

Managers often make significant business decisions based on little more than convincing book jacket blurbs. They should hold themselves—and the experts—to a higher standard.

Why Hard-Nosed Executives Should Care About Management Theory

by Clayton M. Christensen and Michael E. Raynor

Managers often make significant business decisions based on little more than convincing book jacket blurbs. They should hold themselves—and the experts—to a higher standard.

Why Hard-Nosed Executives Should Care About Management Theory

by Clayton M. Christensen and Michael E. Raynor

Imagine going to your doctor because you're not feeling well. Before you've had a chance to describe your symptoms, the doctor writes out a prescription and says, "Take two of these three times a day, and call me next week."

"But—I haven't told you what's wrong," you say. "How do I know this will help me?"

"Why wouldn't it?" says the doctor. "It worked for my last two patients."

No competent doctors would ever practice medicine like this, nor would any sane patient accept it if they did. Yet professors and consultants routinely prescribe such generic advice, and managers routinely accept such therapy, in the naive belief that if a particular course of action helped other companies to succeed, it ought to help theirs, too.

Consider telecommunications equipment provider Lucent Technologies. In the late 1990s, the company's three operating divisions were reorganized into 11 "hot businesses." The idea was that each business would be run largely independently, as if it were an internal entrepreneurial start-up. Senior executives pro-

claimed that this approach would vault the company to the next level of growth and profitability by pushing decision making down the hierarchy and closer to the marketplace, thereby enabling faster, better-focused innovation. Their belief was very much in fashion; decentralization and autonomy appeared to have helped other large companies. And the startups that seemed to be doing so well at the time were all small, autonomous, and close to their markets. Surely what was good for them would be good for Lucent.

It turned out that it wasn't. If anything, the reorganization seemed to make Lucent slower and less flexible in responding to its customers' needs. Rather than saving costs, it added a whole new layer of costs.

How could this happen? How could a formula that helped other companies become leaner, faster, and more responsive have caused the opposite at Lucent?

It happened because the management team of the day and those who advised it acted like the patient and the physician in our opening vi-

nette. The remedy they used—forming small, product-focused, close-to-the-customer business units to make their company more innovative and flexible—actually does work, when business units are selling modular, self-contained products. Lucent’s leading customers operated massive telephone networks. They were buying not plug-and-play products but, rather, complicated system solutions whose components had to be knit together in an intricate way to ensure that they worked correctly and reliably. Such systems are best designed, sold, and serviced by employees who are not hindered from coordinating their interdependent interactions by being separated into unconnected units. Lucent’s managers used a theory that wasn’t appropriate to their circumstance—with disastrous results.

Theory, you say? Theory often gets a bum rap among managers because it’s associated with the word “theoretical,” which connotes “impractical.” But it shouldn’t. A theory is a statement predicting which actions will lead to what results and why. Every action that managers take, and every plan they formulate, is based on some theory in the back of their minds that makes them expect the actions they contemplate will lead to the results they envision. But just like Monsieur Jourdain in Molière’s *Le Bourgeois Gentilhomme*, who didn’t realize he had been speaking prose all his life, most managers don’t realize that they are voracious users of theory.

Good theories are valuable in at least two ways. First, they help us make predictions. Gravity, for example, is a theory. As a statement of cause and effect, it allows us to predict that if we step off a cliff we will fall, without requiring that we actually try it to see what happens. Indeed, because reliable data are available solely about the past, using solid theories of causality is the only way managers can look into the future with any degree of confidence. Second, sound theories help us interpret the present, to understand what is happening and why. Theories help us sort the signals that portend important changes in the future from the noise that has no strategic meaning.

Establishing the central role that theory plays in managerial decision making is the first of three related objectives we hope to accomplish in this article. We will also describe how good theories are developed and give an idea of how a theory can improve over time. And, fi-

nally, we’d like to help managers develop a sense, when they read an article or a book, for what theories they can and cannot trust. Our overarching goal is to help managers become intelligent consumers of managerial theory so that the best work coming out of universities and consulting firms is put to good use—and the less thoughtful, less rigorous work doesn’t do too much harm.

Where Theory Comes From

The construction of a solid theory proceeds in three stages. It begins with a description of some phenomenon we wish to understand. In physics, the phenomenon might be the behavior of high-energy particles; in business, it might be innovations that succeed or fail in the marketplace. In the exhibit at right, this stage is depicted as a broad foundation. That’s because unless the phenomenon is carefully observed and described in its breadth and complexity, good theory cannot be built. Researchers surely head down the road to bad theory when they impatiently observe a few successful companies, identify some practices or characteristics that these companies seem to have in common, and then conclude that they have seen enough to write an article or book about how all companies can succeed. Such articles might suggest the following arguments, for example:

- Because Europe’s wireless telephone industry was so successful after it organized around a single GSM standard, the wireless industry in the United States would have seen higher usage rates sooner if it, too, had agreed on a standard before it got going.

- If you adopt this set of best practices for partnering with best-of-breed suppliers, your company will succeed as these companies did.

Such studies are dangerous exactly because they would have us believe that because a certain medicine has helped some companies, it will help all companies. To improve understanding beyond this stage, researchers need to move to the second step: classifying aspects of the phenomenon into categories. Medical researchers sort diabetes into adult onset versus juvenile onset, for example. And management researchers sort diversification strategies into vertical versus horizontal types. This sorting allows researchers to organize complex and confusing phenomena in ways that highlight their most meaningful differences. It is then possible

Clayton M. Christensen is the Robert and Jane Cizik Professor of Business Administration at Harvard Business School in Boston. **Michael E. Raynor** is a director with Deloitte Research and a professor at the Richard Ivey School of Business in London, Ontario, Canada. This article elaborates on a central theme of their forthcoming book, *The Innovator’s Solution* (Harvard Business School Press, 2003).

to tackle stage three, which is to formulate a hypothesis of what causes the phenomenon to happen and why. And that's a theory.

How do researchers improve this preliminary theory, or hypothesis? As the downward loop in the diagram below suggests, the process is iterative. Researchers use their theory to predict what they will see when they observe further examples of the phenomenon in the various categories they had defined in the second step. If the theory accurately predicts what they are observing, they can use it with increasing confidence to make predictions in similar circumstances.¹

In their further observations, however, researchers often see something the theory cannot explain or predict, an anomaly that suggests something else is going on. They must then cycle back to the categorization stage and add or eliminate categories—or, sometimes, rethink them entirely. The researchers then build an improved theory upon the new categorization scheme. This new theory still explains the previous observations, but it also explains those that had seemed anomalous. In other words, the theory can now predict more accurately how the phenomenon should work in a wider range of circumstances.

To see how a theory has improved, let's look at the way our understanding of international trade has evolved. It was long thought that countries with cheap, abundant resources would have an advantage competing in industries in which such resources are used as important inputs of production. Nations with inexpensive electric power, for example, would have a comparative advantage in making products that require energy-intensive production methods. Those with cheap labor would excel in labor-intensive products, and so on. This theory prevailed until Michael Porter saw anomalies the theory could not account for. Japan, with no iron ore and little coal, became a successful steel producer. Italy became the world's dominant producer of ceramic tile, even

though its electricity costs were high and it had to import much of the clay.

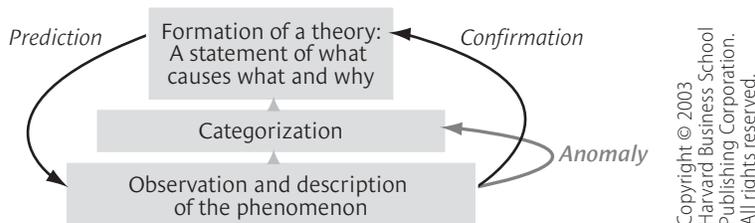
Porter's theory of competitive clusters grew out of his efforts to account for these anomalies. Clusters, he postulated, lead to intense competition, which leads companies to optimize R&D, production, training, and logistics processes. His insights did not mean that prior notions of advantages based on low-cost resources were wrong, merely that they didn't adequately predict the outcome in every situation. So, for example, Canada's large pulp and paper industry can be explained in terms of relatively plentiful trees, and Bangalore's success in computer programming can be explained in terms of plentiful, low-cost, educated labor. But the competitive advantage that certain industries in Japan, Italy, and similar places have achieved can be explained only in terms of industry clusters. Porter's refined theory suggests that in one set of circumstances, where some otherwise scarce and valuable resource is relatively abundant, a country can and should exploit this advantage and so prosper. In another set of circumstances, where such resources are not available, policy makers can encourage the development of clusters to build process-based competitive advantages. Governments of nations like Singapore and Ireland have used Porter's theory to devise cluster-building policies that have led to prosperity in just the way his refined theory predicts.

We'll now take a closer look at three aspects of the theory-building process: the importance of explaining what causes an outcome (instead of just describing attributes empirically associated with that outcome); the process of categorization that enables theorists to move from tentative understanding to reliable predictions; and the importance of studying failures to building good theory.

Pinpointing Causation

In the early stages of theory building, people typically identify the most visible attributes of the phenomenon in question that appear to be *correlated* with a particular outcome and use those attributes as the basis for categorization. This is necessarily the starting point of theory building, but it is rarely ever more than an important first step. It takes a while to develop categories that capture a deep understanding of what *causes* the outcome.

Consider the history of people's attempts to



fly. Early researchers observed strong correlations between being able to fly and having feathers and wings. But when humans attempted to follow the “best practices” of the most successful flyers by strapping feathered wings onto their arms, jumping off cliffs, and flapping hard, they were not successful because, as strong as the correlations were, the would-be aviators had not understood the fundamental causal mechanism of flight. When these researchers categorized the world in terms of the most obvious visible attributes of the phenomenon (wings versus no wings, feathers versus no feathers, for example), the best they could do was a statement of *correlation*—that the possession of those attributes is associated with the ability to fly.

Researchers at this stage can at best express their findings in terms of degrees of uncertainty: “Because such a large percentage of those with wings and feathers can fly when they flap (although ostriches, emus, chickens, and kiwis cannot), in all probability I will be able to fly if I fabricate wings with feathers glued on them, strap them to my arms, and flap hard as I jump off this cliff.” Those who use research still in this stage as a guide to action often get into trouble because they confuse the correlation between attributes and outcomes with the underlying causal mechanism. Hence, they do what they think is necessary to succeed, but they fail.

A stunning number of articles and books about management similarly confuse the correlation of attributes and outcomes with causality. Ask yourself, for example, if you’ve ever seen studies that:

- contrast the success of companies funded by venture capital with those funded by corporate capital (implying that the source of capital funding is a cause of success rather than merely an attribute that can be associated with a company that happens to be successful for some currently unknown reason).
- contend that companies run by CEOs who are plain, ordinary people earn returns to shareholders that are superior to those of companies run by flashy CEOs (implying that certain CEO personality attributes cause company performance to improve).
- assert that companies that have diversified beyond those SIC codes that define their core businesses return less to their shareholders than firms that kept close to their core (thus

leaping to the conclusion that the attributes of diversification or centralization cause shareholder value creation).

- conclude that 78% of female home owners between the ages of 25 and 35 prefer this product over that one (thus implying that the attributes of home ownership, age, and gender somehow cause people to prefer a specific product).

None of these studies articulates a theory of causation. All of them express a correlation between attributes and outcomes, and that’s generally the best you can do when you don’t understand what causes a given outcome. In the first case, for example, studies have shown that 20% of start-ups funded by venture capitalists succeed, another 50% end up among the walking wounded, and the rest fail altogether. Other studies have shown that the success rate of start-ups funded by corporate capital is much, much lower. But from such studies you can’t conclude that your start-up will succeed if it is funded by venture capital. You must first know what it is about venture capital—the mechanism—that contributes to a start-up’s success.

In management research, unfortunately, many academics and consultants intentionally remain at this correlation-based stage of theory building in the mistaken belief that they can increase the predictive power of their “theories” by crunching huge databases on powerful computers, producing regression analyses that measure the correlations of attributes and outcomes with ever higher degrees of statistical significance. Managers who attempt to be guided by such research can only hope that they’ll be lucky—that if they acquire the recommended attributes (which on average are associated with success), somehow they too will find themselves similarly blessed with success.

The breakthroughs that lead from categorization to an understanding of fundamental causality generally come not from crunching ever more data but from highly detailed field research, when researchers crawl inside companies to observe carefully the causal processes at work. Consider the progress of our understanding of Toyota’s production methods. Initially, observers noticed that the strides Japanese companies were making in manufacturing outpaced those of their counterparts in the United States. The first categorization efforts were di-

rected vaguely toward the most obvious attribute—that perhaps there was something in Japanese culture that made the difference.

When early researchers visited Toyota plants in Japan to see its production methods (often called “lean manufacturing”), though, they observed more significant attributes of the system—inventories that were kept to a minimum, a plant-scheduling system driven by kanban cards instead of computers, and so on. But unfortunately, they leaped quickly from attributes to conclusions, writing books assuring managers that if they, too, built manufacturing systems with these attributes, they would achieve improvements in cost, quality, and speed comparable to those Toyota enjoys. Many manufacturers tried to make their plants conform to these lean attributes—and while many reaped some improvements, none came close to replicating what Toyota had done.

The research of Steven Spear and Kent Bowen has advanced theory in this field from such correlations by suggesting fundamental causes of Toyota’s ability to continually improve quality, speed, and cost. Spear went to work on several Toyota assembly lines for some time. He began to see a pattern in the way people thought when they designed any process—those for training workers, for instance, or installing car seats, or maintaining equipment. From this careful and extensive observation, Spear and Bowen concluded that all processes at Toyota are designed according to four specific rules that create automatic feedback loops, which repeatedly test the effectiveness of each new activity, pointing the way toward continual improvements. (For a detailed account of Spear and Bowen’s theory, see “Decoding the DNA of the Toyota Production System,” HBR September–October 1999.) Using this mechanism, organizations as diverse as hospitals, aluminum smelters, and semiconductor fabricators have begun achieving improvements on a scale similar to Toyota’s, even though their processes often share few visible attributes with Toyota’s system.

Moving Toward Predictability

Manned flight began to be *possible* when Daniel Bernoulli’s study of fluid mechanics helped him understand the mechanism that creates lift. Even then, though, understanding the mechanism itself wasn’t enough to make manned flight perfectly *predictable*. Further

research was needed to identify the circumstances under which that mechanism did and did not work.

When aviators used Bernoulli’s understanding to build aircraft with airfoil wings, some of them still crashed. They then had to figure out what it was about those circumstances that led to failure. They, in essence, stopped asking the question, “What attributes are associated with success?” and focused on the question, “Under what circumstances will the use of this theory lead to failure?” They learned, for example, that if they climbed too steeply, insufficient lift was created. Also, in certain types of turbulence, pockets of relatively lower-density air forming under a wing could cause a sudden down spin. As aviators came to recognize those circumstances that required different technologies and piloting techniques and others that made attempting flight too dangerous, manned flight became not just possible but predictable.

In management research, similar breakthroughs in predictability occur when researchers not only identify the causal mechanism that ties actions to results but go on to describe the circumstances in which that mechanism does and does not result in success. This enables them to discover whether and how managers should adjust the way they manage their organizations in these different circumstances. Good theories, in other words, are *circumstance contingent*: They define not just what causes what and why, but also how the causal mechanism will produce different outcomes in different situations.

For example, two pairs of researchers have independently been studying why it is so difficult for companies to deliver superior returns to shareholders over a sustained period. They have recently published carefully researched books on the question that reach opposing conclusions. *Profit from the Core* observes that the firms whose performance is best and lasts longest are, on average, those that have sought growth in areas close to the skills they’d honed in their core businesses. It recommends that other managers follow suit. *Creative Destruction*, in contrast, concludes that because most attractive businesses ultimately lose their luster, managers need to bring the dynamic workings of entrepreneurial capitalism inside their companies and be willing to create new core businesses.

Unfortunately, many management researchers are so focused on how companies succeed that they don’t study failure.

Because they've juxtaposed their work in such a helpful way, we can see that what the researchers actually have done is define the critical question that will lead to the predictability stage of the theory-building cycle: "Under what circumstances will staying close to the core help me sustain superior returns, and when will it be critical to set the forces of creative destruction to work?" When the researchers have defined the set of different situations in which managers might find themselves relative to this question and then articulated a circumstance-contingent theory, individuals can begin following their recommendations with greater confidence that they will be on the right path for their situation.

Circumstance-contingent theories enable managers to understand what it is about their present situation that has enabled their strategies and tactics to succeed. And they help managers recognize when important circumstances in their competitive environment are shifting so they can begin "piloting their plane" differently to sustain their success in the new circumstance. Theories that have advanced to this stage can help make success not only possible and predictable but sustainable. The work of building ever-better theory is never finished. As valuable as Porter's theory of clusters has proven, for example, there is a great opportunity for a researcher now to step in and find out when and why clusters that seem robust can disintegrate. That will lead to an even more robust theory of international competitive advantage.

The Importance of Failures

Note how critical it is for researchers, once they have hypothesized a causal mechanism, to identify circumstances in which companies did exactly what was prescribed but failed. Unfortunately, many management researchers are so focused on how companies succeed that they don't study failure. The obsession with studying successful companies and their "best practices" is a major reason why platitudes and fads in management come and go with such alarming regularity and why much early-stage management thinking doesn't evolve to the next stage. Managers try advice out because it sounds good and then discard it when they encounter circumstances in which the recommended actions do not yield the predicted results. Their conclusion most often is,

"It doesn't work."

The question, "When *doesn't* it work?" is a magical key that enables statements of causality to be expressed in circumstance-contingent ways. For reasons we don't fully understand, many management researchers and writers are afraid to turn that key. As a consequence, many a promising stream of research has fallen into disuse and disrepute because its proponents carelessly claimed it would work in every instance instead of seeking to learn when it would work, when it wouldn't, and why.

In a good doctor-patient relationship, doctors usually can analyze and diagnose what is wrong with a specific patient and prescribe an appropriate therapy. By contrast, the relationship between managers, on the one hand, and those who research and write about management, on the other, is a distant one. If it is going to be useful, research must be conducted and written in ways that make it possible for readers to diagnose their situation themselves. When managers ask questions like, "Does this apply to my industry?" or "Does it apply to service businesses as well as product businesses?" they really are probing to understand the circumstances under which a theory does and does not work. Most of them have been burned by misapplied theory before. To know unambiguously what circumstance they *are* in, managers need also to know what circumstances they are *not* in. That is why getting the circumstance-defined categories right is so important in the process of building useful theory.

In our studies, we have observed that industry-based or product-versus-service-based categorization schemes almost never constitute a useful foundation for reliable theory because the circumstances that make a theory fail or succeed rarely coincide with industry boundaries. *The Innovator's Dilemma*, for example, described how precisely the same mechanism that enabled upstart companies to upend the leading, established firms in disk drives and computers also toppled the leading companies in mechanical excavators, steel, retailing, motorcycles, and accounting software. The circumstances that matter to this theory have nothing to do with what industry a company is in. They have to do with whether an innovation is or is not financially attractive to a company's business model. The mechanism—the resource allocation process—causes the established leaders to win the competitive fights when an

For most managers, trying out a new idea to see if it works is simply not an option: There is too much at stake.

innovation is financially attractive to their business model. And the same mechanism disables them when they are attacked by disruptive innovators whose products, profit models, and customers are not attractive to their model.

We can trust a theory only when, as in this example, its statement describing the actions that must lead to success explains how they will vary as a company's circumstances change. This is a major reason why the world of innovating managers has seemed quite random—because shoddy categorization by researchers has led to one-size-fits-all recommendations that have led to poor results in many circumstances. Not until we begin developing theories that managers can use in a circumstance-contingent way will we bring predictable success to the world of management.

Let's return to the Lucent example. The company is now in recovery: Market share in key product groups has stabilized, customers report increased satisfaction, and the stock price is recovering. Much of the turnaround seems to have been the result, in a tragic irony, not just of undoing the reorganization of the 1990s but of moving to a still more centralized structure. The current management team explicitly recognized the damage the earlier decentralization initiatives created and, guided by a theory that is appropriate to the complexity of Lucent's products and markets, has been working hard to put back in place an efficient structure that is aligned with the needs of Lucent's underlying technologies and products.

The moral of this story is that in business, as in medicine, no single prescription cures all ills. Lucent's managers felt pressured to grow in the 1990s. Lucent had a relatively centralized decision-making structure and its fair share of bureaucracy. Because most of the fast-growing technology companies of the day were comparatively unencumbered with such structures, management concluded that it should mimic them—a belief not only endorsed but promulgated by a number of management researchers. What got overlooked, with disastrous consequences, was that Lucent was emulating the attributes of small, fast-growing companies when its circumstances were fundamentally different. The management needed a theory to guide it to the organizational structure that was optimal for the circumstances the company was actually in.

Becoming a Discerning Consumer of Theory

Managers with a problem to solve will want to cut to the chase: Which theory will help them? How can they tell a good theory from a bad one? That is, when is a theory sufficiently well developed that its categorization scheme is indeed based not on coincidences but on causal links between circumstances, action, and results? Here are some ideas to help you judge how appropriate any theory or set of recommendations will be for your company's situation.

- When researchers are just beginning to study a problem or business issue, articles that simply describe the phenomenon can become an extremely valuable foundation for subsequent researchers' attempts to define categories and then to explain what causes the phenomenon to occur. For example, early work by Ananth Raman and his colleagues shook the world of supply chain studies simply by showing that companies with even the most sophisticated bar code-scanning systems had notoriously inaccurate inventory records. These observations led them to the next stage, in which they classified the types of errors the scanning systems produced and the sorts of stores in which those kinds of errors most often occurred. Raman and his colleagues then began carefully observing stocking processes to see exactly what kinds of behaviors could cause these errors. From this foundation, then, a theory explaining what systems work under what circumstances can emerge.

- Beware of work urging that revolutionary change of everything is needed. This is the fallacy of jumping directly from description to theory. If the authors imply that their findings apply to all companies in all situations, don't trust them. Usually things are the way they are for pretty good reasons. We need to know not only where, when, and why things must change but also what should stay the same. Most of the time, new categorization schemes don't completely overturn established thinking. Rather, they bring new insight into how to think and act in circumstance-contingent ways. Porter's work on international competitiveness, for example, did not overthrow preexisting trade theory but rather identified a circumstance in which a different mechanism of action led to competitive advantage.

- If the authors classify the phenomenon

they're describing into categories based upon its attributes, simply accept that the study represents only a preliminary step toward a reliable theory. The most you can know at this stage is that there is some relationship between the characteristics of the companies being studied and the outcomes they experience. These can be described in terms of a general tendency of a population (20% of all companies funded by venture capital become successful; fewer of those funded by corporate capital do). But, if used to guide the actions of your individual company, they can easily send you on a wing-flapping expedition.

- Correlations that masquerade as causation often take the form of adjectives—*humble* CEOs create shareholder value, for instance, or *venture-capital* funding helps start-ups succeed. But a real theory should include a mechanism—a description of how something works. So a theory of how funding helps start-ups succeed might suggest that what venture capitalists do that makes the difference is meter out small amounts of funds to help the companies feel their way, step by step, toward a viable strategy. Funding in this way encourages start-ups to abandon unsuccessful initiatives right away and try new approaches. What corporate capitalists often do that's less effective is to flood a new business with a lot of money initially, allowing it to pursue the wrong strategy far longer. Then they pull the plug, thus preventing it from trying different approaches to find out what will work. During the dot-com boom, when venture capitalists flooded start-ups with money, the fact that it was venture money *per se* didn't help avert the predictable disaster.

- Remember that a researcher's findings can almost never be considered the final word. The discovery of a circumstance in which a theory did not accurately predict an outcome is a triumph, not a failure. Progress comes from refining theories to explain situations in which they previously failed, so without continuing our examination of failure, management theory cannot advance.

When Caveat Emptor Is Not Enough

In shopping for ideas, there is no Better Business Bureau managers can turn to for an assessment of how useful a given theory will be to them. Editors of management journals publish a range of different views on important issues—leaving it to the readers to decide which theories they should use to guide their actions.

But in the marketplace of ideas, caveat emptor—letting the reader beware—shirks the duty of research. For most managers, trying out a new idea to see if it works is simply not an option: There is too much at stake. Our hope is that an understanding of what constitutes good theory will help researchers do a better job of discovering the mechanisms that cause the outcomes managers care about, and that researchers will not be satisfied with measuring the statistical significance of correlations between attributes and outcomes. We hope they will see the value in asking, “When *doesn't* this work?” Researching that question will help them decipher the set of circumstances in which managers might find themselves and then frame contingent statements of cause and effect that take those circumstances into account.

We hope that a deeper understanding of what makes theory useful will enable editors to choose which pieces of research they will publish—and managers to choose which articles they will read and believe—on the basis of something other than authors' credentials or past successes. We hope that managers will exploit the fact that good theories can be judged on a more objective basis to make their “purchases” far more confidently. 

1. Karl Popper asserted that when a researcher reaches the phase in which a theory accurately predicts what has been observed, the researcher can state only that the test or experiment “failed to disconfirm” the theory. See *The Logic of Scientific Discovery* (Harper & Row, 1968).

Reprint [R0309D](#); *Harvard Business Review* OnPoint [4937](#)

To order, see the next page
or call 800-988-0886 or 617-783-7500
or go to www.hbr.org



Harvard Business Review OnPoint articles enhance the full-text article with a summary of its key points and a selection of its company examples to help you quickly absorb and apply the concepts. *Harvard Business Review* OnPoint collections include three OnPoint articles and an overview comparing the various perspectives on a specific topic.

Further Reading

This article is also part of the [Harvard Business Review OnPoint edition](#), Product no. 4937, which includes the following suggestions for further reading:

[Skate to Where the Money Will Be](#)

by Clayton M. Christensen, Michael E. Raynor, and Matthew Verlinden
Harvard Business Review OnPoint
Enhanced Edition
November 2001
HBR OnPoint 8113

[The Innovator's Solution: Creating and Sustaining Successful Growth](#)

by Clayton M. Christensen and Michael E. Raynor
Harvard Business School Press
forthcoming October 2003
HBR OnPoint 8520

To Order

For reprints, *Harvard Business Review* OnPoint orders, and subscriptions to *Harvard Business Review*:
Call 800-988-0886 or 617-783-7500.
Go to www.hbr.org

For customized and quantity orders of reprints and *Harvard Business Review* OnPoint products:
Call Frank Tamoshunas at 617-783-7626,
or e-mail him at ftamoshunas@hbsp.harvard.edu