

FOREWORD BY

CLAYTON M. CHRISTENSEN

AUTHOR OF *The Innovator's Dilemma*

THE INNOVATOR'S MANIFESTO

DELIBERATE DISRUPTION
FOR TRANSFORMATIONAL GROWTH

Michael E. Raynor

CO-AUTHOR OF *The Innovator's Solution*

FOREWORD

FROM ART TO SCIENCE

by Clayton M. Christensen

In this fascinating book, Michael Raynor tells us that the world of investing to create successful businesses is about to change. Just as theories in the world of biology or physics have allowed us to predictably create desired outcomes in medicine or engineering, Raynor shows here that Disruption promises much greater predictability in the realm of creating successful new businesses. Raynor shows us that there are certain technologies and strategies that succeed much more often than others. He shows us what they are, why they work, and how to apply them. Science—at least in this one instance—truly is making a difference in the practice of management.

The ultimate significance of *The Innovator's Manifesto* will be revealed only over time. I, however, have high hopes for its longevity and impact because Raynor's work falls very neatly into a well-established pattern for the transformation of tacit, intuitive knowledge—art, if you will—into codified, well-understood, explicit rules—in other words, science. I believe that Raynor is playing a central role in transforming the management of innovation from an art to a science. This will truly be a landmark work.

To see the significance of this contribution, consider that in the early stages of any field, our collective knowledge is little more than an assortment of observations collected over many generations.

There are many unknowns, and so the work is complex and intuitive, and the outcomes are relatively unpredictable. Only skilled experts are able to cobble together adequate solutions, and their work proceeds through intuitive trial-and-error experimentation. This type of problem-solving process can be costly and time-consuming, but there is little alternative when our knowledge is still in its infancy.

Creating new, successful innovations still looks very much like this today. Investment decisions and strategic choices are typically based on intuition; learning, if it happens at all, is a very expensive by-product of trial and error. Entrepreneurs and new venture investors alike live a perpetual contradiction, convinced on a case-by-case basis that the venture they have just launched will succeed, even as they cannot escape the fact that 90 percent of all new ventures—including theirs—ultimately fail. In such a world, we can make no clear connection among the attributes of the new business, the oversight provided by the investors, the management methods of the leadership team, and final outcomes. That makes it very hard to learn how to succeed at innovation.

In the face of this uncertainty, some widely accepted rules of thumb have emerged. For example, a mantra for most venture capitalists is that it is folly to make investment decisions based upon the start-up's technology or strategy. The VCs have concluded from their trials and errors that even they—the best in the world—cannot predict in advance whether the technology or strategy described in a start-up's business plan will actually work. As a result, they typically assess—intuitively—whether the management team has the intuition to succeed. If members of the team are experienced and perceptive, the VCs reason, they can develop the right technology and the right strategy—because they and only they will have the instinct to change direction when needed. As far as affecting outcomes in a meaningful and predictable way, however, this approach ranks up there with “feed a cold, starve a fever.” It is little more than an aphorism based on selective memory, the force of repetition, and the hope that at least it does no harm.

Getting beyond myth requires that we first carefully document patterns that repeat over time. This does not provide any guarantee of

success, but it does provide at least some confidence that there is a correlation among factors of interest. Ultimately these patterns of correlation are supplemented with an understanding of causality, which makes the results of given actions much more predictable. Work that was once intuitive and complex becomes routine, and specific rules are eventually developed to handle the steps in the process. Abilities that previously resided in the intuition of a select group of experts ultimately become so explicitly teachable that rules-based work can be performed by people with much less experience and training.

To illustrate, consider the evolution of medical science. At its core, the problem in medicine historically is that the human body has a very limited vocabulary from which it can draw when it needs to declare the presence of disease. Fever, for example, is one of the “words” through which the body declares that something inside isn’t quite right. The fever isn’t the disease, of course. It is a symptomatic manifestation of a variety of possible underlying diseases, which could range from an ear infection to Hodgkin’s lymphoma. Medications that ameliorate the fever don’t cure the disease. And a therapy that addresses one of the diseases that has fever as a symptom (as ampicillin can cure an ear infection) may not adequately cure many of the other diseases that also happen to declare their presences with a fever.

As scientists work to decipher the body’s limited vocabulary, they are teaching us that many of the things we thought were diseases actually are not. They’re symptoms. For example, we have learned that hypertension is like a fever—it is a symptomatic expression of a number of distinctly different diseases. There are many more diseases than the number of physical symptoms that are available, so the diseases end up having to share symptoms. One reason why a therapy that effectively reduces the blood pressure of one patient is ineffective in another may be that they have different diseases that share the same symptom. When we cannot properly diagnose the underlying disease, effective care generally can be provided only through the intuition and experience of highly trained (and expensive) caregivers—medicine’s equivalent of Warren Buffett.

At the other end of the spectrum, we define precision medicine as the provision of care for diseases that can be precisely diagnosed and

for which the underlying causes are understood. This makes it possible to develop a predictably effective therapy. In these circumstances, caregivers such as nurses and technicians can give effective care and at lower cost than is possible today by the best clinicians. Most infectious diseases live here: we have dispositive tests for their presence and well-understood and highly effective treatments for their cure. We can all but guarantee an outcome for an individual; exceptions are rare and noteworthy.

Not all of medicine falls into the “intuitive” or “precision” category, however. There is a broad domain in the middle called empirical medicine. The diagnosis and treatment of a pathology falls into this third category when a field has an incomplete but still very valuable set of causal models and validated patterns. The connections between actions and outcomes are consistent enough that results can be usefully, if imperfectly, predicted. When we read statements like “98 percent of patients whose hernias were repaired with this procedure experienced no recurrence within five years, compared to 90 percent for the other method,” we’re in the realm of empirical medicine. Empirical medicine enables caregivers to follow the odds. They can generally guarantee the probabilistic outcome only for a population.

What makes *The Innovator’s Manifesto* so significant is that it is perhaps the first and in my view the most significant and successful effort yet to move the field of innovation from the intuitive stage into the world of empirical management. Building upon groundbreaking research at Intel Corporation, Raynor has quantified the improvements in predictive accuracy and survival rates that are possible through the careful application of Disruption to early-stage businesses. He has elaborated upon particular elements of Disruption in ways that make clear when and how the theory can be applied. And he has provided frameworks for its application that will enable most any business to reap the benefits that Disruption makes possible.

Achieving such an outcome means that this is not your typical management book. There are no “just-so” stories attributing the success of the latest bottle rocket to a new buzzword. Instead, you will

find the careful collection of real data, considered and circumspect analysis that recognizes shortcomings without being paralyzed by them, a rigorous and reflective treatment of some of the chestnuts of popular management thinking, and a genuine appreciation for the challenges of applying real theory in the real world. You will have to read this book carefully and reflect upon it deeply. But it will be worth it.

As I have said elsewhere, my admiration for Michael Raynor has no end. The integrity of Disruption theory has improved substantially since Michael and I coauthored *The Innovator's Solution*, and much of that improvement I attribute to my continued collaboration with him. I love just to sit in his presence and listen as his magnificent mind goes to work on the complicated puzzles of management. Though I have a busy life, for Michael Raynor I always have time. I hope that you will enjoy being with him as you read this book.

Clayton M. Christensen is the Robert and Jane Cizik Professor of Business Administration at the Harvard Business School in Boston, Massachusetts.

PROLOGUE

THE FIVE- PERCENTAGE- POINT SOLUTION

“Disruption,” used in a technical sense, is a theory of innovation—of how particular types of new products and services, or “solutions,” come to achieve success or dominance in markets, often at the expense of incumbent providers. Disruption was discovered by Clayton Christensen, a professor at the Harvard Business School, in 1992 when he was a doctoral student there. (When using “disruption” or its cognates in a technical sense I will use an uppercase *D*.) Christensen’s 1997 best-selling book, *The Innovator’s Dilemma*, was the first popular expression of his ideas. Christensen and I collaborated on *The Innovator’s Solution*, published in 2003. At least seven more books and hundreds of articles have been published since then exploring the theory’s implications in different contexts.¹ It is in widespread use as an organizing principle for innovation at organizations around the world. Many who have used it have credited it with a significant role in creating successful new businesses.

And yet, thanks to the confusing world of applied management research, Disruption is still seen by many as “just another theory.” One new book after another cascades into the marketplace of ideas, attempting to explain the latest success story or allegedly revolutionary phenomenon with a newly coined term and a fresh set of case studies as supporting evidence. How are practicing managers to de-

cide which frameworks, theories, approaches, or 2x2s are applicable to their circumstances and truly useful to them? How is one to know whether to use Disruption or something else to navigate through the challenges associated with innovating successfully?

EXPLANATION AND PREDICTION

One way to sort out what is useful and accurate from the noise is to take a page from the philosophy of science. In his 2010 book *Nonsense on Stilts*, Massimo Pigliucci points out that the type of evidence one adduces in support of a position depends in large part on the sort of argument one hopes to make.² If, for example, a theory is intended merely to be useful—that is, instrumental in achieving a desired outcome—then one needs to demonstrate predictive accuracy. In other words, theories are useful if they tell us what will happen next, and the most useful theories are simply the ones that do that best.

Assessing predictive accuracy requires very carefully controlled and repeated experiments and at times a remarkably high tolerance for experimental error. Physics, the queen of the hard sciences, has risen to this challenge time and again, and as a result that discipline's long-term project has made enormous progress. We have abandoned theories of phlogiston and the ether for quantum mechanics and the standard model of elementary particles thanks to a careful accumulation of data under increasingly well-controlled conditions. It is a long and complex chain from formulating a theory to controlled experiments testing the theory's propositions to usefulness in the everyday world of middle-sized, middle-distance objects. But every link holds (well enough) for the predictive power of physics to manifest itself in many and repeated successful applications in fields such as engineering.

Predictive power establishes that a theory is useful, but it does not prove that a theory is true; a true theory explains reality. Galileo, for example, would not likely have been in such hot water with the Catholic Church authorities of his day if he had said merely that the

heliocentric view of the solar system was a useful method for *predicting* the future locations of the planets. He got himself in trouble by claiming that it *explained* why the planets moved as they did, namely, because the planets really do orbit the sun and not the earth.

Prediction and explanation require very different sorts of evidence and rules of inference. Experiments to establish predictive power admit of sometimes significant measurement and other sorts of error. Even under the most carefully controlled conditions there remains a great deal that is, well, uncontrolled; indeed, experiments that come out too close to perfect are often suspected of having been fudged. We insist that the theory be specified in advance of the experiments, rather than creating our theory after the fact: our unconscious biases might lead us to create a theory that fits our data perfectly, and since a data set is usually only a sample, this kind of interpolation undermines a theory's broader application. Theories "win" based on the statistical significance of their results over a number of trials and their *parsimony*—their ability to explain the broadest range of outcomes with the fewest and simplest theoretical constructs.

In contrast, explanatory frameworks address a fixed and unchanging past. We cannot test a proposed explanation of what has already happened by turning back the clock and seeing if history plays out the same way again. We must therefore decide what wins based on the completeness of the explanation, the weight of circumstantial evidence, and wherever possible what Pigliucci calls a "smoking gun": one or two critical facts that no other competing theory can plausibly account for.

So, for example, how do we know that an asteroid impact explains the extinction of the dinosaurs sixty-five million years ago? We can reasonably infer from what we know about asteroid impacts in general that an asteroid of sufficient size could trigger a mass extinction. What we need to show is that there *was* an impact by an asteroid of sufficient size at about the right time and that the pattern of extinctions is consistent with the expected consequences. Over the years enough circumstantial evidence has accumulated to convince most informed observers that this was the case. For example, there is a cra-

ter of the right size in the floor of the Gulf of Mexico (which was also an ocean back then), along with evidence of devastating tsunamis along ancient coastlines. We also have a telltale layer of iridium ore of just the right concentration laid down at just the right time in rock strata around the world. Finally, competing theories—such as the rise of egg-eating mammals or climate change due to eccentricities in the earth’s orbit—cannot account for the fact that the dinosaurs were extirpated simultaneously with a great many plant and mammal species as well, nor for the rapidity with which the mass extinctions occurred.

Due to these differences in purpose and hence evidence, establishing explanatory power says nothing about a theory’s predictive power. That the dinosaurs were wiped out by an asteroid implies little about what will cause the next mass extinction. It just turns out that an asteroid strike caused that one.

Consider now the last management book you read. What kind of evidence did it provide in support of its central claims? It very likely relied for evidence on an analysis of case studies, and out of that analysis emerged a framework purporting to explain why events turned out as they did—why a given company succeeded or failed or why a given product was a hit or a flop.

Very often, however, the explicit claim is that the principles that have been extracted from an analysis of the past can be used to shape future outcomes in desired ways. Typically, authors seem to believe that case-study evidence alone supports prescriptive claims. In other words, most every management book I am familiar with—and certainly most of the best sellers—makes *predictive* claims based on *explanatory* power. Whether deliberate or not, it is a most unfortunate and potentially damaging form of conceptual bait and switch.

Is there any way to avoid this, though? After all, the subject matter of management research—actual organizations functioning in the real world—does not lend itself to the kinds of carefully controlled experiments that allow us to test predictive accuracy in the usual ways. Perhaps we can do no better than simply to infer predictive power on the basis of explanatory persuasiveness.

THREE OBJECTIVES

I disagree. The first objective of this book is to demonstrate that Disruption has true *predictive power*. I hope to show this using what is for many people the most persuasive evidence there is when it comes to prediction: controlled experiments. My hope is that you will find these data sufficiently compelling to conclude that Disruption is unique in having evidence to support the claim that it is genuinely *useful*.

Second, I will make the case for Disruption's unique and superior *explanatory power*. I will lay out a definition of Disruption precise enough that Disruptive innovations can be accurately identified in advance of knowing how they ultimately fare and their results in the marketplace explained more fully and parsimoniously than by any other theory. To the extent I succeed, I hope you will conclude that Disruption is far more than merely a useful perspective but is in fact *true*.

Finally, I will offer some thoughts on how one can go about *applying* these concepts to greatest effect at the least expense. To the extent this third objective is achieved, I hope you will conclude that Disruption is *practical*.

And if I can convince you that Disruption is useful, true, and practical, I will go further and hope that you will want and be able to use it in support of your innovation efforts.

PREDICTION: CHAPTERS 1 AND 2

Chapters 1 and 2 describe the design and results of carefully controlled experiments testing the predictive power of Disruption's central claims: that an innovation has the best chance of success when it has a very different performance profile and appeals to customers of relatively little interest to dominant incumbents, and the organization commercializing it enjoys substantial strategic and operational autonomy. In contrast, attempts to introduce better-performing solutions targeted at customers valued by successful incumbents will fail.

To test these propositions I use a portfolio of forty-eight new business proposals that received seed financing from Intel Corporation.

To summarize the results, test subjects improved their predictive accuracy by as much as 50 percent when they applied Disruption theory to make their choices. Specifically, in the actual portfolio of funded businesses just over 10 percent survived. The portfolio chosen by MBA students who did not use Disruption theory had a similar survival rate, while students using Disruption theory to pick winners built a portfolio with a survival rate of up to slightly more than 15 percent. That five-percentage-point gain is a 50 percent improvement. (More recently, Intel reports that the survival rate of its funded businesses has increased, in part due to the application of Disruption theory.)

Of course, neither the data nor the experimental design is perfect (and I will have more to say about the precise nature of the imperfections of this work later on), but perfection is the wrong benchmark. In the mortal realm, all success is relative, and the most important question is not “What are the flaws of this design and these data?” but “Are this design and these data better than what you have seen elsewhere?”

Note also that I am not claiming that I have shown that Disruption theory is better than some other theory. Rather, I am claiming that the evidence in support of Disruption theory’s predictive power is better than the evidence supporting any other relevant theory’s predictive power.

To see the difference between these two claims, consider tests for the efficacy of new pharmaceutical drugs. Imagine that Disruption is a drug that purports to treat a given condition, and some other theory is a different drug making the same claim. The evidence in these first two chapters supports the claim that Disruption actually “treats the condition”: it improves predictive accuracy. I have not shown that Disruption works better than any other drug; that requires comparing the relative effectiveness of two drugs. At the same time, however, as far as I know no one has shown that any other drug actually treats the condition at all.

What I hope to convince you of at the outset, then, is that Disruption can claim more legitimately than any other theory to make you better than you are with respect to one critically important decision: assessing which businesses will live or die.

EXPLANATION: CHAPTERS 3 TO 5

A common challenge in research of any kind, and certainly in the field of applied management, is determining the extent to which one can “generalize beyond the sample.” For example, if someone does a study on large public companies, do the findings apply to small, privately held, family-run businesses?

To extend our pharmaceutical drug testing analogy, consider clinical trials on a drug that treats high blood pressure. Such trials typically include thousands of people and years of observation in order to determine whether a new drug is safe (does no harm) and effective (actually helps in the desired way). Assume for the sake of argument that the drug proved safe and effective, but it turned out that there were no subjects named Phil. Administering the drug to people named Phil with the expectation of safe and effective outcomes is generalizing beyond the sample. One is therefore open to the possibility that the drug could have a different effect on people named Phil than it did on those observed in the study.

Thankfully, we can claim a credible understanding of what will happen in circumstances we have not tested directly if we have a correct understanding of why results turn out as they do. In the pharmaceutical example, if we understand the mechanisms of action for a particular drug and we know with a high degree of certainty that being named Phil has no material impact on a drug’s effect, then we are justified in generalizing beyond the sample. If, however, there are other attributes that we believe might affect the drug’s efficacy—say, a patient’s sex or age or being diabetic—in ways that we do not fully understand, then we are not justified in generalizing beyond the sample.

In reality, as is often the case, such judgments are not binary: one

is more or less justified in generalizing beyond the sample depending on the sample, what one hopes to generalize, and how far beyond the sample one wishes to go. In the large public/small private company example, we might ask what the relationships are between behaviors and outcomes being investigated and whether there are meaningful differences between these types of companies that might affect the relationships we observe in our sample. A study about processes for implementing a quality-management system might generalize across such diverse companies much better than a study on governance processes, for example, since the public or private structure of a company has a direct bearing on the relevant legal and regulatory governance requirements.

With this in mind, the extent to which we can reasonably expect the predictive power of Disruption to be evident in contexts that were not directly tested in the experiments turns on whether Disruption can account for its predictive power by specifying when it should be applied and providing sufficiently powerful and compelling explanations for why it works. In other words, the generalizability of demonstrated predictive power is a function of explanatory power.

The experiments in chapters 1 and 2 test whether Disruption improved the ability of MBA students to predict the survival of very early-stage business plans. Chapters 3 through 5 explore the extent to which other types of people in different circumstances can do anything with these findings by making the case for Disruption's explanatory power. Unlike the tests of predictive power, this entails a direct comparison of the explanatory power of Disruption with the explanatory power of competing theories when accounting for specific outcomes.

The test case, explored in chapter 3, is Southwest Airlines, for although Southwest has been analyzed seemingly ad nauseam, the signal feature of Southwest's performance—its nearly twenty-year run of slow growth and declining profitability from the early seventies to the early nineties, with a sharp turnaround and a decade of record-setting growth, increasing profitability, and share-price appreciation—has had no parsimonious explanation. Disruption, however, explains not

merely why Southwest was successful but also why its growth occurred precisely when it did. I will argue that Disruption explains the salient features of Southwest's performance in a way that no other theory does, and in a way that would have made it possible to predict Southwest's success. This is the sort of "smoking gun" required to establish that Disruption is the *right* explanation, rather than merely a plausible one.

Now, proving that Southwest was a Disruptor says nothing about any other company. Nor am I claiming that every successful innovation is a Disruptive one. So chapter 4 describes how to determine whether or not a given opportunity has even the potential to be Disruptive. For example, I explain how so far the hotel industry, strategy consulting, and the discovery of new patentable pharmaceuticals have been immune to Disruptive innovation, not (to use a phrase you will see repeatedly) as a matter of theoretical necessity but merely as a matter of empirical fact. The key message here is that an integral part of Disruption theory is the criteria for determining when it is applicable.

Having defined the circumstances under which Disruption is possible, chapter 5 addresses how to assess the timing and extent of Disruption. For example, why did Disruption take so long in the automotive sector (Toyota's rise to global leadership took almost seventy years) and so quickly in telecommunications equipment (Cisco was an industry leader less than fifteen years after going public). Chapter 5 explains why these Disruptions played out as they did.

This second section makes the case for generalizing beyond the experimental sample and suggests that Disruption can be used to do more than merely "pick a winner." For example, thanks to its combination of predictive and explanatory power, Disruption can be applied:

- **If you are an investor:** to pick with greater accuracy which businesses have the best chance of survival. This is the most direct application of the experimental results.
- **If you are an entrepreneur:** to shape your ideas and your strategy so that your new businesses have a better chance of surviving, getting additional

funding, and ultimately thriving. Since looking at a new venture from the perspective of the entrepreneur is just the other end of the situation faced by the investor, this is perhaps the most direct extension of Disruption's applicability. In short, if you understand what makes a company successful from an investor's point of view, you have a better shot at building a business with those characteristics.

- **If you are a manager trying to grow an existing business:** to improve materially your ability to identify or create opportunities to innovate successfully. What makes Disruptive innovations successful is their trajectory of performance improvement: the ways and rate at which a product or service gets better. It is because Disruption allows you to assess and determine these variables that it makes for better investment decisions. Consequently, if you want to improve your chances of success in an existing business, Disruption prescribes that you guide your own innovation efforts in ways that make you Disruptive to others whenever possible.
- **If you are in corporate M&A:** to identify viable targets and manage them in ways that are likelier to create value. Although materially different in important ways from launching a new business from within an established company or piloting a going concern, acquiring an existing firm demands that you think carefully about the strategy you hope to advance with the acquisition. Disruption theory provides a way to think about this problem, with important implications for how to manage the integration process in particular.

At the same time, Disruption is not a theory of everything. There are lots of other questions you will have to answer no matter which of these roles you fill. For example, as an investor, you likely have to worry about the risk/return structure of your overall portfolio. If you are an entrepreneur you likely have to worry about how to raise capital. If you are managing an existing business, you probably have to worry about organizational politics and the challenges of head-to-head competition in your core markets. If you are in corporate M&A, you likely have to worry about how best to finance the deal and realize cost synergies. These are important questions, but Disruption does not bear directly on them. What Disruption *can* do is materially and significantly contribute to your overall likelihood of success.

APPLICATION: CHAPTERS 6 TO 8

Whatever scientific rigor and theoretical elegance might characterize Disruption, the proof of the pudding is in the eating. And so how to apply Disruption successfully is addressed next.

Chapter 6 is an exploration of how Disruption can be used to shape a specific product innovation. We follow the evolution of what is now Johnson & Johnson's SEDASYS™ automated sedation system from an early-stage partial equity stake in a small start-up to a commercialized product aimed at revolutionizing a wide and increasing range of surgical procedures the world over.

It is a fact that non-Disruptive innovations can succeed and that breakthroughs by new entrants sometimes revolutionize an industry—something that Disruption theory cannot account for. Consequently, chapter 7 explores the implications of deliberately pursuing this sort of unexpected (to Disruption theory, at least) success for specific management processes. Highlighting the key success factors, probability of success, magnitude of initial investment, time horizon, requisite autonomy, and connections to the established business for each type of success should be helpful when deciding how much to invest in different types of innovation. In other words, where chapter 6 explores how to use Disruption to shape a single project, chapter 7 looks at how Disruption might fit into a broader portfolio of innovations.

Finally, chapter 8 takes a process perspective on the application of Disruption. Is Disruption a theory that can be plugged into existing ways of thinking about and fostering innovation, or is a fundamental shift in mind-set required to make the most of what Disruption implies? The claim here is that the existing paradigm of innovation is evolutionary (variation, selection, retention) and, despite the exhortation to “fail fast,” is unavoidably profligate. Disruption admits of a different tack: begin with a clear focus on areas ripe for Disruption; shape ideas so that they are consistent with the prescriptions of the theory; and persist in the pursuit of a Disruptive strategy, learning and adapting along the way.

The examples and tools in these chapters are intended to start you—whether you are an investor, an entrepreneur, a manager, or a corporate M&A strategist—down the road to using Disruption effectively.

HOW MUCH IS ENOUGH?

The MBA students in the experiments improved their population-level predictive accuracy by up to five percentage points. That does not mean you can expect to do the same.³ What do these results mean for you, then?

There are at least two questions worth asking yourself as you answer this question. First, is the evidence I provide sufficient to support my conclusion? I have attempted to make my case for Disruption's predictive and explanatory power with as full an accounting of its shortcomings as I am able to provide. You might find still other flaws. I would encourage you, however, to assess the significance of these shortcomings in light of the evidence supporting claims made by other investigators or, for that matter, your current views about innovation. Without some sort of critical parity there is a danger that one will end up holding on to existing beliefs not because they are better supported but only because they are existing beliefs. Consequently, whether you personally should accept the claims made here and add Disruption to your arsenal of ideas depends not on the objective merits of my case but on how well my evidence and my argument compare with the foundations of competing views.

Second, even if you believe my findings, are they meaningful? After all, a bump in the survival rate of a portfolio from 10 percent to up to 15 percent across a population is no guarantee of riches for you, personally, on your next endeavor. If I could credibly make such a promise, I would not sell you the knowledge. But five percentage points is still a 50 percent increase over the baseline survival rate of 10 percent.

Putting those five percentage points in a broader context, it is

worth remembering that even physics—so impressive in its predictive and explanatory power—is a long way from having everything figured out. In addition to the long-standing difficulties of reconciling quantum mechanics and general relativity, current thinking is that we actually do not understand what the universe is made of.

Galaxies are rotating so fast that the gravitational force of the stars within them is insufficient to keep those galaxies from flying apart. To account for their coherence, physicists have invoked the notion of “dark matter,” which is really just a label for whatever it is that is generating the additional gravitational force unaccounted for by the mass of the stars. At the same time, the universe is expanding, not contracting, which is what it should be doing thanks to all that dark matter that is supposedly out there. So to counteract the effects of the dark matter, cosmologists have ginned up “dark energy,” which is whatever is overcoming the dark matter and pushing the universe outward.

When you put it all together, according to current estimates, the universe is made up of 24 percent dark matter (whatever that is), 72 percent dark energy (whatever that is), and only 4 percent matter—the bit we actually think we understand, putting aside the schism between quantum theory and general relativity, of course.⁴ And yet, with our arms barely around barely 4 percent of the universe, look what we have been able to accomplish.

Maybe five percentage points is pretty good, after all.

PART I

PREDICTION

CHAPTER ONE

A PROBLEM OF PREDICTION

If the purpose of a theory is to inform our choices today, we must demand more than compelling explanations of the past. For a theory to have a legitimate claim on our allegiance there must be evidence that it improves our ability to predict future outcomes.

Creating and backing winning businesses is by all accounts a low-probability endeavor. Far more new businesses fail, or at least do little better than limp along mired in mediocrity, than actually break away from the pack and create real wealth. There is more to this statement than simply the necessary truth that only 10 percent of all businesses can be in the top 10 percent: the best businesses tend to do fabulously well, while most of the rest, if they survive at all, generate returns that are embarrassingly small in comparison.⁵ We have become collectively resigned, it seems, to the notion that successful innovation is unavoidably unpredictable.

Despite the challenges and the long odds, there is no shortage of players in this great game. Hedge funds and venture capital partnerships channel capital into the businesses they feel will succeed. Many corporations maintain internal venture functions for strategic purposes, some seeking to create ecosystems around a core business or to stake a claim to possible new growth opportunities in adjacent

markets or to establish a line of defense against possible usurpers of a valuable entrenched position, to name only three possible objectives.

Take, for example, Intel Corporation, best known for its significant role over the last thirty years in the global microprocessor industry. In 1998 Intel launched the New Business Group (NBG) in order to coordinate and more effectively manage the company's attempts to diversify beyond the microprocessor industry.⁶ Within NBG, approximately \$20 million was earmarked for the New Business Initiatives (NBI) group, which had the remit to identify, fund, and develop new businesses that were especially far afield, such as Internet-based businesses and consumer products. NBI's mandate included exploring new technologies, new products, new markets, and new distribution channels and had an investment horizon of five to ten years.

NBI operated as a largely autonomous unit within NBG. Unlike the relatively formal and structured annual planning and budgeting processes that drove sustained success in the microprocessor segment, NBI typically committed only seed capital to new business ventures, ramping up its level of commitment as various strategic and financial mileposts were reached. In addition, leadership explicitly accepted the inherent unpredictability of incubating new businesses along with an unavoidable implication of that uncertainty: that some and perhaps many of the ventures that were launched could fail.

Intel Optical Links (IOL) was one of NBI's investments. Thomas Thurston, then an attorney in his midtwenties with an MBA and law degree, joined IOL in 2005, excited at the prospect of helping launch a new venture inside an established company. Although successfully incubated, IOL was sold off following Intel's broader divestiture of optical component and communications businesses. However, Thurston's curiosity was piqued by this initial exposure to the internal venturing process: he wanted to understand better how Intel decided which initiatives to support and why.

Something in excess of seventy business proposals are explored by NBI's investment directors each year. They work with a range of people and sources, both inside and outside Intel, to determine

the potential of a given idea. The constant challenge is to find the “diamonds in the rough”—the concepts that have within them the seeds of sustainable success and perhaps greatness. It is an inherently risky undertaking, and the only way to avoid failure entirely is to do nothing, which of course reduces one’s chance of success to zero as well.

It is this unavoidable uncertainty that leads many observers to prescribe an investment strategy based on “rapid failure”: the willingness to attempt as many different initiatives as possible with an eye to learning what does not work as the inevitable prerequisite to discovering what does. In Intel’s world, however, bone fide initiatives—the kinds of efforts that actually teach you something useful—can get very expensive very quickly. NBI executives are therefore forced to make difficult trade-offs between the need to husband their investment capital and the risk of overlooking the next blockbuster product or service.

For present purposes, the salient features of NBI’s investment process were the Seed Approval Meeting (SAM) and Business Approval Meeting (BAM). Proposals that were approved at the SAM received funding of several hundred thousand dollars to typically less than \$1 million, with an upper range that rarely exceeded \$2 million. This allowed a team to get beyond the idea stage and flesh out a business plan, perhaps by developing a prototype, collaborating with potential customers, doing market research, and so on. BAM funding was contingent on having demonstrated an increased level of viability and brought with it investment capital that ranged from several million dollars to in some cases as much as \$20 million. Ultimately, NBI’s goal was to transition or graduate one new business opportunity per year to an existing or new business unit within Intel. (Not every venture had to pass through both stages of approval: some ventures were graduated directly from SAM to an operating division in light of their strong performance.)

Intel takes a very rigorous approach to understanding competitors, technology, customers, market structure, and a host of other variables when analyzing opportunities for growth. Unfortunately for Intel, and

everyone else who seeks to innovate in order to grow, there are no data about the future, and so there often remained many important but unanswered questions. Consequently, well-informed, experienced executives could look at the same opportunity and come to different conclusions about that venture's challenges, financial potential, and so on. Worse, only when a venture was funded could the merits of the decision-making process employed be assessed, since if something was turned down, it rarely got funded via other channels, and so the opportunity cost of passing on what would have been a winner was almost always incalculable.

Thurston undertook a forced march through the popular management research into innovation in search of a more nearly rules-based approach in the belief that, given the importance of the subject and the wealth at stake, any framework holding even a scintilla of advantage over the others would be readily identified. Yet Thurston discovered that instead of a vibrant marketplace of ideas populated by challengers seeking to unseat the reigning champion, the agora where theoretical dominance is established is characterized by general disarray. There were a great many frameworks supported by compelling evidence, yet when they conflicted and counseled different courses of action, there was little basis in the evidence to guide someone in choosing one approach over the others. When different approaches did not conflict, it was difficult to treat them as cumulative and attempt to follow the sum total of their collective advice, since doing so resulted in a paralyzingly long to-do list.⁷

In light of this theoretical cacophony, in all likelihood NBI executives made their choices in largely the same way most early-stage investors make their choices: do the best you can with the data you have available, while necessarily relying on your experience and your wits to fill in the sometimes significant gaps. The very best practitioners typically do all they can to create a solid fact base, but personal judgment generally figures prominently in making the final choice.⁸ It is simply the nature of the beast that evaluation criteria differ from person to person and project to project. Thurston recounts that at NBI, this meant that sometimes the emphasis was on technology,

sometimes on management expertise, sometimes on the promise of the market opportunity, sometimes on the strength of linkages with Intel's core business. It is a process that seems to have served Intel well, for there is no reason to think that its achievements are anything other than representative of the very best efforts in this space.

The prevalence of this sort of approach is an understandable consequence of the reliance of popular management research into innovation on post hoc case-study evidence to support its claims. What Thurston was looking for was evidence supporting *predictive accuracy* in addition to the requisite *explanatory power*. And no theory he could find provided both.

CLOSE, BUT NO CIGAR

Christensen's first book, *The Innovator's Dilemma*, introduced the world to the notion of "disruptive technology." Christensen described how large, successful incumbent organizations in all types of industries were toppled by much smaller start-ups. Entrants typically succeeded by developing solutions for relatively small and unattractive markets that were of essentially no interest to successful incumbents. These constituted the entrants' "foothold" markets. Sometimes customers in these foothold markets were quite happy with inferior but much less expensive solutions; sometimes they required solutions with a vastly different performance profile. Either way, entrenched players, focused on the needs of their established customers, proved systemically unable to devote investment funds to those markets. In contrast, driven by their desire to grow, the entrants were strongly motivated to improve their initial offerings in ways that would allow them to compete effectively for the larger, more lucrative mainstream markets. This was the entrants' "upmarket march," and entrants that marched upmarket successfully eventually captured the customers that had been the incumbents' lifeblood.

Christensen observed that when entrants attacked successful incumbents by adopting the incumbents' models and technological so-

lutions, they tended to fail. They tended to succeed by combining a business model suitable for a relatively less attractive market—the entrants' foothold—with an ability to improve their original solutions in ways that allowed them to provide superior performance in a manner incumbents were unable to replicate—the upmarket march. Christensen called the union of these two elements a *disruptive strategy*.

The archetypal illustration of this phenomenon is Christensen's all-inclusive study of innovation and competition in the U.S. disk drive industry from 1976 to 1994. In the midseventies, companies such as Storage Tech and Control Data were making fourteen-inch disk drives for mainframe computer makers. These companies, among them Amdahl and Unisys, wanted Storage Tech and Control Data to innovate: greater storage capacity, faster data-retrieval times, and lower costs per megabyte.

When minicomputers were first brought to market by start-ups such as Sun Microsystems and Hewlett-Packard, they required very different disk drives: smaller, more modular, and less expensive. To achieve these outcomes, disk-drive makers found they would have to reduce storage capacity, increase data-retrieval times, and accept higher costs per megabyte. The result, the eight-inch disk drive, was close to the antithesis of what Storage Tech and Control Data would countenance as an innovation; it was, if anything, a technological step backward in the interest of serving a small and highly uncertain new market. That opened the door for start-up drive makers such as Micropolis and Maxtor to develop something that was technologically trivial to Storage Tech and Control Data but strategically impossible for them to launch.

In the short run, no harm done: Storage Tech and Control Data went on printing money in the fourteen-inch disk-drive market while Micropolis and Maxtor eked out a living selling technically inferior eight-inch disk drives to small minicomputer makers. But then Kryder's law—the disk-drive equivalent of Moore's law in microprocessors—asserted itself: the areal density of disk-drive storage space was doubling annually thanks to improvements in record-

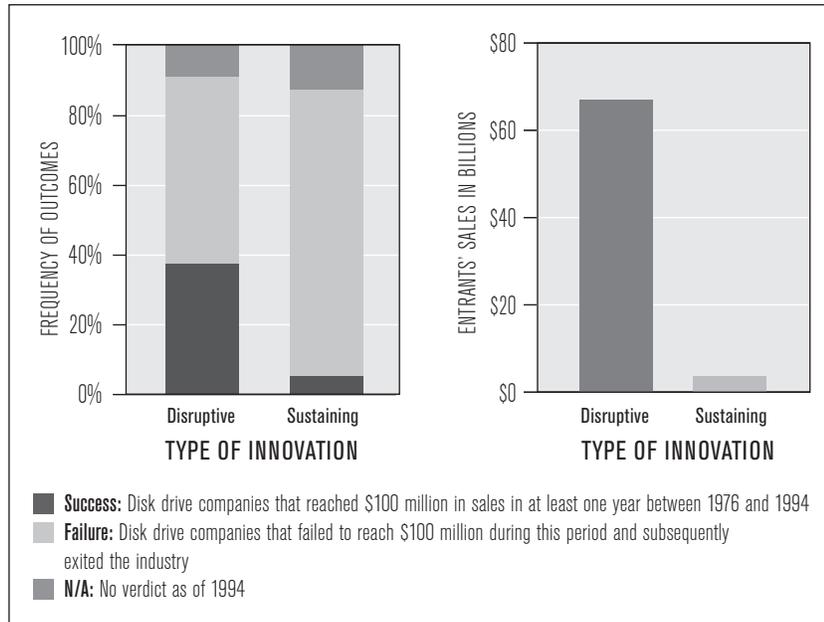
ing media, software correction codes, and other key technologies. In addition, other dimensions of minicomputer performance were improving rapidly, fueled in large part by advances in microprocessor technology and software design. As minicomputers began to encroach on the mainframe market, and ultimately pushed mainframes into decline, the fourteen-inch disk drive makers cast about for new markets but found only the minicomputer makers buying, and they wanted eight-inch drives. Thanks to their relative unfamiliarity with the innovations first commercialized by the eight-inch disk drive makers (e.g., greater modularity and smaller size), the companies making fourteen-inch disk drives were at an insuperable disadvantage. Most went out of business, and none was able to maintain its market dominance in the disk-drive industry.

The start-up eight-inch disk drive makers found a foothold by first *exploiting* trade-offs among different dimensions of performance and appealing to the needs of an economically unattractive market. They disrupted the fourteen-inch disk drive makers by ultimately breaking those trade-offs and remaining the primary disk drive suppliers to the newly dominant minicomputer companies. In other words, as the most lucrative and largest end customers for computers switched from mainframes to minis, the fourteen-inch disk drive makers ended up going down with their chip. (Sorry.)

Accept for the moment that Disruption is a good explanation for a specific phenomenon: the seemingly unlikely ability of entrants to topple well-resourced and well-managed incumbents on their home turf. Still more remarkably, however, Christensen observed that over the eighteen years of competition in disk drives that he documented, Disruptive strategies had a much higher frequency of success, and when successful were much *more* successful than sustaining strategies.

On the strength of this, Thurston felt that Disruption was among the most promising of the frameworks he had studied. He was particularly encouraged by the fact that Disruption lent itself to fairly straightforward predictions of what would work and what would not.

FIGURE 1: THE FREQUENCY OF SUCCESS OF DISRUPTIVE AND SUSTAINING STRATEGIES



Sources: *The Innovator's Dilemma*, p. 145; *The Innovator's Solution*, p. 43

And then Thurston ran into a brick wall. There were no data to support any claims of predictive accuracy for Disruption. Christensen and others had developed a robust library of literally hundreds of cases across dozens of industries that were explained by Disruption—but the same was true of many other theories out there. Worse, for just about every case study explained by Disruption there were competing explanations that drew on entirely different sets of concepts. (Academic journals continue to debate whether Disruption is the best explanation of the disk-drive industry's evolution.) And even if it were possible to win the battle for explanatory-power bragging rights, until there was some evidence in support of Disruption's predictive power it could not claim to be the right theory to use for making decisions about the future. Thurston could have no more con-

fidence in the prescriptions of Disruption than he could in any other theory.

EVERYONE COMPLAINS ABOUT THE WEATHER

Intel has worked with Christensen for some years, and the company has used Disruption theory in its own strategic planning processes. In fact, Christensen and former Intel CEO Andy Grove appeared together on the cover of *Forbes* magazine in January 1999 under the headline “Andy Grove’s Big Thinker.” Consequently, when Thurston approached NBI’s leadership about exploring whether or not Disruption might have predictive power when applied to NBI’s portfolio of investments, divisional leadership provided Thurston the latitude and support necessary to conduct some preliminary investigations.

Thurston began by stating Disruption’s predictions. Specifically, Disruptive innovations are defined as products or services that appeal to markets or market segments that are economically unattractive to incumbents, typically because the solution is “worse” from the perspective of mainstream, profitable markets or market segments. Disruption predicts that leading incumbents with so-called sustaining innovations—innovations targeted at their most important customers—typically succeed. New entrants with sustaining innovations typically fail.

Disruptions typically succeed, whether launched by incumbents or entrants, but only when the ventures launching them are highly autonomous and able to design strategic planning processes and control systems and financial metrics, among other characteristics, independently of systems built for incumbent organizations. This element is important and hardly unique to Disruption: established, successful businesses can and should be held to very different measures of performance and expectations for future performance than start-up organizations, and for at least two reasons. First, a start-up typically has a trajectory of growth and profitability that is very different from that of an established business. Second, start-ups typically must

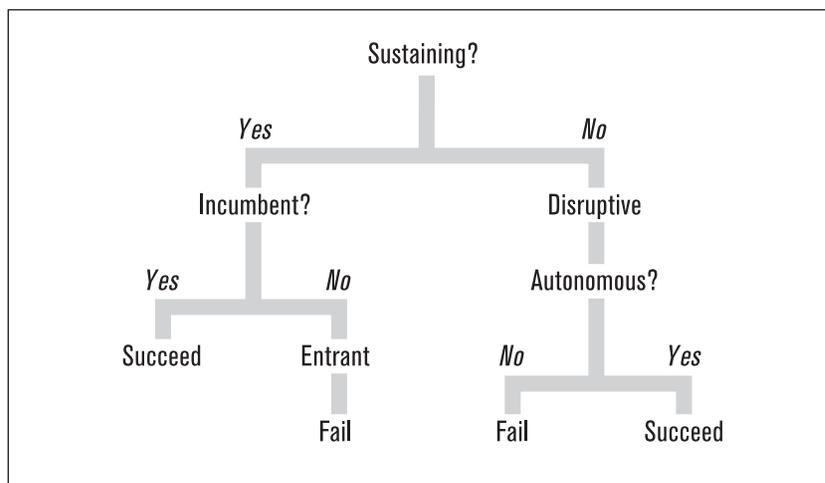
change, sometimes dramatically, material elements of their strategy as they grapple with the unpredictable nature of customer reaction, competitive response, and the performance of key technologies. Consequently, start-ups must find their own way, and that is possible only when they enjoy the requisite autonomy to do so.

In short, Thurston inferred that Disruption predicts that success awaits sustaining initiatives launched by successful incumbent organizations and Disruptive initiatives launched by autonomous organizations. Everything else is predicted to fail. (See figure 2 for a summary of Thurston's hypotheses.)

Now Thurston needed data with which to test those predictions. Fortunately, NBI had retained a robust archive of the materials supporting many of its previous efforts. This allowed Thurston to compile a portfolio of forty-eight ventures that had received at least SAM-level funding over the ten-year period ending in 2007. SAM funding, recall, was very early-stage support, analogous perhaps to "angel" investing. Using the "pitch decks" that were used to explain each business to NBI executives as part of its funding process, Thurston assessed these SAM-approved businesses for "incumbent" or "entrant" status based on the degree of Intel's participation in the market targeted by the start-up and assessed the start-up's product or service as sustaining or Disruptive based on how it compared to existing solutions in that targeted market.

These decks were typically exemplars of business planning and communication. They began with a summary of the technology involved and the benefits to Intel of commercializing it. The most optimistic projections were usually for devices or services that were demonstrably superior to existing solutions offered by competitors. The growth opportunity was often argued to be greatest when Intel did not already compete in that market.

A review of the management team's expertise then followed. It was not uncommon for ventures to be run by an impressive cross section of Intel veterans, new hires with experience in the target market, and others with deep expertise in functions such as marketing or design, depending on what was seen as critical to long-term success.

FIGURE 2: THURSTON'S HYPOTHESES

In framing the predictions implied by Disruption in this way, Thurston was emphasizing two elements of Disruptors: they start out targeting markets or market segments that incumbents do not value, and they have significant autonomy. But he ignored one other element that will prove crucial: Disruptors must improve in ways that allow them to compete for mainstream markets from a position of structural advantage. That is, it is not enough simply to appeal to a market or market segment that is unattractive to incumbents; that is a niche strategy. We will tie off this loose end at the conclusion of chapter 4.

For now, focus on what Thurston was trying to get done: he was looking for actionable advice that would help him predict whether a start-up would succeed or fail, and Disruption—as he interpreted it—provided the kinds of predictive, falsifiable statements that he could test.

Then came a detailed description of the value proposition. This was the team “making good” on its claims of superiority, often including endorsements of prototypes by customers the team was targeting as early adopters. This was followed by an implementation plan: which market segments would be targeted in what sequence, with specific descriptions of how Intel would be successful in each, often accompanied by a multigenerational product road map. Finally, financial projections, complete with sensitivity analysis, described the anticipated economic value of the business to Intel, usually over three to five years.

To keep things as simple as possible, he defined “success” as survival—that is, the venture was still functioning as a going-concern venture, whether or not it was still controlled by Intel—and “failure” as “dead”—that is, no longer a commercial going concern. Without knowing the actual outcomes for these ventures, if Thurston could assess the relevant characteristics of the NBI-backed ventures and predict subsequent “success” and “failure” more accurately than chance alone, he would have solid evidence supporting Disruption’s predictive power.

Here is how it worked with Image Illusions, a disguised NBI-backed venture. Image-processing technologies, such as printers or photocopiers, typically use a large number of application-specific integrated circuits (ASICs) to handle different elements of image manipulation, such as shrinking or rotating an image, prior to printing. ASICs are very efficient, but this efficiency brings with it two drawbacks. First, because each ASIC is highly customized, manufacturing economies of scale are limited, which keeps costs up. Second, ASICs are not programmable, so changing the features of a product typically requires designing and sourcing an entirely new chip, which is costly and slows down development times.

Alternatives to ASICs, such as media processors, digital signal processors, and central processing units, provided vastly increased economies of scale and programmability but sacrificed performance to such an extent that they were rarely viable. In other words, there was a sharp trade-off among performance, flexibility, and cost. Manu-

facturers of image-processing technology—for example, the folks who make printers and photocopiers—would find it very valuable to break that trade-off, for then they could introduce a greater range of more powerful new products faster and at lower cost.

Intel is an incumbent in one of these three alternative technologies mentioned above. Image Illusions sought to leverage this position to create a new solution that provided both efficiency and flexibility. By competing with ASICs, Image Illusions would be leveraging one of Intel's core competencies to expand into a “white space” opportunity to generate new, innovation-driven growth.

In collaboration with a key potential customer—a large, successful manufacturer of digital imaging technology—the Image Illusions team developed a highly sophisticated and demonstrably superior solution based on proprietary intellectual capital. It cost almost twice as much per unit as ASICs, but the team felt (and the customer corroborated) that the higher price was more than offset by the increased performance and flexibility. In other words, the team had broken the critical trade-off that was limiting the performance, cost, and pace of innovation in image-processing technology.

There were, of course, challenges. The largest companies that made image processors—including the one that Image Illusions had collaborated with and all of the targeted early adopters—had their own in-house ASICs design staffs. Many of these people were also on the internal committees that assessed new technologies. To adopt a non-ASICs solution was effectively to put themselves out of a job. That meant Image Illusions would likely have to be *vastly* superior before customers would switch in volume, since the in-house ASICs design teams would be strongly motivated to show that they could up their game and match the new technology.

The Image Illusions team had reason for optimism. The image-processing market was fiercely competitive, and the vast performance improvements Image Illusions could provide meant that all the team needed was one major player to adopt its solution and the rest would follow suit. The ability to leverage Intel's strong brand and customer access made the odds of getting one domino to fall seem

very favorable. The cash-flow projections for Image Illusions estimated a net present value (NPV) between \$9 million and \$100 million over five years, a range that reflected both the team's confidence and the unavoidable uncertainty that comes with launching a new business.

Assessing the prospects of such a venture is reasonably seen as a complex and challenging task. Is the technology *really* that much better? Is it "better enough" to overcome the entrenched interests of the customers' in-house design functions? Is the management team at Image Illusions up to the challenge of overcoming the inevitable and unforeseeable twists and turns on the road to success? Is Intel sufficiently committed to this venture to support it for the one, two, or three years needed to make it to positive cash flow? It would appear that to predict with any confidence what will happen one must have deep experience and expertise in the relevant technologies and markets, strong familiarity with the management processes at Intel, and an intuitive but accurate take on the abilities of the leadership team.

Not if you are Thomas Thurston trying to test the predictive accuracy of Disruption. For him, the only questions that mattered were the following:

- 1. Is Intel an incumbent in this market; that is, does Intel already sell this sort of product to this sort of customer?*
- 2. Is Intel's innovation sustaining or Disruptive in nature? A Disruptive solution makes materially different trade-offs than the existing solutions purchased by mainstream customers; a sustaining solution is straightforwardly better.*
- 3. If the innovation is Disruptive, does the new business launching it enjoy operational and strategic autonomy from Intel's established processes?*

In the Image Illusions case the answers were pretty clear. Intel was a new entrant: it did not sell image processors. The Image Il-

lusions technology was sustaining: it promised better performance than ASICs, as defined by the largest and most profitable customers. According to Disruption, an *entrant* with a *sustaining innovation* can expect to fail.

So that is what Thurston predicted.

Copyright © 2011 by Michael E. Raynor.
Foreword copyright © 2011 by Clayton M. Christensen.
All rights reserved.

Excerpted with permission from Crown Business, an imprint of the Crown Publishing Group, a division of Random House, Inc., New York.