

Chapter One

The Knowledge Funnel: How Discovery Takes Shape

The United States in the years after World War II was a restless place, engaged in an audacious social experiment that would eventually transform how and where Americans worked, played, and consumed. Victorious with its Allies against the Axis, then instrumental in the rebuilding of Europe, the U.S. had become the West's undisputed military and economic champion. Not all was cheery, mind you. Communist powers posed a new threat, the nuclear age had dawned, and Cold War anxieties ran high. Yet Americans felt confident in their place atop the global order and free to invent new ways to enjoy their burgeoning prosperity.

The automobile was central to the sense of open-ended possibility shared by America's rapidly growing middle class. Spurred by the development of the Interstate Highway System, new roads were rolling out from the cities to the suburbs springing up at their edges. The number of cars sold in America leapt from just 70 thousand in 1945 to more than 6 million in 1950. A mobile, moneyed lifestyle was taking root, with the automobile at the center of it.

Some of the first entrepreneurs to see the opportunities in this cultural change planted their flags in California, where so many American trends first take root. Drive-in burger joints began to spring up across southern California, where the nascent car culture cross-fertilized a leisure culture centered on the beach. By 1955, a strong-willed salesman named Ray Kroc was able to

make a good living selling milk-shake mixers to a wide collection of mom-and-pop So-Cal restaurants. His biggest account was the McDonald brothers, who operated a small but thriving chain of drive-ins in the Los Angeles suburbs.

The brothers had opened their first restaurant, a barbeque and burger drive-in in San Bernardino, in 1940. It wasn't much different from other drive-ins, which had been popping up ever since A&W first delivered root beer to car windows in 1923. But it was popular. The McDonald's outpost attracted throngs of teenagers, with harried carhops serving up to 125 carloads at a time.

Within the decade, though, the McDonald brothers realized they had to revamp their restaurant or find a new line of work. Some of their best customers were families giving Mom a night off from the kitchen. But now these families were driving right past, turned off by the loitering toughs that drive-ins attracted. Many of the remaining customers complained that the food got cold on the journey from kitchen to car. The McDonald brothers needed a new approach, but what? How and what did Californians want to eat when they set out in their Fords and Buicks and Studebakers in determined pursuit of sun and surf?

The brothers experimented with different menus and store formats until they arrived at a winning approach. They filled in the barbeque pit, cut the menu to only 25 items, and standardized the burgers—each one was served with ketchup, mustard, onions, and two pickle slices. The carhops were eliminated, replaced with service windows where customers ordered and picked up their own food. Productivity enhancers like Kroc's five-at-a-time milkshake mixers enabled them to

turn food orders around quickly. The brothers called their new concept the Speedee Service System. It was the prototype of the quick-service restaurant.

It wasn't long before the brothers had opened four additional outlets. They may have been content to stop there, but Kroc was not. He looked at the crowds packed into the brothers' stores and imagined the scene repeated from coast to coast—and even around the world. He bought out the McDonalds and set about improving and standardizing the mass-production system they had developed.

Kroc saw that the Speedee Service System, innovative as it was, left too much to chance and judgment. He refined it meticulously, pursuing a vision of a perfectly standardized operation. He simplified the McDonald's system down to an exact science; with a rigid set of rules that spelled out *exactly* how long to cook a hamburger, *exactly* how to hire people, *exactly* how to choose locations, *exactly* how to manage stores, and *exactly* how to franchise them. Under Kroc, nothing in the McDonald's kitchen was left to chance: every hamburger came out of a stamping machine weighing exactly 1.6 ounces, its thickness measured to the thousandth of an inch, and the cooking process stopped automatically after 38 seconds, when the burgers reached an internal temperature of exactly 155 degrees. In every phase of McDonald's operations, judgment was removed; possibilities were removed; variety was removed.

Kroc relentlessly stripped away uncertainty, ambiguity, and judgment from the processes that emerged from the McDonald brothers' original insight. And by fine-tuning the formula, he powered McDonald's from a modestly prosperous chain of burger restaurants to a scale

previously undreamed-of. Within a decade, McDonald's grew from successful local business to a ubiquitous cross-country chain and, in another few decades, to a globe-girdling behemoth.

The path taken by the McDonald's and Kroc — from pinpointing a market opportunity to devising an offering for that market to codifying its operations — is not just a study in entrepreneurship. It's a model for how businesses of all sorts can advance knowledge and capture value. I will argue in this book that the McDonald brothers and Kroc took the same route followed by successful business innovators in every domain. My term for that path is the **knowledge funnel**. My purpose in this book is to map that funnel in detail and investigate its implications for organizations, individuals, and thinking processes. Along the way, we'll meet innovators in business, science, and the arts, all of whom are advancing understanding and creating exciting new opportunities for people and organizations.

<A> **Seeking Reconciliation**

The model for value creation offered in this book requires a balance—or more accurately a reconciliation—between two prevailing points of view on business today. One school of thought, put forward by some of the world's most respected theorist and consultants, holds that the path to value creation lies in driving out old-fashioned practices of gut feeling and instinct, replacing them with strategy based on rigorous, quantitative analysis (optimally backed by decision-support software). In this model, the basis of thought is **analytical thinking**, which harnesses two familiar forms of logic—deductive and inductive reasoning—to declare truths and certainties about the world. The goal of this model is mastery through rigorous, continuously

repeated analytical processes. Judgment, bias, and variation are the enemies. If they are vanquished, the theory goes, great decisions will be made and great value will be created.

The opposing school of thought, which is in many ways a reaction to the rise of analytical management, is centered on the primacy of creativity and innovation. To this school, analysis has driven out creativity and doomed organizations to boring stultification. “The minute you start analyzing and using consumer research, you drive all the creativity out of the product,” the vice-chairman and chief of design for a world-leading American firm told me recently. “No good product was ever created from quantitative market research. Great products spring from the heart and soul of a great designer, unencumbered by committees, processes, or analyses.” To proponents of this philosophy, the creative instinct, the unanalyzed flash of insight, is venerated as the source of true innovation. At the heart of this school is **intuitive thinking**—the art of knowing without reasoning. This is the world of originality and invention.

These two models seem utterly incommensurable; an organization must choose to embrace either analysis or intuition as the primary driver of value creation. This choice then plays out in the structure and norms of the organization. Organizations dominated by analytical thinking are built to operate as they always have—they are structurally resistant to the idea of designing and redesigning themselves and their business dynamically over time. They are built to maintain the status quo. By sticking closely to the tried and true, organizations dominated by analytical thinking enjoy one very important advantage - they can build size and scale. In organizations dominated by intuitive thinking, on the other hand, innovation may come fast and furious but

growth and longevity represent tremendous challenges. Intuition-biased firms cannot and will not systematize what they do, so they wax and wane with individual intuitive leaders.

Neither analysis nor intuition alone is enough. Rather than forcing a binary choice to drive out either analysis or intuition, the burden of this book is to reconcile the two modes of thought. I will argue that aspects of both analytical and intuitive thinking are necessary but not sufficient for optimal business performance. The most successful businesses in the years to come will balance analytical mastery and intuitive originality in a dynamic interplay that I call **design thinking**. Design thinking is the form of thought that enables movement along the knowledge funnel, and the firms that master it will gain a nearly inexhaustible long-term business advantage. The advantage, which emerges from the design-thinking firms' unwavering focus on the creative design of systems, will eventually extend to the wider world. From these firms will emerge the breakthroughs that move the world forward.

Design-thinking firms stand apart in their willingness to engage in the task of continuously redesigning their business. They do so with an eye to creating advances in both innovation and efficiency—the combination that produces the most powerful competitive edge. This is not to suggest that only design-thinking firms pursue innovation. No, the value that business leaders place on innovation is reflected in the wealth of resources that they devote to its pursuit. But in all too many cases, businesses unwittingly work against their own purposes. Even as corporate leaders chase the vital, elusive spark of creativity, their organizations' structures, processes, and norms extinguish it wherever it flares up. Their cultures and routines privilege analysis over intuition and mastery over originality.

As A.G. Lafley has demonstrated at Procter & Gamble, however, even organizations with a deeply ingrained bias toward analysis and mastery can develop powerful capacities for innovation. With determined leadership, they can develop the skills, structures, and processes that generate value by driving valuable insights along the knowledge funnel. Here, briefly, is how knowledge proceeds through the funnel.

<INSERT CHART OF KNOWLEDGE FUNNEL HERE (To come)>

The first stage of the funnel is the exploration of a **mystery**, which takes an infinite variety of forms. A research scientist might explore the mystery of a syndrome such as autism. A hospital administrator might ask what kind of space would improve the condition of cancer patients coping with chemotherapy. Or an ambitious salesman might ask how and what Americans would like to eat on the go.

The next stage of the funnel is a **heuristic**, a rule of thumb that helps narrow the field of inquiry and work the mystery down to a manageable size. The heuristic may be a genetic anomaly, a user-centered approach to the process flow of a chemotherapy patient, or the concept of a quick-service, drive-through restaurant. It is a way of thinking about the mystery that provides a simplified understanding of the mystery and allows heuristic holders to focus their efforts.

As an organization puts its heuristic into operation, studies it more, and thinks about it intensely, it can convert from a general rule of thumb (Americans want a quick, convenient, tasty meal) to a fixed formula (Kroc's totally systematized McDonald's). That formula is an **algorithm**, the last of three stages of the knowledge funnel.

<sidebar>

<A> Hunches, Heuristics and Algorithmics: A Quick Note

By Mihnea Moldoveanu

It is important to be precise about concepts, because they are critical building blocks of any human enterprise, intellectual, and otherwise. One way to analyze concepts is to prescribe use cases for the words that denote them. To wit:

The route out of a mystery begins with a hunch. Hunches are *pre-linguistic intuitions*. You are in a dense fog high up in the Rocky Mountains. Darkness is on its way. You can see no more than five feet ahead. As you worry about your next step – and the rest of the way – your peripheral vision ‘sees’ a slanted spruce at 15:00 and you experience a ‘sense’ that you should turn right. If someone were to ask you at the time “Why do you want to go right?” you could not, of course, answer the question in a way that would seem objective to that person. “Just a hunch,” you might say. “Something beyond words.” You turn and you get safely to the lodge. You never do become aware of the fact that you have seen this tree before, on the way to the woods. The hunch remains a hunch. It remains beyond words, but not, obviously, beyond either reason or sense.

Heuristics are *open-ended prompts* to think or act in a particular way. For instance ‘Look in the rear view mirror before passing’; ‘Go with the first instinct when trying to decide if someone is lying to you in a face-to-face interaction’ (this is a heuristic that recognizes the value of a hunch); or, ‘Buyer beware!’ Heuristics offer no *guarantee* that using them produces a certain result. Rather, they contain the vague promise that, all things being equal, using the heuristic in the context it is meant for *may*, or, on average, *will* be better for you than not using it. Heuristics are different from hunches in that they are *explicit*: they bring intuitions to language.

Algorithms are *certified production processes*. They *guarantee* that, in the absence of intervention or complete anomaly, following the sequence of steps they embody will produce a particular result. For instance, an algorithm (PRIME_SEARCH) designed to figure out if a given number is a prime number - by brute force -will systematically try to divide that number by every number smaller than itself, and return the answer ‘PRIME’ if no divisor is found and ‘Divisor = ...’ if a divisor is found. Algorithms differ from heuristics in that they offer a *performance guarantee* that comes along with using them: you cannot use the algorithm PRIME_SEARCH on the number 209870987403987 and *not* get to an answer, *except* if some catastrophe intervenes and stops you from executing the steps prescribed by the algorithm.

</sidebar>

Each stage of the knowledge funnel has its own unique features that are worth examining in some detail. It is said that the road to wisdom begins with ignorance, and that is where we begin, too.

<A> It Starts with a Question

Over the course of time, phenomena enter our collective consciousness as mysteries—things in our environment that excite our curiosity but elude our understanding. The mystery of what we now know as gravity confounded our ancestors: when they looked around them, they saw that most objects—apples, famously—seemed to fall to the ground quickly; but others, such as leaves, seemed to take forever to reach the ground. And then there were birds, which didn't seem to fall at all. In the visual arts, one of the most enduring mysteries was how to represent what we see in front of us in three dimensions on a two-dimensional surface. In both cases, people struggled for centuries to come to an understanding of the phenomena. Even the most baffling mysteries, though, eventually crumble under the force of human intelligence. With sufficient thought, a first-level understanding emerges from the question at hand. We develop heuristics—rules of thumb—that guide us toward a solution by way of organized exploration of the possibilities.

Consider the falling objects. After a long period of observation and contemplation, human beings in various cultures more or less simultaneously developed the notion of a universal force that tends to pull physical objects earthward. Understanding advanced from a mystery—why do things fall to earth?—to a heuristic, a rule of thumb for explaining why things fall—a force we call gravity causes things to fall to earth. Heuristics represent an incomplete yet distinctly advanced understanding of what was previously a mystery.

In art, after literally centuries of questioning and experimentation, the heuristic of perspective emerged as a solution to the mystery of three-dimensional representation. First, in about the fifth

century BC, came a tool called skenographia, which historians conjecture was developed by Greek dramatists to make their sets appear to have depth. A heuristic had begun to emerge.

Heuristics represent an incomplete yet distinctly advanced understanding of what was previously a mystery. But that understanding is unequally distributed—some people remain stuck in the world of mystery, while others master its heuristics.

The beauty of heuristics is that they guide us toward a solution by way of organized exploration of the possibilities. With a heuristic to guide his further thought and consideration, the great scientist Sir Isaac Newton derived precise rules for determining how fast an object will fall under any circumstance. Newton's rule—that an object dropped from any height will accelerate at a constant rate of 32 feet per second²—advanced the understanding of gravity to the third stage—the algorithm. An algorithm is an explicit, step-by-step procedure for solving a problem.

Algorithms take the loose, unregimented heuristics—which take considerable thought and nuance to employ—and simplify, structuralize, and codify them to the degree that anyone with access to the algorithm can deploy it with more or less equal efficiency.

As with gravity, the algorithm for perspective took centuries to develop. By the eleventh century AD, early physicists had arrived at the understanding that the conical shape of the eye influences how we see three-dimensional objects. A few centuries later, the Florentine painter and architect Filippo Brunelleschi studied the heuristic until he innovated a repeatable method—an algorithm—that allowed him and other artists to reliably create the illusion of three-dimensional space.

As understanding moves from mystery to heuristic to algorithm, extraneous information is pared away—the complexities of the world are mastered through simplification. That is why my graphic model of the advance of knowledge is a funnel that tapers as knowledge moves through its stages of refinement. The gain in understanding comes from picking salient features of the environment and out of them constructing a causal explanation of the mystery. From the inchoate phenomenon of falling objects came the concept of a universal force that pulls things earthward, which in turn was painstakingly developed, through trial and error, into a simple formula that described the unchanging properties of this once-mysterious force.

There's significant value to pushing knowledge to the algorithm stage. It is quite handy to have at one's disposal a logical, arithmetic or computational procedure that, if correctly applied, guarantees success. When Brunelleschi created the precise "vanishing point" algorithm for perspective during the first two decades of the fifteenth century, he provided a significant advantage to the Florentine artistic community until the algorithm became more widely disseminated and understood.

The ultimate destination of algorithms as of the late 20th century is computer code. Once knowledge has been pushed to a logical, arithmetic or computational procedure, it can be reduced to software. Armed with the algorithm for gravity, clever engineers at Honeywell were able to create autopilot systems for giant commercial aircraft so that they could be made to fall out of the sky in a passenger-friendly fashion without human intervention. And what about Brunelleschi's algorithm for perspective? Computers now use the three-dimensional data transferred from a camera to spit out a two-dimensional representation of it based on the

formulae handed down by Brunelleschi and codified in matrix-multiplication software.

Of course, not every mystery can become an algorithm; not all logic can be pushed through to the end of the funnel. Consider the mystery of the oldest art, music. How can certain arrangements of notes, timbres, and rhythms have such a profound effect on our emotions, and how can we harness that power to soothe or rouse our listeners? Norman Greenbaum stumbled upon the answer to that mystery once and once only, coming up with the 1969 smash “Spirit in the Sky.” Wildly catchy and instantly recognizable, the song continues to spin off royalties that provide Greenbaum with a comfortable living. But the mystery of the hit song remains just that for Greenbaum. He has never produced a follow-up to the fuzzed-up hippie spirituality of “Spirit in the Sky.”

Now contrast Greenbaum’s career with that of U2, the band that developed a heuristic—a way of understanding the world and conveying that understanding through harmony, melody, and rhythm—that enables it to write songs that resonate with millions of people around the world, not once but over and over again. From the release of the earnest, anthemic album *Boy* in 1980 to the eclectic pleasures of *Achtung Baby* in 1991, U2’s mastery of heuristics produced a string of industry awards and top 40 hits around the world. But when the band consciously stepped away from the heuristic that had served it so well—experimenting with techno, dance, and electronica on *Zooropa* and *Pop*—fans promptly voted with their feet. When, in 2000, the band reunited with producers Brian Eno and Daniel Lanois to record *All That You Can’t Leave Behind*, they also returned to their pre-*Zooropa* heuristic, leading to Bono’s famous remark at that year’s Grammy Awards: “The whole year has been quite humbling,” he said. “Going back to scratch,

reapplying for the job. What job? The best-band-in-the-world job.” The heuristic still worked; *Rolling Stone* called *All That You Can't Leave Behind* U2's third masterpiece (after *The Joshua Tree* and *Achtung Baby*).

Yet even U2's greatest albums contain some forgettable songs—their mastery of the heuristics of the pop song falls short of a surefire algorithmic formula. The occasional failures of a serial hit-maker like U2 tell us something important about heuristics: They don't guarantee success.

Heuristics can do no more than increase the probability of getting to a successful outcome, or at least getting there more quickly.

Thus far at least, pop music has proven resistant to advance from heuristic to algorithm. But there have been movements in that direction: In the late 1970s, musical innovators like producer Brian Eno experimented with the sound of the human heartbeat and determined that songs with a synthesized heartbeat as their rhythm track are instinctively enjoyed by listeners, no matter what musical setting sits atop the heartbeat. As a producer, he was able to help bands turn out hits in a variety of genres, from the jittery dance-pop of Talking Heads' "Once in a Lifetime" to the orchestral strings of Coldplay's "Viva la Vida" to those massively successful U2 albums. Other producers in search of a success algorithm created a succession of disposable boy bands, pop princesses, or lip-synching electro-pop acts like Milli Vanilli. And even now we have the mass populism of Simon Fuller's *American Idol*, which has produced bona fide stars in Kelly Clarkson and Carrie Underwood, and a few forgettable flashes in the pan. The algorithm remains elusive. There is still nothing close to a formula for producing consistent success in the music business. Yet.

<A> Back to McDonald's

Incongruous as it might sound, the McDonald brothers and Ray Kroc followed the same path that Newton and Brunelleschi trod as they built their business from a single drive-in to a global enterprise. Their journey began with the question that so perplexed the McDonald brothers as they watched a new culture grow up around them: What and how did the mobile, leisured, mass middle class of southern California want to eat? That was their mystery.

The brothers devised an answer by focusing on a specific facet of that emerging culture—the consumers' desired out-of-home eating experience. The heuristic they developed—a quick-service restaurant with strictly limited menu options—emerged when they narrowed the field of possibilities to a manageable set of salient features. In doing so, the McDonald brothers discovered a way to create value from their understanding of their world.

Kroc then picked up the baton, driving that understanding—that heuristic—all the way to an algorithm by continuing to cut away vast tracts of possibility. Hamburgers could be charbroiled or pressure-cooked. The menu could be broad or narrow. Restaurants could be smaller or larger. Ultimately, Kroc plucked one answer along innumerable dimensions to construct McDonald's defining algorithm. Once that algorithm was in place, Kroc pushed it as far as it would go, adapting its elements to changing markets and economic conditions but leaving its essential outlines unchanged.

<A> The Creation of Value in Business

The McDonald's story illustrates important elements of the dynamics of the march of knowledge from mystery to heuristic to algorithm. As a business moves from mystery toward algorithm, information is pared away—the complexities of the world are mastered through simplification. The gain in understanding comes from picking out salient features of the environment and out of them building a causal understanding of it—"I think that Californians would like a quick-service hamburger joint." The heuristic doesn't attempt an encyclopedic understanding of the new Californian beach culture and how the Freeway system brought it into being. It focuses instead on a specific facet of that culture—the consumers' desired out-of-home eating experience.

To create an algorithm from that heuristic requires clear-cutting more vast tracts of possibility. Ultimately, one answer along innumerable dimensions had to be plucked to provide McDonald's' defining algorithm. Judgment was removed; possibilities were removed; variety was removed.

What is the value to a business of driving through the knowledge funnel from mystery to heuristic to algorithm? The reward is a massive gain in efficiency. By paring away possibilities from the mystery of what and how Californians want to eat to the limited menu, drive-through quick-service burger joint, the McDonald brothers could focus on a few important things and to replicate the model several times over, extending its success.

When Kroc converted the heuristic into a precise algorithm, he was able to scale the chain to a size previously unimaginable. Restaurant site selection followed an efficient algorithm, so sites could be found and developed quickly in each desired locale. Staffing the restaurant was easy,

because the procedures for hiring the unskilled labor needed were precisely laid out and the new employees could be readily taught the precise in-restaurant procedures from comprehensive manuals. Supply of food and beverage items to the new restaurant could be easily added to the precisely organized supply chain, making that supply chain even higher-scale and more efficient still.

By solving the mystery before its competitors, McDonald's created an efficiency advantage. By honing and refining the heuristic, it extended that efficiency advantage. By converting that heuristic to algorithm, new owner Kroc drove the efficiency advantage still further ahead of its competitors, creating an enterprise worth billions of dollars—all from one new-style burger joint.

<A> A Fine Balance

Matching McDonald's accomplishment—and that of every other organization that creates value across the knowledge funnel—requires two very different activities: *moving across the knowledge stages* of the funnel—from mystery to heuristic and heuristic to algorithm; and *operating within each knowledge stage* of the funnel—honing and refining an existing heuristic or algorithm. We can map these two different activities onto the theories of the great management theorist Jim March, who posited that organizations may engage primarily in **exploration**—the search for new knowledge (in our terms, seeking movement across the knowledge stages)—or **exploitation**—the maximization of payoff from existing knowledge (refinement within a knowledge stage). Both activities can create enormous value and both are critical to the success of any business organization. But they are hard to engage in simultaneously; most often, organizations choose to focus on one activity— either exploration or

exploitation – to the exclusion of the other and to their own detriment.

An organization exclusively dedicated to exploration will expire in relatively short order.

Typically, exploration alone will not generate the returns needed to fund further exploration.

Imagine Norman Greenbaum as a corporation. After the one random incident of successful intuitive thinking, of exploration, that created “Spirit in the Sky,” Greenbaum Inc. would have gone bust waiting for the next chance event—which never happened. Devotion to exploration is *the invention of business*, a risky proposition and the reason that 9 of 10 entrepreneurial start-ups expire in less than two years. Exploration alone is unstable business.

On the other hand, many organizations flip quickly from an early exploration phase – the generation of the founding idea behind the business – to the steady exploitation of that idea, never returning to exploration. These organizations, solely dedicated to exploitation, might last somewhat longer than exploration-only businesses, but the business that creates value only through exploitation will exhaust itself in due course. It can’t keep exploiting the same piece of knowledge forever. If it tries to do so, the cost to the business can be devastating.

The exploitation of knowledge within a given stage – that is, running an existing heuristic, gently honing and refining it, but not seeking to move knowledge to algorithm or running an existing algorithm and not seeking to explore the next mystery—is the *administration of business*. A high-end Wall Street law firm runs the legal-services heuristic over and over; McDonald’s runs the fast-food algorithm over and over. That is a far different activity from the exploration that

drives knowledge from one stage to the next—from mystery to heuristic, from heuristic to algorithm.

<insert Table 1-1>

The vast majority of businesses follow a common path. The company is birthed through a creative act that converts a mystery to a heuristic through intuitive thinking. It then hones and refines that heuristic through increasingly pervasive analytical thinking and enters a long phase in which the administration of business dominates. And in due course, a competitor stares at the mystery that provided the spark for this company, comes up with a more powerful heuristic and supplants the original business.

A small fraction of those companies generate a second intuitive breakthrough—often, as in the case of McDonald’s, from a new owner rather than the original entrepreneur—and drive the heuristic to algorithm. These exceptional companies grow to massive size thanks to the efficiency advantage gained over competitors left behind in the heuristic stage. But they too can fall prey to a new competitor that returns to the original mystery and generates a new heuristic – one powerful enough to overcome even the enormous efficiency advantages of the algorithm. They too can be supplanted in due course.

<A> Picking up the Option

The business that remains at one stage in knowledge funnel fails to capitalize on the option created when knowledge is advanced quickly through the knowledge funnel. It misses the opportunity to delve into the next mystery and push that mystery through the knowledge funnel

ahead of the competition. To exploit that opportunity, a company can choose to redeploy the personnel who successfully tackled the last mystery and advanced knowledge along the funnel. By putting these resources to work on new mysteries, the company both defends its current position and goes on the offensive by exploring new opportunities

McDonald's more recent history provides a useful illustration. After its transition from the drive-in heuristic to the quick-service restaurant algorithm, McDonald's grew big and strong exploiting that algorithm with burgers, fries, and shakes. But by the 1990s, it had lost touch with its consumers and what they wanted by way of fast food; its original solution to that mystery had grown stale with time. The company's management was so busy running its algorithm it failed to grasp that many consumers wanted the fast turnaround that is McDonald's byword, but with menu offerings that were healthier or more diverse than pressure-cooked beef, deep-fried potatoes, and sugared milk. Many other chains from Taco Bell to Subway explored the mystery of what those consumers wanted, and their solutions drove McDonald's into a tailspin.

Playing defense is essential because there are multiple paths out of virtually any mystery.

McDonald's chose one route out of the mystery and drove it to an algorithm. But when it settled at that algorithm, it gave its rivals an opening to develop alternative solutions to the mystery. Subway, for example, retained the quick-service component, but replaced burgers and fries with submarine sandwiches and fresh, healthful ingredients.

In doing so, Subway took advantage of the blind spot created on McDonald's path through the knowledge funnel. Remember that as an idea moves through the funnel, information is shaved

away. Some of that seemingly extraneous information can in fact prove crucial to the solution of the next mystery. Early on, McDonald's left health issues by the wayside. Subway made healthy eating the centerpiece of its value proposition, touting its fresh ingredients and low-fat specialties in response to consumers' increasing concerns about unhealthy fast food. McDonald's has subsequently made halting progress toward a healthier menu, but its struggles point to the difficulty that companies have in doubling back along the knowledge funnel.

Other companies can spare themselves similar anguish by using the cost savings generated from pushing their current activities through the knowledge funnel to re-explore the mystery whose initial solution drove its original business model. By reengaging with the mystery and considering information sliced away during the previous trip through the funnel, a company can avoid being blindsided, as McDonald's was by Subway. Only when McDonald's began to explore new approaches to satisfying the consumer's changing desires did it start to climb out of its trough.

The company that gains efficiencies by pushing current knowledge through the funnel also gains an offensive advantage. It can redeploy the savings and redirect its freed-up personnel toward consideration of entirely new mysteries. Proctor & Gamble realized enormous efficiencies by refining its knowledge of household cleaning products. The equity generated by those efficiencies was deployed toward the mystery of baby diapering. The result: Pampers, now one of P&G's biggest businesses. Only 50 years ago it was a complete mystery. In other words, they used the equity realized by becoming more efficient to pursue innovation. They exploited the gains of previous advances to fund their exploration of new mysteries.

<INSERT IMAGE OF ITERATIVE CYCLES OF THE KNOWLEDGE FUNNEL (To Come)>

<A> **Design Thinking and the Design of Business**

Very few companies balance exploration and exploitation by continuously looking back up the knowledge funnel to the next salient mystery (or back to the original mystery) and driving across the knowledge funnel, in a steadily cycling process. These few businesses come to be defined by their balanced approach—they become design-thinking businesses.

Why do so many companies fall into the trap of choosing either exploration or exploitation rather than balancing both? The reason, I believe, is that as companies grow, they become more comfortable with the administration of business. They like and encourage analytical thinking. They embrace a very specific way of arguing and thinking that includes a narrow definition of what constitutes reasonable grounds for moving ahead with a project—a very narrow definition of proof. For analytical thinking, all proof emanates from the past—a general rule handed down from the past, or a set of observations of events or behaviors that have already happened. The average manager has been trained and rewarded to look to the past for proof before making the big decisions.

And to these analytically trained managers, the alternative appears quite frightening: the knowing without reasoning of intuitive thinking. It is no wonder that organizations slowly but surely shift their structures, processes, and cultures to be friendly to only analytical thinking and, without realizing it, to only exploitation of existing knowledge. Their goal is not to drive out

innovation but rather to protect the organization against the randomness of intuitive thinking.

They do not realize that they worship at the altar of **reliability**, a concept that I will return to in Chapter Two. But drive out innovation they do. It is a trap and a pernicious one.

The answer is not to try to get corporations to embrace the randomness of intuitive thinking and eschew analytical thinking entirely. They just won't do it: it is too scary. What's more, analytical thinking is absolutely essential to the exploitation of existing knowledge. No, we cannot do without analytical thinking or intuitive thinking entirely.

The answer lays in embracing a third form of thinking—design thinking—which helps a company both hone and refine within the existing knowledge stage and generate the leap from stage to stage, continuously, in a process I call **the design of business**. I will explore design thinking at length in Chapter Three. For now, let me say that at the heart of design thinking is abductive logic, a concept originated by turn-of-the-20th-century philosopher C. S. Peirce. His important insight was that it is not possible to prove any new thought, concept, or idea *in advance*: all new ideas can be validated only through the unfolding of future events. To advance knowledge, we must turn away from our standard definitions of proof—and from the false certainty of the past—and instead stare into a mystery to ask what could be. The answer, Peirce said, would come through making a “logical leap of the mind” or an “inference to the best explanation” to imagine a heuristic for understanding the mystery.

The McDonald brothers didn't know that their Speedee Service System would work. They had imperfect data, but not irrelevant data. They knew that car-hop restaurants had the appeal of

relatively quick service, but had some drawbacks—loitering toughs and cold food. Their logical leap, their inference, was that patrons liked the basic concept but would like it a lot more if the restaurant were a drive-through with a narrower, more standardized menu.

The brothers had no “proof,” but they did not lack logic. At the time a heuristic is first tentatively proposed, no one can prove whether it is useful or valid at all. Proof comes only if the heuristic is tried and found to be helpful in producing the desired, or valid, result. The same holds for turning the heuristic into an algorithm. Neither of these steps into new knowledge can be proved in advance; all are validated—or not—through the passage of time.

As such, abductive logic sits squarely between the past-data-driven world of analytical thinking and the knowing-without-reasoning world of intuitive thinking. Rather than being confined to regressing the past to hone and refine within the current knowledge stage, the design thinker can add abductive logic to the reasoning repertoire to drive the organization through the knowledge funnel. And rather than being confined to the knowing without reasoning of intuitive thinking, the design thinker uses an explicit form of logic and a process that, while less certain and clear than analytical thinking, has promise for producing advances with greater consistency and replicability than pure intuition.

The design thinker therefore enables the organization to balance exploration and exploitation; invention of business and administration of business; originality and mastery. Design thinking powers the design of business—the directed movement of a business across the knowledge funnel from mystery to heuristic to algorithm and then the utilization of the efficiencies from

doing so to tackle the next mystery and the next and the next. That is the most powerful formula for competitive advantage in the 21st century—the velocity of movement through the knowledge funnel powered by design thinking.

To get there, businesses must acknowledge that they implicitly favor exploitation over exploration, because most businesses, whether they know it or not, favor **reliability** over **validity**. In the next chapter, we will examine those concepts in detail and investigate the forces that converge to reinforce an organizational bias toward reliability. Subsequent chapters will deepen our understanding of the contrasting modes of logic that produce exploration and exploitation, and the implications of both for the way we work. Most firms are dominated by declarative logic, or deductive and inductive reasoning (the logic of what should be or is operative). But new knowledge comes about by way of abductive reasoning, the logic of what might be. I offer a guide to developing the capacity for abductive reasoning—the essential core capacity for design thinkers—in the individual and the organization.

<A> A Different Kind of Organization

That different kind of thinking demands a different way of organizing work. Reliability-oriented firms view it as their overarching task to apply a heuristic or run an algorithm on a continuous basis. So they build up permanent departments staffed by fungible people in permanent slots. They devote the bulk of their energies and resources to rigorous planning and strict budgets. Those processes, which are applications of inductive and deductive logic, drive out initiatives that can't produce near-certain future outcomes.

To balance administration and invention, a business needs to shift the weighting of its structure, processes, and culture. While some aspects of the organization can and should continue to be structured as permanent jobs or tasks, significant parts of the organization should be structured as projects—that is, with teams and processes designed to move knowledge forward a stage—with a definite endpoint. While planning and budget management can't be thrown out the window, they have to be loosened to incorporate initiatives and investments whose outcomes can't be predicted in advance. And culturally, it's imperative that people know it is safe and rewarding to bring forward an abductive argument.

<A> A Different Kind of Leader

Without committed leadership, no business can realize the structural, process, and cultural adjustments needed to become a design-thinking organization. Given the enticing short-term financial rewards of reliability, most organizations will pursue reliability out of simple self-interest. But given those same rewards, validity has a good chance of being squeezed out if someone at the very top of the organization does not champion its value. CEOs must learn to think of themselves as the organization's balancing force—the promoter of both exploitation and exploration; of both administration and invention.

The CEO can perform that function in a number of different but successful ways. Some CEOs, such as Guy Laliberté, founder of Cirque du Soleil, and Mike Lazaridis, founder of Research in Motion, the company that gave the world the Blackberry, take it on themselves to lead the search for innovation. Laliberté spearheads the design of circuses the likes of which the world has never seen before. Lazaridis keeps creating new devices that define the future of mobile

communications. Both keep their organizations moving forward by taking the lead in moving knowledge forward.

At the other end of the spectrum, we find CEOs who build design-friendly organizations. Procter & Gamble's A.G. Lafley is the poster boy for transforming a large, reliability-biased enterprise into a design-friendly organization that maintains a balance between analytical thinking and abductive reasoning. Jim Hackett of Steelcase acquired the design firm IDEO to infuse design thinking across the entire Steelcase organization.

Between the extremes represented by Laliberté and Lazaridis at one end and Hackett and Lafley at the other, there are numerous intermediate alternatives. Steve Jobs, for instance, co-founder and returned CEO of Apple Inc., is probably the CEO most widely viewed as a design thinker, thanks to elegant, customer-pleasing products like the Macintosh, iMac, iPod, and iPhone, among many others. But he is not the solitary design genius of popular imagination. It was Apple's designers, led by Jonathan Ive, who realized those innovative products. Jobs played a different, equally crucial role: He created an organization that placed "insanely great" design at the top of its hierarchy of values and he gave the green light to spend the resources necessary to make lasting successes of his designers' innovations.

<A> A Different You

But what if you are not a CEO? Worse, what if the CEO you work for doesn't have the first clue about the importance of design thinking and the need to encourage it? Are you then powerless to improve your circumstances and your organization's prospects?

Definitely not. You are far from powerless, and you are anything but alone. The business world is dominated by reliability. But that need not impede your efforts to sharpen your own design-thinking capability. While you may have to fight your organization's baked-in biases against design thinking, the effort will help you sharpen your innovation skills and prepare you to be your company's champion of design.

To become a design thinker, you must develop the **stance**, **tools**, and **experiences** that facilitate design thinking. Stance is your view of the world and your role in it. Tools are the models that you use to understand your world and organize your thinking. Experiences are what build and develop your skills and sensitivities over time.

Rather than being cowed by a reliability-oriented world and becoming a prisoner of it, the design thinker develops a stance that puts a priority on seeking validity and making advances in knowledge, even if that stance places the thinker at odds with the organization's culture. In addition to mastering tools for analyzing the past and using that analysis to predict the future, the design thinker develops the capacity for observation, for seeing features that others may miss.

The design thinker, in the words of novelist Saul Bellow, is "a first-class noticer."

Along with developing tools for honing and refining the status quo, design thinkers develop tools for moving knowledge forward. They build their capacity for the unique configuration of designs that transform their insights into viable business offerings. And design thinkers use their experiences to deepen their mastery of the current knowledge domain and exercise originality in

moving knowledge forward to the next stage. In combination, this approach to your stance, tools, and experiences will create a virtuous cycle, reinforcing your design-thinking approach over time.

Toward the end of the book, I will return to advice for the individual design thinker. First, though, let us turn our attention to the organization and to the powerful currents that the design thinker must swim against to transform it. We need to understand how reliability, which first appears to ensure success for any business that cultivates it to its highest point, turns out to be the chief limiter of success. And how validity, which at first seems to be the enemy of reliability, is the force that, when paired with reliability, creates a winning advantage.