started identical and were sealed from one another.

The results in the eight social worlds diverged. A song called "Lockdown," by a band called 52metro, finished first in one world and fortieth in another. In the social worlds, the gap between the most and least downloaded songs was far wider than in the world without social information. Which songs ended up at the top was also far less predictable. Quality still counted for something: the songs that visitors in the independent world rated highest rarely did badly in any social world, and the ones rated lowest rarely did well. But for the songs in between, which were most of them, an early run of downloads that happened to land on one song rather than another was enough to

carry it to the top of its world and leave an equally good song near the bottom of the next.

Nobody who looked at the final chart in a single world could have told which part of a song's position was the song and which part was the order in which the first few hundred visitors happened to arrive. Chapter 3 was about that limit. But the experiment shows something else as well: the early downloads, however they came, were real. In the world where "Lockdown" led, it kept leading, because every new visitor saw it leading. An accident at the start had become a position that kept the song on top from then on.

The Mechanism Has a Name

The sociologist Robert Merton described the pattern among scientists in 1968 and borrowed its name from the Gospel of Matthew (unto every one that hath shall be given). A discovery made jointly by a famous scientist and an unknown one is credited to the famous one. The credit brings the famous scientist more resources and more attention; those produce more discoveries, and the discoveries bring more credit. Merton called it the Matthew effect. The general term is cumulative advantage: a process in which an early difference in outcomes, from whatever cause, increases the probability of later favorable outcomes.

Stated as a mechanism, in the manner chapter 1 asked for, cumulative advantage runs through four channels,

each of which can be observed without knowing how the story ends. The first is resources: an early success brings money, and money buys the next attempt (a second fund, a second product, a larger lab). The second is attention. An early success makes the actor visible, and visibility lowers the cost of every later transaction; customers call instead of being called. The third, and I think the most important of the four and the hardest to see, is access. An early success changes the set of opportunities the actor is offered. Better deals come to better-known investors, better students to better-known professors, and harder problems to the engineer whose last project worked. The actor's later record is then drawn from a better pool, and a better pool produces better outcomes with no change in the actor at all. The fourth channel runs through the actor's rivals: those who lose early often stop trying, and the field thins.

None of the four requires the early success to have been earned; all four require it to have happened. That is the sense in which luck becomes an asset: the position that an early outcome creates passes the two tests of chapter 2. It improves expected results through the four channels, and a rival cannot easily acquire it, because the only way to get it is to have had the early outcome.

Luck Given at Random

The music experiment showed cumulative advantage inside an artificial market. Two later studies looked for

it in real ones, using the cleanest method for separating the position from the person: giving the early success away at random, or finding a setting where something close to chance gave it away.

Arnout van de Rijt and his colleagues, in 2014, chose four online settings where success is visible and accumulates. They donated to randomly selected projects on a crowdfunding site, gave randomly selected editors on a collaborative encyclopedia a status award, rated randomly selected product reviews favorably, and signed randomly selected petitions. In every setting the recipients of the random early success went on to attract more subsequent success than the matched controls who received nothing. Although the gifts were small and the recipients did not know they were arbitrary, they worked. The experiment also found a limit: a larger initial boost did not produce a proportionately larger later gain. The second helping did less than the first.

The second study used a natural experiment in science funding, made possible because the Dutch research council awards early-career grants on a ranked score with a fixed cutoff. Thomas Bol, Mathijs de Vaan and van de Rijt compared applicants who scored just above the line with those who scored just below it, two groups whose proposals were, by construction, almost indistinguishable in quality. Over the following eight years the winners accumulated more than twice as much research funding as the near-miss losers. Part of the gap came from evaluators, who looked more favorably on later applications from peo-

ple with a grant on their record. Part came from the applicants, since a substantial share of the near-miss losers stopped applying for the next grant at all.

Because the experimenters assigned the early boost in the first study, and a cutoff assigned it among nearly equal proposals in the second, the later advantage cannot have been a hidden signal of ability in either.

Venture capital, where I think confusing the position with the person is more expensive than in any other business, shows all four channels at once.

Consider a two-partner firm whose first fund, through a combination of judgment and the particular year in which it was raised, contains one company that returns the whole fund several times over. (This firm is hypothetical, assembled from how the industry works rather than from any firm's record.) The second fund is larger and raised faster, because limited partners, the investors in the fund, allocate on records. The partners are now invited to see companies before other investors do, because founders who can choose their investors choose the ones with a famous hit. When the firm's name is on a round, the company finds its next round easier to raise and so is more likely to succeed, and its success adds to the firm's record. Rival firms that missed the hit lose partners and sometimes fold. All four channels are now running, and the partners' second fund is drawn from a better pool of companies than their first.

How much of a venture firm's apparent skill is that better pool, and how much is what the firm does once a deal is in hand? Morten Sorensen, in a 2007 paper in the *Journal of Finance*, set out to measure it. His estimate, from a structural model of how investors and companies match, is that about two thirds of the persistence in top-quartile venture returns comes from better firms getting access to better companies, and about one third from the firms' influence on the companies they fund. The split depends on the model's assumptions, so I would not repeat it as a law. Still, if the estimate is even roughly right, most of what a limited partner is paying for when it pays for a top firm's record is the firm's position. The position is a different asset from the partners' judgment. The position may attach to the franchise rather than to the people, in which case it need not survive partner departures, successions and increases in fund size, precisely the occasions on which a limited partner decides whether to reinvest.

The founder evidence has the same shape: Paul Gompers and his colleagues found, in a 2010 study, that founders with a prior venture-backed success were about 30 percent likely to succeed again, against 18 percent for first-time founders. So an investor who backs the serial founder is making a better bet. But skill is only one of the channels the paper identifies. A founder with a prior success attracts better co-founders, better early employees, and better investors, and is given more time by all of them. Part of the second success comes through the

position the first one created, whether or not the first one reflected the founder's skill.

Three Questions, Not One

People usually ask about an early success as if it raised one question, when it actually raises three, with different answers.

The first is origin: did the actor earn the early success? Chapter 3 argued that in most consequential domains the record cannot answer this. Even where it could, the answer would not settle the other two.

The second is current effect: does the position created by the early success improve expected results now? This one can be answered, the same way chapter 2 tests any alpha asset, by stating the mechanism and checking it against observations that do not depend on the outcome. Whether the venture firm sees deals before its rivals can be checked by counting the inbound. Does the professor's name on a grant application raise its score? The Dutch study measured it. Whether the salesperson from chapter 1 has a territory full of customers whose equipment is due for replacement is in the service records. These are facts about the present.

The third is portability: does the advantage belong to the person or to the position? The cleanest evidence is Boris Groysberg's study of star securities analysts, the ones ranked at the top of their sector by institutional investors. When a star moved to a different firm, their

ranked performance usually fell and often stayed down for years. The fall was smaller when they moved with their team and when the new firm had capabilities comparable to the old one. Much of what had looked like the analyst's skill was the firm's research platform, its relationships with the companies they covered, and the colleagues who fed them. The moves were not random, so the study does not give a universal penalty for changing employers. It does show that what travels is a question with an empirical answer. I suspect the answer is often less than the star believes, and less than the firm that hires them is paying for.

The same three questions apply to the salesperson with the inherited territory. Whether they earned it is unknowable and, for most purposes, beside the point. Whether the territory is an advantage now is a matter of record. Their next employer should ask whether the territory goes with them before matching their number. So should they, before they assume that the number is theirs.

My Own Case

In 1995 Steven Eppinger and I published a textbook on product design and development, which engineering and business schools adopted in numbers that surprised us. Neither of us could have predicted those numbers from its contents, a clear organization of what good product development teams already did. What followed ran through the four channels. The book made me visible to people I

had never met. Executive education programs asked for me because their participants had read it, and companies asked me to consult. Doctoral students chose to work with me, and the ones who chose me were, on average, stronger than the ones I would have been able to recruit without the book. Each of those brought the next project, and the next project was drawn from a better pool than the one before.

I can state the mechanism and point to observations that do not depend on the outcome: the adoption numbers, the invitations, the composition of the students who came. What I cannot do is separate the book's success from the fact that it appeared when it did, in a field that was about to grow, at a school with a name that opened doors. A different book of equal quality, published three years later or from a less visible address, might have sold a tenth as many copies, and I would have been a different professor. I believe the book was good, though I do not know whether its reception was a measure of how good. I do know that most of what the reception brought me was position rather than skill. The position was real, and I used it.

Then there is portability. Some of what the book brought traveled with me (the students, the method, the relationships with the companies), and some stayed with the address. When I left the university's administration and the daily company of its colleagues for a company I had co-founded, the invitations did not stop, but they changed in kind. Some of what had looked like my reputa-

tion turned out to have been the institution's. I had been warned, by Groyberg's analysts among others, and the warning was accurate.

A board, an investor, or a hiring manager who cannot know whether a record was earned can still learn whether the position the record created is an asset now, and can find out whether it travels. Pay for position when the position comes with the deal; pay for the person only when there is evidence that the person is what you are buying. The limited partner paying premium terms for a franchise's deal flow may be making a sound investment. It is a different investment from paying for the partners' judgment, with a different risk.

And what should the actor do with an early success? Others will read it as skill, and so will the actor. Chapter 3 gave the reasons that reading is mostly unwarranted. The actor's advantage is the position the success created, and positions decay. The deals keep coming only while the record is fresh, the recruits keep choosing the professor only while students keep hearing the name, and the territory keeps yielding only while the customers keep needing the equipment. So the useful question is what to build while the position lasts that will remain after it fades. The venture firm converts a hit into a platform of relationships and a brand that outlast the partners who made the hit. The professor converts a book into students who carry the method. The salesperson converts an inherited territory into expertise in the customers' problems, which is theirs.

That conversion is the subject of the next chapter. The mechanisms by which advantage accumulates, which I call flywheels, work the same way whatever produced the first push. A customer base grows from the customers a firm already has, whatever brought them; learning accumulates with volume, whatever produced the volume; and a brand strengthens with every satisfied customer, whether the first customers came from merit or from a famous name on the label. Although in most cases nobody knows the answer to the origin question, for the flywheels it makes no difference.

How Does Advantage Accumulate?

In 2015, taking what I thought was the advice of the investor Warren Buffett, I bought shares in Wells Fargo, one of the four largest banks in the United States. Oops. Over the next five years the bank admitted to opening millions of accounts its customers had not asked for, paid fines, lost its chief executive, operated under an asset cap imposed by its regulator, and watched its share price fall while the rest of the sector rose. I sold (and the share price recovered after I was gone). Yet Wells Fargo has operated continuously since 1852 and is still there; it seems pretty nearly impossible to kill this organization. Few companies last as long as Wells Fargo, but incumbents in general are remarkably resilient. A company that survives infancy is very likely to live for decades, either as an independent firm or as a division of whoever buys it.

An organization is a machine for doing a job. It exists because it can do the job for its customers better or more cheaply than they could do it for themselves. For a company, economic viability is a pretty simple formula: revenue above cost, with revenue rising and cost falling over time relative to rivals. But what allows incumbents to keep that formula running through scandals, recessions and

their own mistakes? I call the mechanisms that sustain them flywheels. In a machine, a flywheel is a heavy wheel that accumulates energy from an uneven power source and smooths the output. In an organization, the same word fits the mechanisms that store the results of past effort with customers, cost, quality, reputation and capability, and release them as needed.

Most organizations, I believe, have five flywheels. Each is a stock that grows as the business runs and a candidate alpha asset, since each improves results and each takes years to build. Which of the five counts for most depends on the industry. As chapter 4 put it at its close, the flywheels do not ask where the first push came from.

Customer Network

Very few businesses serve each customer just once; cataract surgeons and MCAT tutors come to mind. Nearly everyone else acquires a customer and then earns from the relationship over years. This is why marketing teaches customer acquisition cost and lifetime value as a pair.

When we were just starting out at MakerStock, a company I co-founded that sells pre-cut acrylic and plywood to people who fabricate things with laser cutters, we bought search advertising. Two or three of every hundred people who clicked placed a first order, and about one in seven of those placed a second and became a repeat customer. So for every three hundred clicks or so we acquired one repeat customer, and most weeks we added a handful.

By the second year we had a few hundred of them, and we worked hard to keep them: we'd special-order materials, send samples of new materials, and offer cutting services. There is really no shortcut to a customer base, since customers are added one at a time. Once accumulated, though, the stock persists even when the effort stops; satellite television companies still bill customers who installed a dish on the roof twenty years ago.

In some categories the customer network does something the other flywheels do not. It accelerates. When the value of a product to a new customer rises with the number of customers who already use it, the product has a positive network effect, and each customer acquired makes the next one cheaper to acquire. A payment app is worth little when only your children use it and a great deal when the man who plows your driveway does. Platforms that assemble many sellers for many buyers, such as an app store or an online marketplace, have network effects on both sides. In a category with strong network effects the market tends to tip toward one provider. Once it has tipped, the installed base is among the most powerful alpha assets a firm can hold. A rival cannot buy an installed base, nor can it build one against a tipped market. Good luck launching a mobile operating system to rival Apple's iOS and Google's Android.

I have two cautions. First, network effects are mostly a property of the product, and a product that lacks them usually cannot be given them. I have tried to think of a way for MakerStock to benefit from one, and I can't, because

the value of the right sheet of acrylic on the right day does not depend on how many other people are buying acrylic. Second, switching costs, often listed beside network effects as a source of lock-in, are a weaker asset than they look. Customers resent them, competition tends to push markets toward interoperability, and a customer who stays only because leaving is painful is already looking for the exit.

Brand Equity

My education is in mechanical engineering, my father is a chemical engineer, and my brothers are engineers. We worship at the church of science and reason, and at first glance brands are ridiculous. In one of my first papers as a professor, Taylor Randall and I examined prices in the mountain bike industry. At the time nearly every frame was welded in a handful of Asian factories, and models from different brands often differed in nothing but the decals. Even so, we found that the brand name alone accounted for hundreds of dollars of price. The same thing can be observed in a Costco every day. Shoppers pay half again as much for Tito's vodka as for Kirkland, Costco's own brand, and even the Kirkland bottle commands a premium over the unbranded one. (My preferred cocktail is vodka mixed with a potent home-made kombucha, and I am 100 percent certain I could not taste the difference in that use.)

There are three logical reasons people pay for brands. First, a brand is a credible promise about qualities the buyer cannot inspect, since the owner of a valuable name has more to lose from an impure batch than an anonymous producer does. Second, a brand reduces search cost. Yes, the shopper could read reviews and run tastings, or they could spend fifteen seconds in the aisle and grab the Tito's. Third, a brand is a signal to other people; the bottle on the shelf where guests can see it is chosen differently from the bottle in the cupboard. Put the three together, and even for goods of identical intrinsic quality, the stronger brand commands some combination of higher price and higher share.

Brand equity is the accumulated value of the name, and for many companies it exceeds the value of everything physical they own. It is a flywheel because it grows with time and volume. Awareness rises with every unit sold and every year in the market, longevity itself becomes a signal (count the firms that print their founding year on the label), and improved products feed the name. Bad conduct drains it, as Wells Fargo shows, but the drain is slow enough that an established brand can absorb a scandal that would end a newcomer.

Cost Efficiency: Learning and Scale

Have you ever hung a sheet of drywall? Drywall is a paper-faced gypsum board, about twenty dollars a sheet, heavy, fragile and dusty to cut. A novice takes about an hour to

hang the first sheet, while a crew of three professionals hangs twenty an hour. It's crazy how they do it. I once watched an installer with short stilts strapped to his boots balance a sheet on his head, screw it to the ceiling, find the electrical box hidden behind it from the bulge in the sheet, and cut the opening with a hand router, all in about three minutes.

The learning curve is the empirical regularity behind that crew: with each doubling of cumulative units, the time per unit falls by a constant percentage. On an 80 percent curve (and such curves are common), if the first sheet takes sixty minutes, the second takes 80 percent of that, or forty-eight minutes. The fourth takes thirty-eight, the eighth thirty-one, and the sixteen-thousandth, fourteen doublings later, about two and a half minutes.

I use the fancy academic term empirical regularity because the learning curve is not a law derived from anything; it just works out that drywall, coronary bypass surgery, airframe assembly and insurance claims processing all follow it. The individual improvements are lumpy and unpredictable: the stilts and the router are big steps, positioning the sheet with two nails before driving the code-mandated screws is a small one, and a mechanical pencil gives a barely detectable edge. Blur your eyes, though, and the steps merge into a smooth curve.

Two features of the curve bear on advantage. First, because each percentage improvement needs a doubling of cumulative volume, the curve has diminishing returns, and a firm with a million units behind it improves slowly.

Second, the rate of descent depends on deliberate effort. Frederick Taylor's time studies, the statistical quality methods that Walter Shewhart developed and W. Edwards Deming carried to Japan, and the production system that Taiichi Ohno and Eiji Toyoda built at Toyota are all techniques for descending the curve faster than repetition alone would. A firm that combines a head start in volume with a better improvement process has a cost advantage that compounds.

Scale economies are distinct from learning and often confused with it. For twenty-four years I have owned a small business, Xootr, which makes adult kick scooters. When we first started, my brother machined the aluminum hubs for our wheels in a barn in New Hampshire. Those hubs cost 6.14 dollars each when he was making ninety-six a week. At our peak we used 3,200 hubs a week, and at that rate they cost just 2.14 dollars. Some of the reduction was learning; for instance, we figured out that sandblasting the aluminum gave a better finish than polishing it, for less work. But most of the reduction came from switching from machining to die casting, which requires a custom mold (at the time about forty thousand dollars) that makes no sense at ninety-six hubs a week and every sense at 3,200. Scale economies arise when a higher rate of output justifies a different process, usually one with more capital and less labor per unit. They also extend upstream, since a large buyer gets better prices from its suppliers, who enjoy scale economies of their own. Services have weaker scale economies than facto-

ries, but specialization still produces some, until coordination costs overtake them.

Both learning and scale are flywheels: lower cost permits lower prices, lower prices bring volume, and volume brings more learning and more scale. Both are hard to acquire, because the only route to volume is volume. They are also stickier than they look. Can a way of hanging drywall really be hard for a rival to acquire? The strategy researcher Gabriel Szulanski studied the transfer of proven practices between units of the same company, where everyone wanted the transfer to succeed. He found that even simple practices moved slowly or not at all. Now imagine a rival trying to copy the same know-how from the outside, against its owner's resistance; that takes years.

Product Quality

When I was in college I owned a series of beat-up used cars, none of which cost me more than a thousand dollars: a 1969 Plymouth Fury, a 1970 Plymouth Valiant, a 1975 Toyota Corolla and a 1979 Chevrolet Nova. Each was a little over ten years old when I bought it, and I spent a dozen Saturdays a year under the hood of whichever one I owned. Today a ten-year-old car runs like a new one, and I'm pretty sure most of my students do not own the deep socket needed to change a spark plug.

Some of that improvement came from Toyota, Honda and Nissan making reliability a dimension of competition

in the 1970s and 1980s. Most of it came from the ordinary process that every car company and every supplier has run for half a century: collect failure data, find root causes, rank the fixes by cost and impact, and implement them in order. Product quality is a flywheel because every unit delivered produces information about the next one, so the firm with the most units and the best process reaches the highest quality first. Like learning, quality improvement has diminishing returns: in 1970 reliability was a decisive advantage, whereas today new cars vary too little in reliability for it to be one. So the product-quality flywheel confers its largest advantage while a category is growing, and the advantage erodes as the category matures.

Organizational Capability

Of the five flywheels, this one is the most exposed to the halo of chapter 3, in which a successful firm's culture gets described by its success. So what would make an organizational capability more than a halo?

The Ritz-Carlton hotels, a unit within Marriott, have for years been at or near the top of the luxury segment in independent guest-satisfaction surveys. Part of the capability behind that is a system, in which any employee can record a guest's preferences and the record follows the guest to the next property. Part is a culture of service, built over decades, that attracts and keeps people who like delighting strangers. Marriott operates several other luxury brands, none of which has matched the Ritz-

Carlton's results, even though the sister brands have the same parent, the same capital and access to the same system. Because that failure to match is the pattern Szulanski found inside companies, an organizational capability, though intangible, can pass the hard-to-acquire test.

The capability is a flywheel for two reasons. Like the other four, it improves with repetition; a service organization learns from every complaint. And it has a positive feedback loop of its own, in which a stronger culture of service attracts and retains employees who want to work in one and who in turn strengthen it. The same loop can run on other capabilities; for instance, a company known as a good place to work attracts the engineers who make it a better product company, and a company known for environmental seriousness attracts people with expertise in it. Diversity, process discipline, speed and innovation can each run the loop.

Phil Rosenzweig, author of *The Halo Effect*, would raise the obvious objection, and I raised it against myself in chapter 3: a capability named after the fact explains anything. A firm succeeds and we discover its culture of customer focus; a firm fails and we discover its culture of complacency. My defense is procedural: name the capability as a set of observable practices, not an adjective. Point to an observation that does not depend on the result, such as the sister brand that could not copy it, the survey conducted before the financial result was known, or the practice the firm invested in at a visible cost before it paid. And include the counter-case.

Sierra Trading Post, a retailer of discounted outdoor gear, was founded by a man who printed Bible verses on its order forms and declared its ethics to be those of its owners' faith. The practice was real, costly and sustained. It produced a loop of its own, attracting employees who shared the faith and who in turn reinforced the practice. But it was not an alpha asset. Its customers were back-country skiers and climbers who, judging by the online forums of the day, found it somewhere between irrelevant and off-putting, and when the company was acquired, the verses disappeared. Cotopaxi, a gear company founded a generation later around poverty reduction in Latin America, runs the same kind of loop with values its customers share. The fifth flywheel is an alpha asset only when the capability feeds one of the other four.

How many such loops can one organization run? One or two, not five. The psychologist Anders Ericsson's research on expertise (the real version, not the ten-thousand-hour slogan) makes the point for people: being the best in the world at one thing consumes everything, and being the best at two is nearly unheard of. Organizations face the same limit, with finite management attention and finite cash. A single capability is also easier to communicate; the sign James Carville hung in Bill Clinton's 1992 campaign headquarters read "the economy, stupid," and almost nobody remembers the other two lines. And a single objective is easier to act on. A Ritz-Carlton bellhop asked to store twelve bags knows what to do when the orga-

nization's value is guest satisfaction, and does not when it is guest satisfaction balanced against cost.

The best evidence that organizational capability is causal rather than a halo comes from the experiment chapter 3 cited. When the economist Nicholas Bloom and his colleagues randomly assigned months of management consulting to textile plants in India, the treated plants adopted practices in quality control, inventory and efficiency. They became about a sixth more productive in the first year. When the researchers returned nine years later, some of the practices had eroded and some of the gain remained, so the intervention acted more like a single injection than a flywheel. A practice is an asset. A capability that keeps adopting and refining practices is a flywheel.

Flywheel Physics

Whether a later entrant can catch an incumbent depends on how each flywheel responds to speed.

Brand equity shows roughly constant returns. It is built brick by brick, with no reliable shortcut, and the firm that started earlier and invested continuously has more of it (think of Starbucks in coffee, Adidas in sportswear, Sony in electronics). A later entrant can match it only by investing for as long.

Cost efficiency and product quality show diminishing returns, by the arithmetic of the learning curve: the twenty percent that comes easily between one hundred

units and two hundred comes glacially between one million and two million. So a newcomer with some initial boost, usually a better product, stands a chance of reaching competitive cost and quality against an incumbent whose own improvement has slowed.

Customer networks show constant returns in most industries, where the ten-thousandth customer takes about as much effort to win as the hundredth. In industries with network effects they show increasing returns, since each customer makes the next one cheaper. Increasing returns are the exception and also the most important case. A firm that gets a toehold early in a category with strong network effects and keeps its base growing becomes nearly impossible to dislodge, for as long as the category lasts.

Organizational capability can show increasing returns as well, through the talent loop: the better the culture, the better the people it attracts, the better the culture. Systems thinkers call loops of this kind snowballs. Chinese has a four-character phrase for the general idea, which my first teachers used to encourage me: the more, the better it gets, and the better it gets, the more.

The single word flywheel covers three different things. Some of what I have described is accumulation of a stock: more customers, more cumulative units, more years of brand. Some is learning, in which later units are made better than earlier ones, and some is self-reinforcement, in which the existing stock makes the next addition cheaper. Although the three can coexist, a customer base can be

large without network effects, in which case it is a stock that protects but does not accelerate. Stocks also depreciate: learning that is not used decays, as the management scholar Linda Argote and her colleagues found in shipyard production data; a brand that is not fed fades. So for each flywheel in a particular business, I would ask what grows, why that makes the next round easier, what has to be invested to keep it growing, and what eventually limits it.

Not every alpha asset is a flywheel. If a famous singer founded a fragrance line, the singer would be an alpha asset. But the boost they gave the business would rise and fall with their celebrity rather than grow with the business. A patent is an alpha asset for as long as the patent runs, and a regulatory license is one until the regulation changes. These are endowments, one-shot injections of advantage of the kind chapter 4 described. Who wouldn't love to receive one? But I know of no process for acquiring geese that lay golden eggs. Flywheels, in contrast, can be built from nothing by any firm willing to spend the years, and once spun up they store energy that an endowment cannot.

Then there is the tool that everyone now has. A frontier language model is performance enhancing, but because it is not hard to acquire, it fails chapter 2's second test, as electricity does. What it changes is friction. A product that can be imitated in a week has a shorter life as an alpha asset than one that took a rival a year to copy. So the

product-quality flywheel protects less than it did, while the flywheels that a model cannot shorten are worth relatively more: the customer base, the brand, the relationships, and the organization's capability to deploy the tool well. Is a firm's use of the model a fifth flywheel or a halo? That is settled the same way as for any other capability, by naming the practices, pointing to the sister division that could not copy them, and waiting for the result.

Where Should You Apply Your Strengths?

In July 2004 the New York Times reported from the Milan men's wear shows that the posse of American editors and buyers in the front rows at Dolce & Gabbana, Gucci and Burberry, the men whose job is telling other men what to wear, had, to a man, worn their shirttails untucked. Seven years later a company called Untuckit was founded to sell shirts cut to be worn that way. In 2013 the chief executive of Apple appeared on stage at the company's developer conference with his shirt untucked. Google searches for untucked shirts peaked in 2017. By 2020 the Banana Republic factory outlets had a section labeled "untucked oxford shirts."

Fashion is the fruit fly of commerce: the generations are short, so you can watch the life cycle of a category in a decade. Every category goes through the same four phases, which I call sense, solve, scale and sustain. The alpha asset that wins in one phase is not the one that wins in the next. Sensing is the recognition of a disequilibrium: a need without a solution, or a new way to meet an old need. At this point the idea, which usually occurs to several people at once, is a hypothesis clouded by uncertainty. Solving is the scramble of rival attempts, most of

which we never hear about, that lasts until one or a few find a fit between the solution and the need and, in hardware, a dominant design emerges. Scaling is when the survivors grow, the supply chain forms, the customer learns the category, and the flywheels of chapter 5 begin to turn. Sustaining is the long period, often decades, in which an oligopoly of survivors runs the flywheels and earns most of the category's profit. (A bell foundry in Italy has been sustaining since 1339.)

The most consequential transition is the one from solve to scale. During solve, the only alpha asset a new firm has is a solution better than the alternatives for some customers, plus the speed to get it to them. During scale and sustain, every surviving rival has a solution, and the assets are brand, scale, the customer network and capability. A firm that is good at one phase is usually bad at the next. The organization theorist James March named the general tension exploration and exploitation (looking for the oil versus extracting it) and observed that organizations optimized for one are poorly built for the other. The habits that run a sustaining business (key performance indicators, process control, structured hiring, global expansion) would have prevented the business from being founded. Chapter 7 is about what a founder should do during the solve phase. The prior question is where to be at all.

Fit

The principle is short: seek the setting in which the assets you have, or can build, are rare and necessary. It holds for a company choosing a market and for a person choosing a job, for the same reasons.

By the definition in chapter 2, an asset is an advantage only relative to the people competing for the same work in the same setting. One consequence is easy to state and hard to accept: you do not need to be the best in the world at anything. You need a combination of capabilities that is unusual among the people competing for this particular work, in a setting where that combination improves results. Economists call the general form comparative advantage: a person who is worse than a rival at two tasks can still be the right person for one of them, if the gap is smaller there. A mechanical engineer who writes clearly is an ordinary engineer in a room of engineers and an ordinary writer in a room of writers. In a room of people who must explain engineering decisions to executives, they have no competition.

Fit also runs the other way, because a setting supplies complements. The portability finding in chapter 4 was that much of a star's performance is the platform they stand on. The star securities analysts in Boris Groysberg's study lost rank when they moved to new firms, because the research platform, the company relationships and the colleagues stayed behind. So the question for a person is not only where their capabilities are scarce but

where their complements are present, and who owns them. A firm supplies capital, customers, reputation and colleagues; the person supplies judgment and effort. How the returns are divided, as Russell Coff observed, is a matter of bargaining power. The engineer whose results depend on the firm's platform has less of it than they think; the engineer who holds the relationships has more.

For a Company

The five flywheels of chapter 5 are built in one market and can then be used to enter another. Which directions are open depends on where the firm's alpha assets reach.

Geographic expansion usually comes first. The jobs to be done are similar everywhere; everyone eats, moves, communicates and wants to be entertained. What differs is the friction: regulation, logistics, language, taste. My coauthor Lele Sang and I studied eight multinational companies entering China and found a short list of preconditions that generalize to any new geography. A firm needs legal permission to operate and a job to be done that exists there. It needs a solution adapted to local preference, which is nearly always required (even Coca-Cola adjusts its formula by region, and a twelve-pack of paper towels is not a product for a flat of forty-five square meters in Hong Kong). It needs enough autonomy in local decision making to move at local speed. And, decisive, it needs alpha assets that confer advantage against local competitors who have everything else: scale, a global brand, or a

solution they cannot match. A newcomer without one of these has brought a knife to a gunfight.

The other directions follow the same logic. A firm can sell something more to the customers it already has, as Microsoft did when it attached Teams to Office. It can extend a brand into a category where the brand's promise still means something: Starbucks into packaged coffee and breakfast food, not into footwear. It can use its scale, as Interface, a maker of carpet tile for airports and offices, did when it launched a consumer brand, FLOR, that never mentions its parent and is made on the parent's lines at the parent's cost. A firm can also push a proven solution into a new job; Amazon's internal computing infrastructure became Amazon Web Services, the most valuable thing the company owns. Because a solution pushed into a job it was not designed for is often the wrong solution for that job, this direction takes care. Can the fifth flywheel point to a direction for growth? Sure. Toyota's production system is a culture of engaged people solving problems. Toyota now supplies that culture, through a dedicated unit, to organizations as far from automobiles as a children's hospital.

The direction to avoid is the one a firm's assets do not reach. Around 2000 the book retailer Borders, facing Amazon, had a respected brand, excellent locations, a strong balance sheet, and real capability in managing a temporary workforce and a high-variety inventory. None of those assets helped online. Borders handed its online store to Amazon to run, spent the decade in stores that

mattered a little less each year, and filed for bankruptcy in 2011. Barnes and Noble, with similar assets, spent heavily on the Nook to fight the Kindle, where its assets did not reach, and found its durable business running college bookstores, where all of them did. The test I give executives is the one a venture investor would apply: knowing the firm's assets and not its history, would you fund its plan to enter this market? If the plan depends on assets the firm does not have, the answer is no, however attractive the market.

For a Person

The same tests apply to a career, with two differences. A person cannot spread the work across divisions or replace themselves, so the limits of a finite life bind harder. A person's objective is also not return on capital (chapter 9 takes up what it is), so for now I assume you want to do work you are good at and valued for.

Finding the setting begins with an inventory of capabilities, but most people cannot see their own combination from the inside. My own effort is vivid to me and the alternatives are not. Herminia Ibarra's research on career change, built on interviews with people who had made one, found that people rarely discover a better fit by analysis. They discover it by trying: taking on a project outside the current role, working with a new group, joining something on the side, and noticing where the work comes easily and is welcomed. These experiments are cheap rela-

tive to the commitment they inform. If you want to know whether your combination is rare in some other room, go and stand in the room.

Two more considerations follow from the mechanics of chapter 4. The first is that early choices compound: the setting a person enters first shapes the opportunities they are offered next, through every channel of cumulative advantage. So the first choice of setting has consequences out of proportion to the information available when it is made. The second is that switching has a cost that rises with the investment already made. A specialist who has spent a decade building an asset that is valuable in one setting and ordinary in another faces a real loss in moving. If the setting is wrong, though, the loss is not a reason to stay. It is a reason to decide early, and to run the cheap experiments before the investment is large.

The three questions that I ask a founder to answer in an advantage thesis in chapter 7 apply equally to a person's next five years. What is my advantage now, and who pays for it? What will it be in five years, given what I am building and what the setting will do to it? What must I do in the next two years to get from the first answer to the second? The questions are harder for a person, who has to answer them from inside their current setting. A person who cannot answer the second question is building an asset for a setting they have not chosen.

Comparative fit also applies among collaborators, inside a team as well as across settings. A founder who is a good engineer and a poor seller is better off with a seller

than with a sales course. Together the two hold a combination that neither has alone. Most of the complements a person needs are other people.

The Person and the Model

What does the framework say about a frontier language model, a tool that in 2026 anyone can use? The model is performance enhancing and not hard to acquire, so it fails the second test of chapter 2, as electricity does. The economist Paul David made the historical point about electricity: having it conferred no advantage on a factory, and the gains went, decades later, to the factories redesigned around it and to the people who knew how. The returns to a universally available tool go to whoever holds the complements, and the question for a person is which complements those are.

The first of the four I would name is access to the customers, the problems, the decision rights and the data that make a model's output worth something. The second is accountability: someone must sign the work, and the signature carries a reputation that took years to build and that the model does not have. The third is judgment, which comes in two parts that are different assets. One is knowing what to ask, which requires enough domain knowledge to pose the problem; the other is knowing whether the answer is any good, which is taste. The fourth is the social capability to get an answer adopted by people who did not produce it. I suspect that when generation

becomes nearly free, the bottleneck moves to selection, and the person who can select well captures the surplus. The model is available to everyone, so it captures none. Picture a person with unusual taste and an unusual ability to carry a room, paired with a model that generates without tiring. Each part can be found separately; the combination is rare. By the chapter 2 definition, the asset is the pairing.

The evidence so far is narrower than that claim. In the studies that exist, the model compresses the differences between workers within a task. In a customer-support operation studied by the economist Erik Brynjolfsson and his colleagues, productivity rose about fourteen percent on average, with most of the gain going to the least experienced agents and almost none to the best. In an experiment with consultants, the below-average performers gained far more than the above-average ones on tasks inside the model's competence, while on tasks outside it, consultants using the model did worse. Both studies show the model leveling execution within a task. Does the model also widen the gap at the top, by rewarding the judgment that lets a person exploit it? That is the hypothesis this section rests on, and so far it remains a hypothesis. Chapter 3 told the story of the one study that seemed to prove it.

For young people the framework gives an uncomfortable answer. The entry-level worker's historical alpha asset was a willingness to do codifiable work cheaply while accumulating the judgment that would later be the

real asset. Now the model does the codifiable work. Early payroll evidence suggests that since 2022, in the occupations most exposed to the tools, employment of workers in their early twenties has fallen relative to older workers in the same occupations, while the older workers' employment has not. If that holds, the model is removing the apprenticeship that built the complement it makes valuable. Two things cut the other way. The first is that starting from nothing is cheaper: one person with taste and initiative can now build and ship what used to take a team. For people with the complements, the set of opportunities has widened, even as the entry-level jobs for people without them have thinned out. The second is that friction has fallen. Lower friction shortens the window in which a better product is an alpha asset and raises the value of the assets a model cannot build any faster: a customer base, a reputation, relationships. A young person who understands this will spend less of the window polishing the product and more of it acquiring customers.

What Travels

So the answer is the same for the firm and the person. Go where your assets are rare and necessary and where the complements you lack are supplied. Before you go, find out which of your assets travel, since the ones that do not belong to the setting and stay behind when you leave. The firm entering a new geography carries its brand and its scale and leaves its local relationships at home. The ana-

lyst changing banks carries their method and leaves the platform. The young engineer leaving a famous employer carries their judgment and leaves the name on the business card, and then discovers, as I did, how much of the inbound mail had been addressed to the name.

How Do You Build an Advantage Deliberately?

In 2019 MakerStock's first customers were the prototyping shops of university architecture, design and engineering programs. The business was simple: we bought plywood and acrylic by the pallet, cut it into sizes that fit the laser cutters those shops had installed, and shipped it. One nice thing about the segment was that we could assemble a list of every potential customer by reading university websites. Unfortunately, knowing who your customers are does not make them your customers. We acquired shops one or two a week, by email, by phone and sometimes by visiting. It was slow. Only about half of those shops became repeat customers. But two years in, we had a couple of hundred shops, enough for a steady weekly flow of orders, enough to buy material in economical quantities, and enough to justify better cutting machines. By then the customer base improved our results every week. A rival would have needed two years and a great deal of patience to match it. Performance enhancing and hard to acquire: by the two tests, we had an alpha asset that had not existed when we started.

My experience with MakerStock illustrates the basic physics of spinning up a flywheel from a standstill. The

engine is a value proposition offered to a customer with a job to be done. Because a flywheel accumulates energy incrementally, a small push every day for years produces a store of advantage that no single effort could. So the most fundamental variable is time. Firms build advantage by being there every day, acquiring and keeping customers. That chugging engine in turn spins the other wheels: volume, learning, quality and brand.

The hard part is the first push. When the wheel is stone cold, what value proposition can move it? In most cases, only a solution that is dramatically better than the alternatives for some customers, on some dimension they care about. Peter Thiel, in his book *Zero to One*, puts a number on it: a startup's technology should be ten times better than the nearest substitute. I take the ten as rhetoric rather than a measurement, but the underlying claim is right. When some disequilibrium in technology or in taste opens a gap, an organization that moves fast can offer a product far better than what existed. For a while that product is its only alpha asset. At the start of the pandemic, Zoom offered video calls that required no registration, no download and no special equipment; against a field of a dozen established products it took the market, almost entirely because its product worked when the others did not. Within two years Microsoft had Teams at parity on most features and ahead on some, and Google had rebuilt Meet. Zoom remains a large company, though its alpha assets now are its brand and its customer base; the product is table stakes.

A few firms skip the cold start. An established company can hitch a new business to an existing flywheel, as Microsoft hitched Teams to Office. A new firm with a famous founder starts with a brand, as the wine label LVE started with the singer John Legend. Chapter 2 used both cases as examples of endowments. For the rest of us the sequence is fixed. For a while the better product is the alpha asset; I call that period the window. The job during the window is to spin up the flywheels that will be the assets when it closes.

The Window

The window is short. I think it is the most important period in the life of a business, since what a firm builds during it determines whether the firm has anything left when the product advantage is gone. Founders naturally want to keep improving the product; that is usually the least useful thing they can do with the window. The firms that use it well do something that looks, from the inside, like premature optimization: they invest in brand, in customer relationships, in cost and in the organization's own capabilities at a time when the product is selling itself and those investments seem unnecessary.

Consider Keith Belling, who founded the snack brand Popchips and sold it. He then founded Right Rice, a rice-shaped pasta with a nutritional profile that many health-conscious eaters prefer to white rice. From Popchips he knew the playbook. Before the product was in a single

store, he hired the design firm IDEO to create the packaging and the identity, because brand would be the asset when the product was copied. He built distribution in the proven sequence (direct to consumer, then specialty retailers, then the mainstream chains), since each channel is a customer network that the next one does not replace. He could move many times faster than General Mills, but every move was toward a machine with flywheels that would matter once the business reached its sustaining phase.

Chapter 4 made the same point about a person: an early success is a window. Like a better product, a young professional's early success (a project that worked, a paper that was noticed, a client who was pleased) will be imitated or forgotten. What they have when the attention moves on is what they built while they had it: the relationships, a reputation for a specific kind of problem, a body of work that becomes a method.

The Advantage Thesis

The founders who use the window well can say, before they have anything to show, what their advantage is now, what it will be in five years, and what they will do in the next two to get from one to the other. I ask every student in my entrepreneurship course to write this down as what I call an advantage thesis. It has three parts.

The first is the advantage now: why anyone would pay today for what the venture offers, which for a new firm is

nearly always a solution that is better than the alternatives for some customers, plus the speed to reach them before anyone else does. The second is the advantage at year five, if the firm survives. By then the solution will have been copied, so the question is which of the five flywheels will have spun up far enough to be hard for a rival to acquire: a customer base, a brand, scale, a cost position, a capability. The third is the work of the next twelve to twenty-four months that converts the first into the second. The thesis is the logic that connects the three. I predict that a venture with a good product and no thesis will be copied at the moment it becomes worth copying.

The thesis is a prediction rather than a pitch. It is written before the result, the first of chapter 3's three habits. Camuffo's trials (later in this chapter) found that the habit improves decisions. Revenue in the first two years is noise, but the thesis can be checked on the assets rather than on the outcomes. Its measurements are the number of repeat customers, the cost per unit at the current volume, the share of new customers who arrive by referral, and whether the capability the firm claims to be building shows up in any observable practice. All of these move months before the financial results do. A founder who tracks the assets learns whether the thesis is working while there is still time to change it; a founder who tracks revenue learns at the end.

Here is MakerStock's thesis as I would have written it in 2019. Advantage now: nobody else sells plywood and acrylic cut to the sizes that fit a university shop's laser cut-

ters, and we can ship this week. Advantage at year five: a customer base of a few hundred shops, acquired one at a time, that a rival would need years to match; purchasing scale that lowers our cost per sheet; and a name that makers recognize. Work for the next two years: acquire shops one at a time, build the catalog around what the shops actually cut, and keep every repeat customer. The thesis was mostly right about the customer base and the scale. It was silent about the capability, the part that turned out to be hardest to build and most worth having.

The fifth flywheel belongs in the thesis too. The version of this exercise I assign asks founders to apply two tests to it. First, the capability must be credibly rooted in the founders; a culture of customer obsession declared by founders who dislike customers will not survive the first bad quarter. Second, the capability must reinforce one of the mechanistic flywheels, since, as chapter 5 said, a culture that does not feed cost, quality, brand or the customer base is a club. A thesis that passes both tests names one capability, says which flywheel it feeds, and says what the founders will do in the first year, visibly and at some cost to themselves, to build it.

The Levers

Four of the five flywheels respond to ordinary management: sell more, learn faster, fix the product, feed the brand. The fifth, organizational capability, needs partic-

ular levers. The firms that have built a durable capability have pulled the same ones.

The example I use is DuPont and safety. My father, who worked as a consultant to DuPont, came home one day with a story: someone at the laboratory had cut a finger on a box cutter, and the unit had stopped to analyze the root cause and plan against a recurrence. For more than a century DuPont has listed safety first among its values. It has even turned the capability into a product, a consulting business that sells safety practice to other companies. The man who ran that business put the connection to the firm's performance in one sentence: good safety means control of operations; control of operations is what drives improvement in everything else.

The first lever is visible, costly action aligned with the capability. Employees are not stupid. Because it costs something, stopping a laboratory to investigate a cut finger persuades everybody; a slogan at the plant entrance persuades nobody. When I became Wharton's vice dean for innovation, the dean set aside about half a percent of the school's revenue for exploration with no specified return, appointed a senior faculty member to lead it, and announced the budget. The money was the signal; without it the title would have been decoration.

The second lever is an origin story. DuPont employees will tell you that the company began in 1802 with one product, gunpowder; that explosions were frequent; and that the founder built his own house within the blast radius of the mill. Amazon employees will tell you that the first

desks were doors on sawhorses. Whether every detail is true is less important than whether the stories are told and retold, each one reinforcing a single capability: safety at DuPont, frugality at Amazon.

The third lever is codification: DuPont writes safety into performance reviews and has structured processes for risk assessment that it has refined for decades and now sells. Processes and systems aligned with a capability reduce the friction on the wheel, while processes at cross purposes with it stop the wheel entirely, whatever the slogans say.

The fourth lever is a manifesto, such as the slide deck on its culture that Netflix published in 2009 and that has since been viewed tens of millions of times. Two of its lines say exactly what capability the company wanted and what it would pay to keep it: "A great workplace is stunning colleagues," and "Adequate performance gets a generous severance package." (Whoa. Got it.) A company that says something that specific, in public, has committed itself in a way that a mission statement about excellence never does.

The fifth lever is training. DuPont implemented its safety system with multi-day programs for plant managers and supervisors, coaching for every leader, a half-day for every worker on the floor, and recurring workshops after that. The Indian textile plants in the Bloom experiment that chapter 5 described were trained for months and improved, but over nine years without rein-

forcement they let some of the practices slide. A capability lasts as long as people keep teaching it.

Learning When the Score Is Noisy

The levers above build capabilities whose effect can be seen. The harder problem is building an advantage when the outcomes that would tell you whether you are improving are too noisy to read, which chapter 3 argued is the normal condition in most consequential domains. Three findings from research on expertise and on young firms suggest what to do.

The first finding is about the conditions under which expertise develops. The psychologists Daniel Kahneman and Gary Klein, who disagreed about most things, agreed on this: intuition becomes reliable only in an environment regular enough to have patterns, and only when the learner gets prompt, unambiguous feedback. Chess, fire-fighting and anesthesiology qualify; picking stocks, forecasting politics and hiring executives do not. In those domains confidence grows with experience while accuracy stays put. So build the feedback where the environment does not supply it: write down the prediction, with its reasoning, before the result, and keep a record of decisions and their expected outcomes, and read it. A decision journal turns an environment with no feedback into one with some.

The second finding is about the kind of practice that builds expertise. The psychologist Anders Ericsson's

research on expert performance showed that the kind of hours, not the number, is what produces expertise: deliberate practice, at the edge of current ability, with immediate feedback and repeated correction. The popular version, ten thousand hours of anything, is not what he found. Later work found that the share of performance that practice explains varies a great deal by domain and is far smaller in the noisy ones. Still, the structure of deliberate practice transfers. A salesperson who records every call and reviews it with a colleague is practicing deliberately; one who makes a thousand calls and counts the orders is accumulating experience.

The third finding comes from experiments with young firms. Arnaldo Camuffo and his colleagues ran four randomized trials in which several hundred young firms were trained in what they called a scientific approach: state the theory behind the business, derive what must be true for it to work, test it, and revise. The treated firms were not more successful by every measure. They were more likely to shut down projects that were not working and more focused in how they changed direction, and on average they brought in more revenue over the period. So better process led to more abandonment, not more persistence. In a noisy environment, the deliberate builder of an advantage specifies in advance what would make them stop, and then they stop.

An AI model can supply feedback where none existed, in domains where its output can be checked: code that runs or does not, a draft against a standard, an analysis

against a known answer. In those domains it is a practice partner of a kind no previous generation had. Where the model produces the work instead of checking the person's, it builds no capability in the person. I suspect a career spent asking the model for answers is a career spent practicing nothing. Ericsson drew the same line between deliberate practice and experience.

One more distinction heads off a common mistake. The population findings that chapter 3 recommended for evaluation (a process feature that predicts outcomes across hundreds of firms or managers) are evidence for prediction but weaker evidence for prescription. For instance, active share, the fraction of a fund's holdings that differs from its benchmark index, predicts fund returns across the population of funds. It does not follow that a closet indexer (a manager who charges for active management while holding something close to the index) will earn more by raising their active share, because the funds that chose high active share may differ in other ways. The textile experiment is the exception: because the practices were assigned at random, we know that adopting them helps. For most process features no such experiment exists. A manager who adopts a population-validated practice is making a reasonable bet. They should call it that and watch for the moment the practice becomes a target and stops being a measure.

Bets That Did Not Pay

An account of deliberate building that reports only the wins is the halo that chapter 3 described, the result written back into its explanation. I have more failures than I have room for.

At Xootr, our scooter company, we spent two years and a meaningful fraction of our capital developing a product line adjacent to the scooter, on the theory that our brand and our manufacturing capability would carry it. The brand did not reach that far. Once the volumes were different, the manufacturing capability was not the asset we thought it was. We shut the line down. At MakerStock we assumed for a year that the shops we had won in universities would lead us to the professional fabrication shops that use the same machines, and that the customer network would extend itself. It did not; the professional shops bought on price and lead time from suppliers we could not match, and the university relationships gave us no entry. In both cases we treated an asset that was performance enhancing and hard to acquire in one setting as if it would be the same in the next. And in both cases we had not written down, before we started, what would make us stop, so we stopped late.

When Does Advantage Stop Working?

In April 1860 a company began carrying mail from St. Joseph, Missouri, to Sacramento, California, by relay: a rider covered roughly ten miles at a gallop, swapped horses at one of nearly two hundred stations, and after several stages handed the mail to the next rider. A letter crossed the continent in about ten days. When the transcontinental telegraph was completed in October 1861, the Pony Express, after eighteen months of operation, closed within days. The job it did, moving information across a continent, is done today more than ever, but the solution was obsolete the week a better one appeared. The horses, the stations, the riders and the brand were worth nothing against the telegraph, which needed none of them.

A flywheel keeps a machine running when its power is interrupted, and the flywheels of chapter 5 do the same for a firm. From roughly 2009 to 2015 Volkswagen sold diesel cars whose software switched on the emissions controls only during a test. Executives knew. The company paid tens of billions of dollars in fines and settlements, and its market value fell by a third in a week. Today it is still among the two or three largest car makers in the world.

Incumbents with spun-up flywheels are nearly impossible to kill. What, short of a telegraph, does kill them, or, more often, kills the advantage and leaves the firm alive?

Imitation, and the Friction That Slows It

Most advantage ends the ordinary way: a rival copies it. In the 1989 Tour de France, Greg LeMond's aerodynamic helmet and clip-on bars were alpha for one afternoon; by 1995 they were standard equipment. The mobile ordering app of the Chinese coffee chain Luckin Coffee was alpha for a few years. Patents run out, trade secrets leave with the engineers who know them, and the proprietary process becomes the subject of a conference talk. An alpha asset lives only as long as it stays hard to acquire.

What sets the length of that life? Friction, of which there is more than economists' models assume and less than incumbents hope. Customers switch slowly. Electronic toll collection was demonstrably better than a coin basket for a decade before most drivers had a transponder, and LED lighting was cheaper over its life than the incandescent bulb for years before it was in most sockets. Everett Rogers spent a career cataloguing why good ideas diffuse slowly, from the British navy's two-century delay in adopting citrus against scurvy onward; his answers in *Diffusion of Innovations* (relative advantage, compatibility, complexity, trialability, observability) amount to a list of frictions. Know-how moves slowly even between units of the same firm, as Gabriel Szulanski found (chapter 5), and more

slowly between rivals. So an advantage with a theoretical life of a season often has a practical life of years, and a firm that understands its own frictions can estimate how long it has.

Friction is also what a language model reduces. A product that took a year to reverse-engineer can be approximated in a week. A process that lived in the heads of a few engineers can be described, by them, to a tool that documents it. Advantages that rest on a better solution now have shorter lives than they used to, and the assets that never depended on friction for protection (relationships, installed bases, brands) gain in relative value.

Two Conditions for Disruption

Some advantages do not erode; they end outright, either because the solution to the job changes, as it did with the telegraph, or because the job itself changes, as it did with the decline of church bells. Executives fear the first kind most, so let's start there.

A new solution technology destroys an incumbent only when two conditions both hold. First, the new solution must actually prevail: a substantial fraction of the market must come to prefer it. Second, the incumbent must be unable to adopt it profitably, perhaps because the value is redistributed to the consumer and nothing is left to capture, or because someone else controls the new solution, or, most often, because the incumbent's alpha assets are irrelevant to it. Consider Kodak. Digital imaging pre-

ailed, and nothing Kodak owned (its chemistry, its film plants, its relationships with photo finishers) had any use in it. The job survived; the firm's assets did not. The entire value of consumer photography did not move to a successor firm but to the consumer, who now takes pictures for nothing.

Most threats fail one condition or the other. Although electric cars have taken share for fifteen years, the incumbents that make combustion engines also make electric cars, since the assets that matter in the sustaining phase of the automobile business (brands, dealer networks, scale, manufacturing capability) are mostly relevant to the new solution. Electric cars fail the second condition; e-books fail the first. They exist and are convenient, and their sales plateaued at a fraction of the market because a substantial fraction of book buyers prefer paper.

The threat that passes both conditions is rare. I once tallied, informally and from a bankruptcy database, the public companies that failed in the decade from 2005 to 2015 because a new technology made their solution obsolete. I found about sixteen, out of several hundred bankruptcies among several thousand firms: Blockbuster, Borders, Kodak, Circuit City, a handful of newspaper and printing companies, Nortel. That is a base rate for bankruptcy from technological disruption on the order of a few hundredths of a percent per firm per year.

That base rate comes with two cautions. First, bankruptcy is a narrow measure, since margins, share and independence usually go before the firm does, and a company

can survive while its advantage disappears entirely, as the newspaper companies that did not go bankrupt have. Second, a tally of failures is not a theory. Clayton Christensen's account, in *The Innovator's Dilemma*, is that a cheaper, worse product improves until it displaces the incumbent from below. In a 2015 article, Andrew King and Baljir Baatartogtokh went through seventy-seven cases that the disruption literature uses and found that most did not fit the theory's own conditions. Disruption by technology is real, rare, and poorly predicted by any theory I know, including my own two conditions.

What the Incumbent Can Do

Executives are told that only the paranoid survive and that a firm must disrupt itself before someone else does. I disagree. For most incumbents, most of the time, the right response to a technological threat is to track it, deliberately do nothing about it, and keep spinning the flywheels every day. Most threats fail one of the two conditions, and chasing each one costs more than the rare one that materializes would have.

When a threat does look real, I apply the two tests of chapter 2 by asking whether our alpha assets will remain alpha against the new solution, and if not, whether we can acquire the newly relevant ones by building or buying. A third question is easier to forget: for what other jobs might our assets, no longer alpha for this one, still be alpha?

The menu of responses is short. A firm can commit fully to the new solution, which works when the solution does prevail and the firm has the capabilities to win in it; Bill Gates's 1995 memo turning Microsoft toward the internet is the example everyone cites. It can hedge with a project or an autonomous unit, as Netflix did with streaming while the DVD business still paid the bills, although the legacy business usually starves the hedge that threatens it. It can fight, by acquiring and killing, by suing, by lobbying. (Paris taxi drivers fought ride-hailing for years and bought themselves years.) It can concede and manage the old business for cash. That is the right call more often than managers admit, because the new solution often does not dominate, or takes decades, or leaves a profitable remainder. Or it can push its existing assets into adjacent jobs where they remain alpha, the subject of chapter 6 and the response most often neglected in favor of the more dramatic ones.

The test I use is the one a venture investor would apply: knowing the incumbent's assets, would you fund its plan to launch the rival solution? In 2000 the book retailer Borders could not have raised money to beat Amazon at internet retail from any investor who looked at its assets, and it should not have tried. Its shareholders could buy Amazon themselves. Management's job was to maximize the cash flow from the assets the firm had. That pointed toward the college bookstore business its rival Barnes and Noble

eventually built, not toward the fight with Amazon that Borders lost.

In November 2021 I put a question to executive audiences. Toyota and Tesla had revenues and market values that invited comparison: Toyota's market value was about ten times its earnings, Tesla's about seven hundred times, on revenue a fifth of Toyota's. How would you bet? The question was a test of the two conditions. Electric drive was clearly going to take a large share. Toyota's alpha assets (its brand, its dealer network, its manufacturing capability and its scale) were mostly relevant to electric cars, while Tesla's were a head start in software and batteries and a brand built on being first. For Tesla's valuation to be right, Toyota's assets had to stop being alpha against the new solution; for Toyota's to be right, they had to remain alpha.

Anyone can look up the two companies today, and the figures will have moved again since I asked. A firm whose assets remain alpha against a new solution is not disrupted by it, however loud the solution is. A firm whose assets are irrelevant to the new solution is disrupted by it, however large the firm. Which of the two Toyota turned out to be is a fact about its assets, knowable in outline in 2021 from the worksheet in the appendix, without a forecast of either company's share price.

When Your Own Advantage Stops

The same two conditions apply to a person, where they are easier to see because the firm around them supplies the complements.

Start with a skill at executing a particular kind of work, which lasts as long as friction keeps others from acquiring it. Friction has fallen. The skills most exposed are the ones whose output a model can verify and therefore produce: a clean translation, a working function, a standard analysis. The least exposed are the skills around the execution: knowing which problem to solve, judging whether the solution is right, getting it adopted. If your asset was the execution, its life is getting shorter; if it was the judgment around the execution, its life is getting longer, since the execution you judge is cheaper than it was.

Assets that come with a position stop differently: they belong to a setting and end when the setting does. The analyst's rank was the bank's platform; the professor's inbound mail was the institution's name; the salesperson's number was the territory. Chapter 4 asked whether an advantage belongs to the person or to the position, which is this chapter's question asked in advance. If you know which of your assets are yours and which are the setting's, you can see the end of a position coming and build, before it comes, the assets that travel. If you don't, you learn the difference on the day you leave.

Accumulated advantage, for a person as for a firm, also makes changing direction costly. Ten years invested in an

asset that is alpha in one setting are ten years of real reasons to stay, even when the setting is dying. The deliberate builder of chapter 7 wrote down, before starting, what would make them stop. The holder of an advantage should do the same: name, while the asset is still alpha, the conditions under which it will not be, and watch for them.

What Should Success Mean?

The investor Warren Buffett has for decades put a thought experiment to audiences of successful people. Imagine that twenty-four hours before you are born, a genie offers you the chance to design the economic and social rules of the world you will enter, on one condition: you will then draw a ticket from a barrel containing every person who will be born that day, and you will be that person. You might be born in the United States or in a country with a tenth of its income, healthy or ill, clever or slow, to parents who read to you or to none at all. Buffett calls this the ovarian lottery. His point is that people inclined to credit their position to their own merit have usually drawn a very good ticket without noticing.

The device is older than Buffett's use of it. In 1971 the philosopher John Rawls asked what principles of justice people would choose for a society if they had to choose from behind a veil of ignorance, not knowing which position in that society they would occupy, which talents they would have, or which generation they would belong to. Rawls called the distribution of native talents and family circumstances the natural lottery. He held that the outcome of that lottery is arbitrary from a moral point of

view (so no one deserves their starting place) and that a just society is one the losers of the lottery could also have agreed to. I have been circling the first of those claims for eight chapters. Buffett's genie is Rawls's original position (his name for the choice made behind the veil) retold for an audience of winners. I do not know whether Buffett took it from Rawls. The resemblance is exact; both men conclude that the winner's obligations are larger than the winner's instincts suggest.

So far I have written about what advantage is, how it accumulates, where to apply it, how to build it, and when it stops. I have argued that advantage is real, that most of what makes it decisive was built or inherited rather than willed into being in the moment, and that a record of results tells us less than we think about who built what. What does any of that entitle a successful person or firm to claim, and what does it ask of them? I take the question in four parts: what advantage is for, what luck means, what to do about the fact that none of us chose our starting point, and what we owe each other.

What Advantage Is For

A firm's objective is, to a close approximation, return on the capital invested in it; every chapter so far has treated advantage as whatever raises that return. A person's objective is not given. I have put off the question until now because the answer changes everything before it.

I am a capable administrator. I ran a school's innovation activities and a group of thirty people for years, and by the ordinary measures I did it well. I also dislike meetings; administration is mostly meetings. By the chapter 2 tests, administrative capability was an alpha asset for me, since it was rare enough in a faculty and necessary to the job; it also brought the position and the invitations that chapter 4 described. It was an advantage in a contest I did not want to win. Because a capability that would prevail in a setting a person does not want to be in is not an advantage in any sense that should guide them, the inventory in chapter 6 is incomplete until it includes what the person wants the work to feel like.

Psychologists who study motivation have converged on three things that make work satisfying across cultures and occupations: autonomy, competence and relatedness (in plain terms, a person chooses the work, is good at it, and does it among people they are connected to). Competitive success delivers the second and often undermines the first and third. Someone optimizing for distinction takes the role in which they are best and most visible, which is frequently the role with the least autonomy and the most meetings.

The market, for its part, rewards distinction out of proportion to the ability behind it. The economist Sherwin Rosen showed in 1981 that where a performance can be reproduced at low cost and buyers prefer the best, small differences in ability become enormous differences in earnings. The second-best opera singer in the world earns

a fraction of what the best one earns. Chapter 3 showed that records of results, income among them, are noisy measures of ability. I think income is also a noisy measure of a worthwhile life.

So the objective comes first. Before asking whether an asset is an advantage, say what the advantage is for: a particular kind of work, a level of autonomy, a set of people, a problem worth solving, and, yes, money, in whatever proportion each person sets for themselves. The worksheet in the appendix begins with that question.

Two Meanings of Luck

I have used the word luck in two senses. When the two are run together, a statistical argument turns into a moral one without anyone noticing.

The first sense is outcome luck: the variation in results, for a given level of ability, that chapter 3 called noise. Think of a venture partner's returns, a founder's exit, or a chief executive's three years of results. Outcome luck is a property of the domain and can be estimated from populations; about the person, it says only that their record should be shrunk toward the average.

The second sense is constitutive luck: the fact that a person's ability, and the circumstances that let them use it, were not chosen by them. Consider the cyclist Greg LeMond's lungs, the swimmer Michael Phelps's wingspan, the year a graduate first looks for a job, or the parents who read to them. Constitutive luck is not noise. A capability

can be perfectly real, perfectly measurable and decisive in every contest, and still be unearned, in that its owner did nothing to acquire it. A free-throw shooter who makes 90 percent of their shots is skilled by any statistical test, though they did not choose their hands.

The two senses call for different responses. Outcome luck is a reason for humility about what a record proves; the response to it is the population method of chapter 3. Constitutive luck is a reason for humility about what a real advantage entitles its holder to claim.

Determinism Without Fatalism

Constitutive luck is usually discussed as a matter of degree: how much of a person's success is their own, and how much is their genes, their family and their era. I hold a stronger view, although nothing practical in this chapter depends on it.

I think a person is a physical system, and a physical system does what its state and its inputs determine. Whatever I decide, I decide because of the state of my brain at that moment, a state produced by everything that happened to my brain before, none of which I chose at the start. I have the feeling that there is a separate self behind my eyes, weighing the options and freely choosing. In my view that feeling is an illusion, installed by evolution because it makes a creature deliberate, feel accountable and learn from regret, all of which improved its ancestors' prospects. Experiments in which brain activity predicting

a choice is detectable before the subject reports deciding are suggestive of this, though their interpretation is disputed. The majority position among philosophers belongs to the compatibilists, who accept determinism and still call the deliberating self free; I have never been able to see their view as anything but a redefinition of the word. Serious people disagree on this question; the scientific findings constrain the answer without settling it. Someone who rejects my view can still accept everything that follows.

Nothing that follows depends on my view because of a distinction that people routinely miss. Determinism is the claim that what happens is caused by what came before. Fatalism is the claim that what a person does makes no difference to what happens. Fatalism does not follow from determinism; it is also false. Imagine a train driver who sees a child on the track ahead. If they pull the brake, the child lives. Whether their pulling the brake was freely willed or fully determined by their history has no bearing on the child. Now suppose two men enter the cab and force the driver's hand onto the lever. The driver plainly lacks free will; the child lives all the same. Actions are links in the chain of cause and effect whether or not anyone wills them, just as each domino in a row falling in a fully determined sequence still knocks down the next.

A few studies once suggested that people primed to doubt free will behaved worse. A meta-analysis of a hundred and forty-five experiments found that the manipulations changed what people said they believed, and found

no reliable evidence of downstream effects on behavior. That result fits the distinction, since people who stop believing in free will do not, in general, become fatalists.

So the practical question for a person who accepts that they did not choose their starting point is unchanged. They still have a brain that evaluates options and acts on the one it expects to value most, the evaluation still improves with information, practice and feedback, and their actions still have effects. Every prescription in chapters 5 through 8 survives; what changes is how they interpret the result.

Credit, Incentives and Desert

Once constitutive luck is admitted, three questions that usually travel together come apart: whether a person caused a good result, whether they should be rewarded for it, and whether they deserve the reward in some deeper sense.

Causal contribution is a fact, often an unknowable one (as chapter 3 said), but a fact. The salesperson did make the calls; the founder did persist when the money was short. Incentives are a separate matter. We justify them by what they will produce rather than by what happened: we pay for results because paying for results produces more of them, whether or not the results were earned in the deep sense. A society that stopped rewarding the free-throw shooter because they did not choose their hands would have fewer good free-throw shooters.

Constitutive luck undermines the third claim, desert: the claim that the successful person is entitled to their reward as a matter of justice. Most people have met the word only in the phrase "just deserts" and have assumed it was misspelled. Desert is the noun from deserve (what a person deserves) and is pronounced exactly like the course after dinner. That is no accident: both come from the French for serving. To deserve was originally to serve well enough to earn something; the dessert was what the servants brought once they had cleared the table. (The desert of the Sahara is a different word that happens to be spelled the same way; it comes from the Latin for abandoned.) Moral desert, then, is what your service has earned you. The practical issue is how much of what you have was service and how much was the ticket. If the capability was not chosen, the circumstances were not chosen, and the outcome was partly noise, then the reward is a fact about the system's incentives rather than a verdict on the person's worth.

The strongest objection to that conclusion came from Robert Nozick, another philosopher, three years after Rawls. Grant that nobody deserves their talents; it does not follow, Nozick argued, that nobody is entitled to what flows from them. A person may hold their native assets undeservedly and still have a legitimate claim to what they make with them, in the way that a person who did not earn their eyesight is still entitled to what they see. I find the objection serious but not decisive. Because it shows that the chain from unchosen talent to undeserved

reward is not automatic, I keep causal contribution and entitlement under the rules separate from desert in the deep sense. A person is entitled to what the rules of their society assign to their effort; the rules in turn are justified by what they produce. What the natural lottery removes is the further claim that the rules merely recognize what the person was always owed.

I see three consequences of this view that I find useful rather than merely consoling. The first concerns judging others. A board that fires a chief executive after three years of poor relative returns, which boards routinely do, is acting on a record that chapter 3 says cannot support the judgment. The board is also acting as if the executive deserved the firing, which is a different claim again. The alternative is not to stop making decisions but to decide on the evidence that exists (the population evidence about process) and to drop the language of desert from the decision. Fire the executive if the process predicts bad results; do not pretend the three years proved anything about them. The same holds for the limited partner allocating on two fund records, and for the voter who punishes the government for a recession it did not cause, which voters reliably do.

The second concerns judging oneself. Regret is useful because it is a form of learning: recognizing that a decision produced a bad result, and resolving to decide differently, is how a brain that deliberates gets better. Self-reproach is not useful. Given everything that led to it, a past action could not have been different. Energy spent

on blame is energy not spent on the next decision. For an individual, the practical content of determinism is the difference between learning from a mistake and hating oneself for it. I do not think a person needs any more content than that.

The third concerns compassion. If I had drawn a different ticket, I would be a different person with a different record. The person sleeping on the street outside my office is who I would be with a few changes in DNA and neighborhood. None of this means letting dangerous people do as they please. A computer infected with a virus is disconnected from the network, not punished; a society is entitled to protect itself from people whose decision processes produce harm. The argument is for doing so without the satisfaction of moral superiority, and for extending to the unlucky the same recognition the lucky extend to themselves.

Advantage and Everyone Else

So far I have treated advantage as a good thing for its holder. Whether it is good for anyone else depends on its source, which is usually easy to see.

An advantage from scale or learning lowers cost, and competition passes much of the saving to customers; an aluminum smelter that produces at the lowest cost in the world makes teapots and electric cars cheaper for everyone. An advantage from a better product or a proprietary process delivers something customers prefer. The patent

that protects it for a period is what society pays to get the next one. An advantage from a brand is a promise about quality that customers pay for willingly. These advantages create value and share it. Most of the alpha assets in this book are of this kind. Other advantages transfer value rather than create it: a permit obtained by bribery, a regulation written by the incumbent, a network effect maintained by blocking interoperability, a scarce drug priced at what desperate patients will pay. The two tests of chapter 2 do not tell these cases apart, because a bribed permit is performance enhancing and hard to acquire. The difference is whether the customer is left better off. I think a firm that wants to be allowed to keep its alpha assets should be able to say which kind it holds.

A society faces a version of this question when it decides what to do about the natural lottery. In chapter 1 I stated two beliefs that nothing since has given me reason to retract. The first is that every person and every organization should have the chance to prove themselves and thereby earn access to further chances; that is equality of opportunity, and no society fully provides it. The second is that because so much of the inequality that determines well-being is unchosen, people acting together are obliged to reduce it where fundamental needs are at stake. Constitutive luck is the reason for both beliefs. Rawls gave the second a precise form that is consistent with paying for results. His difference principle permits inequalities, including large ones, when they work to the benefit of the least advantaged; that is, when the incentive that pro-

duces the inequality also produces something the losers of the lottery share in. An advantage from scale that lowers everyone's prices passes; an advantage from a permit obtained by bribery does not. So a society can apply the same test I proposed for a firm's alpha asset, which is whether anyone other than its holder is better off for it. The lucky are not entitled to their luck. They are responsible for what they do with it.

In chapter 1 my mother, a historian, asked what evidence I had that any of this was true. The book has made four kinds of claim, each resting on a different kind of evidence. Some were definitions: advantage, alpha asset, the two tests, the three questions about an early success. Some were mechanisms argued from logic: the flywheels, the four channels of cumulative advantage, the two conditions for disruption. Some were findings from populations: the persistence of returns, the leveling effect of the tools, the Matthew effect in grants, the management experiment in India, the arithmetic of what a record can reveal. And some were stories, including my own, which illustrate the mechanisms and prove nothing. They deserve different weights; the claims I hold with the most confidence are the ones about what cannot be known.

What remains is a short set of questions for the next decision, whether a firm's or a person's. What is the advantage for? What asset would improve expected results in this setting, and would the rivals in this setting find it hard to acquire? Who holds the complements, and who will

capture the return? What does the record actually show, given the noise and the number of outcomes, and what does the population say that the record cannot? Which of the assets travel, and which belong to the position? What would make this stop working, and what will I watch for? And when it works, what part of this was mine, and what do I owe for the rest?

A rider born with exceptional lungs did not choose them; they still have to decide how to train, which races to enter, and whom to ask for help. The rest of us face the same decisions with less convenient finish lines. Like that rider, we start with what we were given and decide what to build from it.

The Alpha Asset Worksheet

I use this worksheet with executives, founders and students to assess a firm, a business unit, or a person. It is a set of prompts for an argument; the numbers it asks for are positions to be defended and disputed, not measurements. Someone who fills it in already knowing how the story turned out will find alpha assets everywhere. So the habits from chapter 3 apply, and I have ordered the questions to enforce them.

Before the Assets

State the objective. For a firm, the usual answer is return on invested capital over a stated horizon; for a person, say what the advantage is for (the kind of work, the autonomy, the people, the problem, and the money, in whatever proportion you choose). If an asset would win a contest you do not want to enter, leave it off the list.

State the setting and the rivals: which job, for which customers, against whom? Because the same asset can pass the tests in one setting and fail them in another, name the setting before you score anything.

The Assets

For each candidate resource, record four things. First, what it is, in observable terms: a practice, a relationship, a stock of customers, a patent, a skill, a name (an adjective is not an asset). Second, the mechanism: how it raises willingness to pay or lowers cost, stated so that someone could check it without knowing the result. Third, an observation independent of the result that supports the mechanism, such as the sister division that could not copy the practice, the rival that tried to build the customer base and failed, or the price premium measured before the firm's success was known. Fourth, an alternative explanation for the firm's results that does not involve the asset.

Then score the asset on two scales, each from one to ten.

Rarity: how hard would it be for a rival in this setting to acquire the resource? One means a commodity available to all in an open market; ten means essentially immobile, obtainable only over many years or at a cost comparable to the value of the enterprise. Score by what the rival would have to spend to get it, not by your own history with it.

Necessity: how much does the resource improve results in this setting, and how few substitutes exist? One means many substitutes and a small effect; ten means a decisive effect with no substitute.

The product of the two is the alpha score. A resource that scores above about forty is a candidate alpha asset,

and one below about fifteen is table stakes or decoration. These thresholds are my own rules of thumb.

The Categories

Most alpha assets fall into one of the following categories, which I offer as a menu of places to look rather than as a theory.

Category	Examples
Solution	Product designs; patents, trade secrets and know-how; proprietary standards; regulatory approvals; a causal model of why the solution works
Process	Process designs; accumulated efficiency; metrics and information systems; a causal model of process performance
Brand	Trademarks, domains, app names; brand associations; awareness
Scale	Relative volume or share in the category
Customer network	Customer relationships and data; the size of the user network; network effects where they exist
Organizational processes and relationships	Talent acquisition and development; demand-generation know-how; supplier relationships; distribution relationships; relationships with regulators
Capability and culture	A named set of practices that the organization performs better than rivals and that reinforces one of the above
Other	Location, natural resources and access to inputs; balance sheet and access to capital; idiosyncratic endowments such as a founder

For a person, the categories translate: the solution is a skill or a body of work; the process is a method; the brand is a reputation; scale is tenure and the volume of work done; the customer network is relationships. The organizational category becomes the platform the person stands on, which may belong to the person or to the employer.

After the Assets

People most often skip the three questions that complete the worksheet.

What would it cost to acquire or build this asset, and to maintain it? If you buy an asset at a price that reflects its value, you earn no return on the purchase. Some assets must also be fed, and the cost of feeding a brand, a culture or a learning curve belongs in the score.

Who holds the complements, and who captures the return? If the asset is valuable only in combination with something another party controls, the other party has a claim on the return, and so does a person who embodies the asset.

What would make this asset stop being alpha? Name the disequilibrium, the imitator, or the change in the job that would reduce its rarity or its necessity, and then name the signal you will watch for.

A Worked Example

For a management consulting firm of the first rank, with one person's opening scores, the worksheet might read as follows. Frameworks and tools: rarity three (they are published and leave with every departing consultant), necessity five, score fifteen: table stakes. Project workflow: rarity two, necessity five, score ten. The brand: rarity ten, necessity five, score fifty; the mechanism is that a buyer of advice cannot inspect its quality and pays for

a name that has more to lose than the buyer does; the independent observation is the premium the firm commands on identical engagements. Scale, some fifty thousand professionals: rarity seven, necessity three, score twenty-one. Share of the largest companies as clients: rarity seven, necessity three, score twenty-one, with little network effect. Recruiting, development and mentoring: rarity seven, necessity seven, score forty-nine. Client relationships, actively cultivated through an alumni network and senior access: rarity eight, necessity eight, score sixty-four. Culture of work ethic and quality: rarity seven, necessity seven, score forty-nine. On this reading, the firm's alpha assets are its client relationships, its recruiting and development system, its culture and its brand; its frameworks are not among them. End conditions: a language model that reproduces the frameworks for nothing lowers their necessity (they were already table stakes) and raises the relative value of the client relationships, which the model cannot reproduce.

The Thesis

The worksheet so far describes assets a firm or a person already holds. A new venture, or a person at the start of something, holds almost none and needs the prospective version that chapter 7 calls an advantage thesis. Chapter 7 gives it three parts; here it has five questions, the last two about evidence.

What is the advantage now? Why would anyone pay today for what you offer, and how will you reach them before a rival does?

What will the advantage be in five years, if you survive? Name the flywheels that will have spun up by then and would be hard for a rival to acquire, along with the one capability the organization will be known for and the mechanistic flywheel it feeds.

What must happen in the next twelve to twenty-four months to get from the first answer to the second? Name the work, not the intention.

What observable evidence, within a year, would show that the assets are forming? Count repeat customers, cost per unit at volume, referral share, or a practice visible in the organization, but not revenue.

What would show that the thesis is wrong? Write it down before the result, and decide now what you will do if it appears.

Two More Examples

For a failed company, fill in the worksheet with only what was known before the failure. For the book retailer Borders around 2000, it might read as follows. Brand: rarity seven, necessity four, score twenty-eight. Retail locations: rarity eight, necessity three, score twenty-four. Capability in managing a temporary workforce: rarity five, necessity four, score twenty. High-variety inventory management: rarity six, necessity five, score thirty. Balance sheet: rar-

ity five, necessity three, score fifteen. None scores above forty, and the scores would have been the same in 2000 as in 2011. Every one of these assets had the same end condition: as book buying moved online, the necessity of each fell toward one. Someone filling in the worksheet in 2000 would not have predicted bankruptcy. What they would have seen is that the firm's assets were worth most in jobs that did not involve competing with Amazon, the same jobs into which the surviving rival took its own similar assets.

Now consider a person: a (hypothetical) mid-career engineer at a large manufacturer who wants more autonomy and less travel. Their solution-level asset, deep knowledge of one product family's failure modes, scores rarity eight and necessity seven inside their firm, and rarity eight and necessity two at any other firm, since no other firm makes that product family. Their method, a way of running design reviews that their division has adopted, scores rarity five and necessity six wherever they work. Their reputation scores rarity six inside the firm and two outside it, where nobody has heard of them. Their relationships with the firm's suppliers score rarity seven and necessity six and would partly go with them, since the suppliers serve other customers. So their highest-scoring asset is the one they cannot take with them; what they would carry is their method and their supplier relationships. That combination is rare at a supplier, or at a rival that uses the same supply base, but not in a different industry. Their end condition is the product family's

retirement, which they can see on the firm's roadmap. They still have to choose what to do, but they now know what they have, where it counts, and when the choice is due.

Notes

Notes are keyed to phrases in the text and gathered by chapter. In this working proof, notes migrate here from each chapter's verification queue as claims are confirmed against primary sources.

Chapter 1. What Explains Success?

passed it over more than a hundred bicycles Union Cycliste Internationale, "The UCI announces Disciplinary Commission decision in the case of Femke Van den Driessche" (<https://www.uci.org/article/the-uci-announces-disciplinary-commission-decision-in-the-case-of-femke-van-den-driessche-176133/6GpCo3zqNjX8fmHTBYEwgr>), April 26, 2016. The release states that the bike was scanned with the UCI's new magnetic resonance testing in the rider's pit area during the women's under-23 race, that the motor and battery were concealed in the seat tube with a Bluetooth switch under the handlebar tape, and that over 100 bikes were scanned at the event. The page as currently served omits the decision text; the sanction details are as reported in "Breaking News: Van den Driessche Receives Six-Year Suspension" (<https://www.cxmagazine.com/breaking-news-van-den-driessche-uci-disciplinary-committee-motorized-doping-technological-fraud>), Cyclocross Magazine, April 2016, and "Mechanical doping: six-year ban for Femke Van den Driessche" (<https://road.cc/content/news/187658-mechanical-doping-six-year-ban-femke-van-den-driessche>), road.cc, April 2016.

The presence of one in a bike at a race is the offense The UCI charged technological fraud under articles 1.3.010 and 12.1.013bis of its regulations, which make the rider liable for the presence of a motor in a bike at a race regardless of intent or use. Cyclocross Magazine, as above.

banned her for six years Suspension from October 11, 2015 through October 10, 2021; fine of CHF 20,000; disqualification of all results from October 11, 2015 onward, including the 2015 European and 2016 Belgian under-23 titles; return of medals and prize money; payment of costs. Cyclocross Magazine and road.cc, as above.

In the HERITAGE Family Study Claude Bouchard, Ping An, Treva Rice, James S. Skinner, Jack H. Wilmore, Jacques Gagnon, Louis Pérusse, Arthur S. Leon, and D. C. Rao, "Familial aggregation of VO₂max response to exercise training: results from the HERITAGE Family Study," *Journal of Applied Physiology* 87, no. 3 (1999): 1003–1008, doi:10.1152/jappl.1999.87.3.1003. The abstract reports 481 sedentary adult Caucasians from 98 two-generation families, a 20-week standardized program, a mean increase in VO₂max of approximately 400 ml/min, individual responses ranging from little or no gain to more than 1.0 l/min, and a maximal heritability estimate of 47 percent. The gains are absolute, not weight-normalized.

An untrained young woman typically measures in the high thirties Typical values of about 38 ml/kg/min for untrained young women and about 45 for untrained young men, and the LeMond (92.5) and Dæhlie (96.0) figures, are as compiled at "World Best VO₂max Scores" (<https://www.topendsports.com/testing/records/vo2max.htm>), Topend Sports. The compilation warns that not all values are verified and that results depend on equipment, calibration, and protocol; another source gives Dæhlie's best as 90. The text treats them as a range for that reason.

won eight Olympic gold medals in the 1990s Guinness World Records, "Most Winter Olympic gold medals won in cross-country skiing (male)" (<https://www.guinnessworldrecords.com/world-records/77783-olympic-games-nordic-skiing-most-golds-men>): eight golds for Bjørn Dæhlie of Norway from 1992 to 1998, twelve medals in all.

reported as 80 inches against a height of 76 "What Makes Michael Phelps So Good?" (<https://www.scientificamerican.com/article/what-makes-michael-phelps-so-good1/>), Scientific American, 2008 (updated August 2016). The article attributes the wingspan figure to the Toronto Sun and the shoe-size and ankle-flexibility claims to a Detroit News blog, and quotes exercise physiologist Weiner doubting that Phelps is markedly more flexible than rival swimmers and suggesting fitness as the main difference.

earned less for about ten years Philip Oreopoulos, Till von Wachter, and Andrew Heisz, "The Short- and Long-Term Career Effects of Graduating in a Recession," American Economic Journal: Applied Economics 4, no. 1 (2012): 1–29, doi:10.1257/app.4.1.1. The abstract reports persistent earnings declines lasting ten years, initial placement at lower-paying employers, and partial recovery through moves to better firms, with less advantaged graduates affected more. The Canadian sample is identified in the paper's data section.

holds it to be self-evident that all men are created equal Declaration of Independence, July 4, 1776, second paragraph.

Chapter 2. What Is an Alpha Asset?

which Jay Barney set out in 1991 Jay B. Barney, "Firm Resources and Sustained Competitive Advantage," Journal of Management 17, no. 1 (1991): 99–120. The VRIN conditions (valuable,

rare, imperfectly imitable, non-substitutable) are developed there. The reframing into two tests is mine.

a large part of it belongs to the individual firm rather than to its industry Gabriel Hawawini, Venkat Subramanian, and Paul Verdin, "Is Performance Driven by Industry- or Firm-Specific Factors? A New Look at the Evidence," *Strategic Management Journal* 24, no. 1 (2003): 1–16. Their finding that firm effects dominate, and that a few exceptional firms account for much of the firm effect, was disputed by Gerry McNamara, Federico Aime, and Paul M. Vaaler, "Is Performance Driven by Industry- or Firm-Specific Factors? A Response to Hawawini, Subramanian, and Verdin," *Strategic Management Journal* 26, no. 11 (2005): 1075–1081, with a reply by Hawawini, Subramanian and Verdin in the same issue. The decompositions are statistical, not causal; the text claims only that persistent firm-level differences exist.

Barney made this argument in 1986 Jay B. Barney, "Strategic Factor Markets: Expectations, Luck, and Business Strategy," *Management Science* 32, no. 10 (1986): 1231–1241.

Russell Coff made this argument in 1999 Russell W. Coff, "When Competitive Advantage Doesn't Lead to Performance: The Resource-Based View and Stakeholder Bargaining Power," *Organization Science* 10, no. 2 (1999): 119–133.

Chapter 3. Lucky or Good?

which I call the detection budget The random-effects model, the detection budget $\text{SNR} \times \sqrt{N}$, the threshold $2.5/d$, the reliability weight (the budget squared, divided by one plus the budget squared), and the equicorrelation correction are developed in Karl T. Ulrich and Linda Zhao, "Lucky or Good? Outcome Noise, Effective Sample Size, and the Attribution of Skill," working paper, The Wharton School, 2026 (draft of

September 28, 2026, prepared for Judgment and Decision Making). The domain calibrations summarized in The Map, including the sources for each placement, are reported there.

are paid for exactly such industry-wide luck Marianne Bertrand and Sendhil Mullainathan, "Are CEOs Rewarded for Luck? The Ones Without Principals Are," Quarterly Journal of Economics 116, no. 3 (2001): 901–932.

roughly 30 percent against 18 percent for first-timers Paul Gompers, Anna Kovner, Josh Lerner, and David Scharfstein, "Performance Persistence in Entrepreneurship," Journal of Financial Economics 96, no. 1 (2010): 18–32. The figures are success rates for venture-backed founders with a prior venture-backed success versus first-time founders in their sample.

looked very much like what chance alone would produce Eugene F. Fama and Kenneth R. French, "Luck versus Skill in the Cross-Section of Mutual Fund Returns," Journal of Finance 65, no. 5 (2010): 1915–1947.

methods for controlling false discoveries Laurent Barras, Olivier Scaillet, and Russ Wermers, "False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas," Journal of Finance 65, no. 1 (2010): 179–216.

the most extreme performances are disproportionately produced by high-variance, unreliable processes Jerker Denrell and Chengwei Liu, "Top Performers Are Not the Most Impressive When Extreme Performance Indicates Unreliability," Proceedings of the National Academy of Sciences 109, no. 24 (2012): 9331–9336.

about two thirds of the persistence in top-quartile returns Morten Sorensen, "How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital," Journal of Finance 62, no. 6 (2007): 2725–2762. The allocation between sorting and influence is a structural estimate and depends on the model's matching assumptions.

Phil Rosenzweig made the case in The Halo Effect in 2007 Phil Rosenzweig, *The Halo Effect ... and the Eight Other Business Delusions That Deceive Managers* (New York: Free Press, 2007). The halo effect is delusion one; connecting the winning dots, lasting success, and absolute performance are among the eight others.

Richard Priem and John Butler argued in 2001 Richard L. Priem and John E. Butler, "Is the Resource-Based 'View' a Useful Perspective for Strategic Management Research?" *Academy of Management Review* 26, no. 1 (2001): 22–40.

the firms' profitability and investment moved Morten Bennedsen, Francisco Pérez-González, and Daniel Wolfenzon, "Do CEOs Matter? Evidence from Hospitalization Events," *Journal of Finance* 75, no. 4 (2020): 1877–1911.

MIT's economics department announced MIT Department of Economics, "Assuring an accurate research record" (<https://economics.mit.edu/news/assuring-accurate-research-record>), May 16, 2025. The statement says MIT "has no confidence in the provenance, reliability or validity of the data" and "no confidence in the veracity of the research" in the paper "Artificial Intelligence, Scientific Discovery, and Product Innovation," and that it asked arXiv to mark the paper withdrawn and contacted the journal to which it had been submitted. The statement does not name the author.

Chapter 4. How Does Luck Become an Asset?

More than fourteen thousand people took part Matthew J. Salganik, Peter Sheridan Dodds, and Duncan J. Watts, "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market," *Science* 311, no. 5762 (2006): 854–856. The experiment used 14,341 participants, 48 songs, one independent world and eight social-influence worlds; the authors report

that social influence increased both the inequality and the unpredictability of success, with the best songs rarely doing poorly and the worst rarely doing well.

Merton called it the Matthew effect Robert K. Merton, "The Matthew Effect in Science," *Science* 159, no. 3810 (1968): 56–63.

giving the early success away at random Arnout van de Rijt, Soong Moon Kang, Michael Restivo, and Akshay Patil, "Field Experiments of Success-Breeds-Success Dynamics," *Proceedings of the National Academy of Sciences* 111, no. 19 (2014): 6934–6939. The four settings were a crowdfunding site, a collaborative encyclopedia, a product-review site, and a petition site; initial successes were randomly assigned, recipients attracted more subsequent success than controls, and larger initial endowments did not yield proportionately larger gains.

more than twice as much research funding Thomas Bol, Mathijs de Vaan, and Arnout van de Rijt, "The Matthew Effect in Science Funding," *Proceedings of the National Academy of Sciences* 115, no. 19 (2018): 4887–4890. The study compares early-career applicants to the Netherlands Organisation for Scientific Research just above and just below the funding threshold, over the following eight years, and attributes the gap to both evaluators' later judgments and the losers' reduced application rates.

about two thirds of the persistence in top-quartile venture returns Morten Sorensen, "How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital," *Journal of Finance* 62, no. 6 (2007): 2725–2762. The allocation between sorting and influence is a structural estimate and depends on the matching model's assumptions.

about 30 percent likely to succeed again, against 18 percent Paul Gompers, Anna Kovner, Josh Lerner, and David Scharfstein, "Performance Persistence in Entrepreneurship," *Journal of Financial Economics* 96, no. 1 (2010): 18–32.

Boris Groysberg's study of star securities analysts Boris Groysberg, Linda-Eling Lee, and Ashish Nanda, "Can They Take It With Them? The Portability of Star Knowledge Workers' Performance," *Management Science* 54, no. 7 (2008): 1213–1230. Analysts ranked by Institutional Investor; performance generally declined after a move, less so when the analyst moved with colleagues or to a firm with comparable capabilities. Moves were not randomly assigned.

Chapter 5. How Does Advantage Accumulate?

Taylor Randall and I examined prices in the mountain bike industry Taylor Randall, Karl Ulrich, and David Reibstein, "Brand Equity and Vertical Product Line Extent," *Marketing Science* 17, no. 4 (1998): 356–379. The study uses mountain bicycle price data to estimate brand premiums.

the production system that Taiichi Ohno and Eiji Toyoda built at Toyota Jeffrey K. Liker, *The Toyota Way: 14 Management Principles from the World's Greatest Manufacturer* (New York: McGraw-Hill, 2004), for the system's history and principles.

Gabriel Szulanski studied the transfer of proven practices Gabriel Szulanski, "Exploring Internal Stickiness: Impediments to the Transfer of Best Practice Within the Firm," *Strategic Management Journal* 17, Winter Special Issue (1996): 27–43. Causal ambiguity, the recipient's absorptive capacity, and arduous relationships between source and recipient were the main predictors of stickiness.

Nicholas Bloom and his colleagues randomly assigned months of management consulting Nicholas Bloom, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts, "Does Management Matter? Evidence from India," *Quarterly Journal of Economics* 128, no. 1 (2013): 1–51 (about 17 percent productivity gain in the first year). The nine-year follow-up is Nicholas

Bloom, Aprajit Mahajan, David McKenzie, and John Roberts, "Do Management Interventions Last? Evidence from India," *American Economic Journal: Applied Economics* 12, no. 2 (2020): 198–219.

Linda Argote and her colleagues found in shipyard production data Linda Argote, Sara L. Beckman, and Dennis Epple, "The Persistence and Transfer of Learning in Industrial Settings," *Management Science* 36, no. 2 (1990): 140–154. In ship production, knowledge depreciated rapidly enough that cumulative output overstated the persistence of learning.

Chapter 6. Where Should You Apply Your Strengths?

the New York Times reported from the Milan men's wear shows "Flying Shirttails: The New Pennants of Rebellion," *New York Times*, July 20, 2004. The 2022 draft quotes the passage on the untucked shirttails of the American men's wear press and buyers.

James March named the general tension exploration and exploitation James G. March, "Exploration and Exploitation in Organizational Learning," *Organization Science* 2, no. 1 (1991): 71–87.

Lele Sang and I studied eight multinational companies entering China Lele Sang and Karl T. Ulrich, *Winning in China: 8 Stories of Success and Failure in the World's Largest Economy* (Philadelphia: Wharton School Press, 2021). The preconditions for success in a new geography are adapted from that book.

Herminia Ibarra's research on career change Herminia Ibarra, *Working Identity: Unconventional Strategies for Reinventing Your Career* (Boston: Harvard Business School Press, 2003; updated edition 2023). Her findings rest on qualitative inter-

views with people who changed careers and emphasize experiment and new relationships over analysis.

Paul David made the historical point about electricity Paul A. David, "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *American Economic Review* 80, no. 2 (1990): 355–361.

a customer-support operation studied by the economist Erik Brynjolfsson and his colleagues Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," *Quarterly Journal of Economics* 140, no. 2 (2025): 889–942. Average productivity gain of about 14 percent, concentrated among less experienced and lower-skilled agents.

Chapter 7. How Do You Build an Advantage Deliberately?

a startup's technology should be ten times better Peter Thiel with Blake Masters, *Zero to One: Notes on Startups, or How to Build the Future* (New York: Crown Business, 2014). The rule of thumb is stated in the chapter on building a monopoly.

Daniel Kahneman and Gary Klein, who disagreed about most things, agreed on this Daniel Kahneman and Gary Klein, "Conditions for Intuitive Expertise: A Failure to Disagree," *American Psychologist* 64, no. 6 (2009): 515–526.

Anders Ericsson's research on expert performance K. Anders Ericsson, Ralf Th. Krampe, and Clemens Tesch-Römer, "The Role of Deliberate Practice in the Acquisition of Expert Performance," *Psychological Review* 100, no. 3 (1993): 363–406; and Anders Ericsson and Robert Pool, *Peak: Secrets from the New Science of Expertise* (Boston: Houghton Mifflin Harcourt, 2016). The finding that the variance explained by deliberate practice differs widely by domain is Brooke N. Macnamara, David Z.

Hambrick, and Frederick L. Oswald, "Deliberate Practice and Performance in Music, Games, Sports, Education, and Professions: A Meta-Analysis," *Psychological Science* 25, no. 8 (2014): 1608–1618.

Arnaldo Camuffo and his colleagues ran four randomized trials Arnaldo Camuffo, Alfonso Gambardella, Danilo Messinese, Elena Novelli, Emilio Paolucci, and Chiara Spina, "A Scientific Approach to Entrepreneurial Decision-Making: Large-Scale Replication and Extension," *Strategic Management Journal* 45, no. 6 (2024): 1209–1237. Treated firms terminated more projects and showed higher average cumulative revenue over the observation period; the simple claim that the treatment increased pivoting did not replicate consistently.

Chapter 8. When Does Advantage Stop Working?

Everett Rogers spent a career cataloguing why good ideas diffuse slowly Everett M. Rogers, *Diffusion of Innovations*, 5th ed. (New York: Free Press, 2003). The scurvy case (James Lind's 1747 trial; the British navy's adoption of citrus in 1795) opens the book's discussion of the slow diffusion of a demonstrably effective innovation, and the five attributes of innovations that affect their rate of adoption are developed in chapter 6.

Andrew King and Baljir Baatartogtokh went through seventy-seven cases Andrew A. King and Baljir Baatartogtokh, "How Useful Is the Theory of Disruptive Innovation?" *MIT Sloan Management Review* 57, no. 1 (Fall 2015): 77–90.

Clayton Christensen's account Clayton M. Christensen, *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail* (Boston: Harvard Business School Press, 1997).

Bill Gates's 1995 memo turning Microsoft toward the internet Bill Gates, "The Internet Tidal Wave," internal Microsoft memo-

random, May 26, 1995, made public in the United States v. Microsoft proceedings and reproduced at Letters of Note.

Chapter 9. What Should Success Mean?

Buffett calls this the ovarian lottery Buffett has told the ovarian-lottery thought experiment on many occasions; a published version appears in Carol J. Loomis, *Tap Dancing to Work: Warren Buffett on Practically Everything, 1966–2012* (New York: Portfolio/Penguin, 2012).

In 1971 the philosopher John Rawls asked John Rawls, *A Theory of Justice* (Cambridge, Mass.: Harvard University Press, 1971; revised edition 1999). The original position and the veil of ignorance are set out in chapter 1, sections 3 and 4; the “natural lottery” and the judgment that its outcome is “arbitrary from a moral point of view” are in chapter 2, section 17; the difference principle in sections 12 and 13.

The strongest objection to that conclusion came from Robert Nozick Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), chapter 7, section II, on Rawls’s treatment of natural assets; the argument that the undeservedness of a talent does not entail the undeservedness of what flows from it is developed there.

Psychologists who study motivation have converged on three things Richard M. Ryan and Edward L. Deci, “Self-Determination Theory and the Facilitation of Intrinsic Motivation, Social Development, and Well-Being,” *American Psychologist* 55, no. 1 (2000): 68–78.

Sherwin Rosen showed in 1981 Sherwin Rosen, “The Economics of Superstars,” *American Economic Review* 71, no. 5 (1981): 845–858.

Experiments in which brain activity predicting a choice is detectable before the subject reports deciding Benjamin Libet, Curtis A. Gleason, Elwood W. Wright, and Dennis K. Pearl, "Time of Conscious Intention to Act in Relation to Onset of Cerebral Activity (Readiness-Potential)," *Brain* 106, no. 3 (1983): 623–642. For the case that such findings imply the absence of free will, Robert M. Sapolsky, *Determined: A Science of Life Without Free Will* (New York: Penguin Press, 2023); for the case that they do not, Alfred R. Mele, *Free: Why Science Hasn't Disproved Free Will* (New York: Oxford University Press, 2014).

the compatibilists Daniel C. Dennett, *Elbow Room: The Varieties of Free Will Worth Wanting* (Cambridge, Mass.: MIT Press, 1984), is the classic statement.

A board that fires a chief executive after three years of poor relative returns, which boards routinely do Dirk Jenter and Fadi Kanaan, "CEO Turnover and Relative Performance Evaluation," *Journal of Finance* 70, no. 5 (2015): 2155–2184, find that chief executives are significantly more likely to be dismissed after poor industry- and market-wide performance that they did not cause.

the voter who punishes the government for a recession it did not cause, which voters reliably do Christopher H. Achen and Larry M. Bartels, *Democracy for Realists: Why Elections Do Not Produce Responsive Government* (Princeton: Princeton University Press, 2016), chapters 5 and 6, on retrospective voting and blind retrospection.

About the Author

Karl T. Ulrich won the ovarian lottery. He was born in 1960, a boy, in Boston, to two healthy, college-educated parents who stayed married and read to their children. He was not the smartest of the five. Whether from birth order, DNA, or both, he finishes what he starts, which has been worth more than it should be. About eighty percent of his venture investments went to zero, and the portfolio has roughly matched an index fund.

He is the CIBC Professor at the Wharton School and a co-founder of Xootr, Terrapass, Belle-V Kitchen, and MakerStock. He is also an author of *Product Design and Development*, *Innovation Tournaments*, *Design: Creation of Artifacts in Society*, *The Innovation Tournament Handbook*, *Winning in China*, *China Speed*, and *The Hundred-Year Chair*.

Colophon

Most of this book was written between 2020 and 2022 as a book about firms. The project stalled, and was revived in 2026 once highly competent language models became widely available. The 2026 revision, which rebuilt the book around advantage for people as well as organizations and added the chapters on what records can reveal, on how luck becomes an asset, and on what success should mean, was drafted with Claude Fable 5.1 from the author's manuscript, papers and notes, and then revised by the author. Tim Gray, the author's long-time and human editor, then scrubbed the full text.

The manuscript lives in a GitHub repository, one folder per chapter, with every citation held in a visible verification queue until it has been checked against its source. The chapters were rendered automatically at alphaas-sets.com as they were written, and the book was drafted in public with comments invited from friends and colleagues. The production files were generated by the repository's own build, which typesets the Markdown chapters in XeLaTeX using the memoir class and the TeX Gyre Pagella typeface, at six by nine inches, on every commit.

Published by Pontifica Press and distributed by Amazon and Ingram.

Proof pages: headings, a list, a chart, a table

This chapter exists so that the physical proof contains every kind of page the finished book will contain. It is set in the same measure, type scale and grid as the chapters, and it should be removed from the file before the final build by setting the switch at the top of `main.tex`.

A section heading, with a subsection beneath it

The section heading above sits 22.5 points below the preceding text and 7.5 points above this paragraph, so that the heading line and its two spaces together occupy three lines of the body grid. The first paragraph after a heading is not indented. The paragraphs that follow carry the ordinary 11 point indent, and nothing is added between them.

A subsection heading

Subsection headings are one size down from section headings and sit 15 points below the preceding text. With the 7.5 points after, a subsection occupies two and a half lines of the grid, which is the one place in the specification where the grid is deliberately broken; the alternative, 15

points before and nothing after, keeps the grid but reads as cramped. Either is defensible, and the proof is the place to decide.

Lists are brief and parallel. Three questions follow from the motor at Zolder:

- Whether an advantage exists is a question about mechanism: does the thing add power?
- What explains a particular result is a question about a mixture that no single result reveals.
- What credit the winner is due depends on the first two and on where the advantage came from.

A block quotation is indented 12 points on both sides and set one half point smaller than the body, with the same 15 point leading so that it stays on the grid:

A resource with both properties, one that improves performance and that rivals cannot easily get, I will call an alpha asset. Competent is not the same as advantaged. An advantage the rivals can match is not an advantage.

The quotation keeps its attribution with it, and short quotations are kept together on one page.

A chart at the full measure

Charts are drawn for the 279 point text width at final size, in black and grays, with labels of 9 to 10 points. The exam-

ple below is a bar chart of the kind chapter 3 will use: outcomes for a population, with the right tail shaded.




Figure 1. Outcomes for a population of firms, placeholder values. Bars above the ninetieth percentile are solid.

The caption is set in the sans at 9 on 12, six points below the artwork, and the figure is never separated from its caption.

A table

Tables are set in the sans at 9.5 on 12 with tabular lining numerals, right-aligned where the column is numeric, with horizontal rules only. The example is the kind of comparison chapter 5 makes between mechanisms of accumulation.

Flywheel	Stock	Years to build	Decay, %/yr
Customer network	relationships	5 to 10	8
Brand equity	reputation	10 to 20	4
Cost efficiency	learning, scale	3 to 7	12
Product quality	design, process	2 to 5	15
Organizational capability	routines, people	7 to 15	6

Numbers in the table are illustrative. The table sits inside the measure with a rule above and below and one rule under the header, and the header is set in the sans semi-bold.