The Opportunity and Threat of Disruptive Technologies

Clayton M. Christensen

The following article is an edited transcript of the plenary presentation given by Clayton M. Christensen (Harvard Business School) at the 2001 Materials Research Society Fall Meeting on November 26 in Boston.

Introduction

I am honored to be with you today. I will give context for where the puzzles developed that have led to the research that I summarized in the book *The Innovator's Dilemma*. Companies that at one point were very successful and widely admired, fell from their leadership position within 5 to 10 years. This was very poignant for minicomputer companies, for example, which collapsed after their peak in the world economy during the 1970s and 1980s. While the companies were thriving, the business press attributed their successes to the abilities of their management teams. However, when these companies began to quickly unravel around 1988, the business press attributed their failures to the ineptitude of the very same management teams that had been in place during the years of success. I wondered how good managers could get that bad that fast. I was also suspicious because the minicomputer companies fell in unison. It was not just Digital, but also Data General, Prime, Wang, and Nixdorf. While the companies could have colluded on price, they would never have colluded to collapse together. Something more fundamental had to be happening. Oddly, I found that the reason these companies failed, not just in the computer industry, but over and over again in a variety of industries, is that they were well managed.

The Disruptive Technology Model

Certain paradigms of good management that we teach in business schools sow the seeds of a company's eventual failure. Figure 1 presents a model that consists of three elements. The performance of a product or service is plotted over time. The first element of the model, represented by the dotted line, shows a trajectory of improvement that customers can absorb or utilize over time. This line represents the median in every market. At the high end of every market, customers are very demanding and will never be satisfied with the best products available. At the low end are customers who can be satisfied by very little.

The steeply sloping lines of the figure represent a distinctly different trajectory of improvement in the market. They represent the improvement that innovators make available as, generation after generation, they introduce new and improved products. This trajectory of technological progress almost invariably outstrips the ability of customers to utilize those improvements; companies can almost always innovate faster than people's lives can change to fully utilize those innovations. In pursuing more profitable customers who will pay higher prices for better products in more demanding tiers of the market, companies are prone to overshoot what the original set of customers is able to utilize at a later point in time. For example, in the mid-1980s when we were first learning to do word processing on those early personal computers, it was not uncommon for us to have to stop our fingers and wait for the Intel 8088 processor to catch up with us. But as Intel has pursued more profitable customers, to whom it could sell better products for higher margins, it has way-overshot what customers in the mainstream of the business computing market are able to utilize. What this also means is that technology that is initially not good enough to be embraced by customers in the mainstream of the market can improve at such a rapid rate that it later intersects with the needs of mainstream customers.

The third piece of the model relates to "disruptive technology," which refers to technology that initially underperforms the incumbent technology, but typically offers more flexibility, convenience, or a lower price. A disruptive technology can serve to create new markets, but it can also upset the equilibrium of companies with good reputations and good products.

Disruption in Action

Computer Industry: From Disk Drives to Intel

I initiated my study of what was driving the failure of some groups of firms with the disk drive industry. In five of six product generations in the industry, a new company would enter and rise to the top only to be toppled by another company that entered and rose to the top. I wanted to develop a model that would

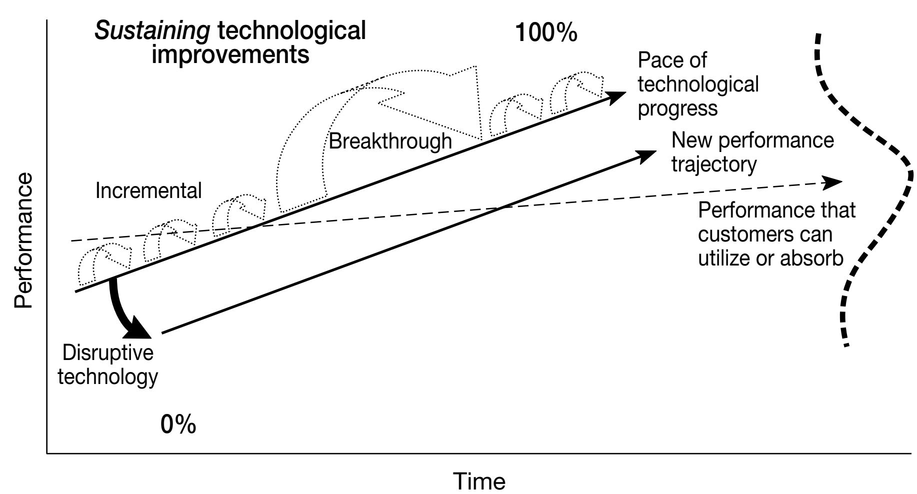

Figure 1. The disruptive technologies model.

https://doi.org/10.1557/mrs2002.81 Published online by Cambridge University Press

explain this phenomenon, which I could then apply to other industries. The disk drive industry served as the "fruit fly" in which it was possible to study many "generations" quickly.

Through this process of the disk drive study, we built a database of every model of every disk drive introduced in the industry's history. For each of the products, we obtained data on each of the components that was used in the drive. This allowed us to locate the point in the industry's history at which each new technology was first used. In this process, we identified 116 distinctly different new technologies. Of the 116, 111 sustained the trajectory of performance improvement in the market, as it existed at the time. In other words, the purpose of the technology was to make a good product better. Of these 111 sustaining technologies, a good number were the simple, year-to-year engineering advancements expected of good companies. Additionally, a remarkable number were truly breakthrough innovations that made good products *much* better, as depicted in Figure 1. What was interesting is that it did not matter technologically how difficult the innovation was. In all 111 of these cases, 100% of the time, the companies that led the industry in developing and adopting the new technology had been the leaders of the old technology. As long as the innovation enabled the company to make a better product that they could sell for higher margins to their best customers, it was on management's radar screen and it received the necessary resources to get it done. Only five technologies disrupted this trajectory of improvement, which means they brought to the market a product that was not as good as what historically had been available. Oddly, this type of technology always killed the industry leader; in none of these cases did the company that had been the leader in the prior generation survive into the next generation. Even more remarkable, these were relatively simple technologies.

Let us take this phenomenon to the computer world. When we consider the sequence of minicomputers that Digital introduced into its markets through the 1970s and 1980s, it had a perfect record in developing the technologies—whether incremental or radical—that *sustained* the trajectory of improvement in many computers. The early personal computers that emerged in the late 1970s and early 1980s—such as the Apple II, marketed to children as a toy—were a disruptive technology. Digital received no signal from the market that the personal computer mattered because none of its customers could use personal computers. They were not good enough. Yet, because the trajectory of technological progress outstrips the ability of customers to use it, what was at one point a toy improved at such a rapid rate that it intersected with the customers' needs of many computer users at a later point in time. A whole population of minicomputer companies was knocked out of the market together as a result of listening to their customers.

> Something more fundamental had to be happening. The reason these companies failed, not just in the computer industry, but over and over again in a variety of industries, is that they were well managed.

When I was consulting recently with the senior management at Intel, the chair, Andy Grove, puzzled over my description of "disruptive technology," and recast it. He said the concept could be more accurately characterized as trivial technology that disrupts the business model of the leaders, which is what makes it difficult to handle. Grove said that Digital did not have a technology problem, but rather a problem with its business model. For example, a minicomputer is fairly expensive and complex, and it has to be sold directly to the customer. Digital had to make gross margins of about 45% just to be acceptably profitable in the market. The senior management faced proposals to invest in two types of products. Some entailed developing better computers than Digital had ever made before that it could sell for \$200,000 and 60% gross margins by penetrating the market where customers had historically only been able to use mainframes. Other proposals entailed investing in personal computers that could be sold for \$2,000 and, in the very best of years, promised gross margins of only 40% that were headed to 20% rapidly. Grove concluded that the choice management had to make was whether to invest in better products that would serve the best customers, which would improve the company's profit margins, or invest in worse products that their customers could not use and that would erode the company's profit margin.

This is the innovator's dilemma. For 111 of the 116 cases in the disk drive study, these paradigms of good management—listening to your best customers and focusing investments on those innovations that promise the highest profit margins—provided good guidance for the innovators. However, in the five cases in which the disruptive technology came in, those same paradigms of good management paralyzed the leading companies and made it impossible for them to respond.

From Steel Mini-Mills to Sony

Let us look now at a very different industry, which is the steel industry. The question is how did the steel mini-mills in North America grow to account for half of steel production? Most of the world's steel is made in large integrated mills that require billions of dollars to build. Mini-mills, in contrast, melt scrap in electric-arc furnaces and can make steel of any given quality for a 20%-lower cost than integrated mills.

While it seems that every leading steel company worldwide would take on mini-mill technology because the technology is straightforward and the chance to reduce the costs of making a commodity like steel by 20% is a sure path to profit. Right? But not a single integrated steel company in North America or Europe has yet successfully invested in a minimill. Only recently, a few companies in Japan have acquired already successful mini-mills. When mini-mills became technologically viable in the mid-1960s, the quality they could produce was very low. Only the concrete reinforcement-bar market accepted their product because rebar had pretty loose specifications. As the mini-mills attacked the competitive rebar tier of the market, the integrated mills happily dropped out of that tier because their gross margins in that market segment were only about 7%. As the mini-mills expanded their capacity, the integrated mills shut the lines down or reconfigured them to make more profitable products. By about 1975, the minimills had driven the integrated mills out of the rebar market. The mini-mills, with a 20% cost advantage, made *lots* of money making rebar in competition with the high-cost integrated mills. But when they finally succeeded in driving the last integrated mill out of the rebar market in about 1975, the price of rebar collapsed by 20%—because you now had low-cost mini-mills slugging it out with low-cost mini-mills in a commodity market. How could they begin making money again?

Material Matters is a forum for expressing personal points of view on issues of interest to the materials community.

They had no option but to turn to products that were larger and of higher quality, which means they attacked the next tier up of the market: angle iron and thicker bar and rod. The mini-mills little by little figured out how to make better products. As they hit the next tier of the market, the integrated mills were happy to get out of it, too, because the 12% margins were still relatively low. As the mini-mills expanded capacity, the integrated mills shut the lines down and reconfigured them to make better products, and while the mini-mills were competing against the high-cost integrated players, the mini-mills once again made a lot of money—until 1984, when they succeeded in driving the last high-cost integrated player out of the angle-iron market. The price of angle iron then dropped by 20%.

The mini-mills then moved up into structural beams. The margin that the integrated mills were making in structural beams was about 18%, and the market was three times larger than those the mini-mills had already conquered. In the early 1980s, beams could not be rolled in a mini-mill format to meet specifications. The mini-mills, however, figured out how to roll high-end structural bar—like 12-, 16-, and 24-inch I-beams and H-beams. By 1992, they had driven U.S. Steel out of the structural-beam market. By 1996, Bethlehem Steel had also been driven out of that market.

Integrated steel companies apparently practiced good management. Every time they lopped off the low end of their product line, their reported gross-profit margins improved. Also, every time the mini-mills added the very same products to the high end of their product lines, their reported gross-profit lines improved. As a perfect symmetry of motivation, the high-end players were motivated to leave the very markets that the low-end players were motivated to enter.

Other companies that recently experienced this kind of phenomenon include Sony, which came into the low end of the market with low-quality transistor pocket radios in the mid-1960s. Sony then became the highest-quality electronics maker worldwide. Seiko started with the inferior plastic digital watch and then became the highest-quality watchmaker worldwide. In the photocopier industry, Canon competed against the huge machines by Xerox with a low-end, limited-feature, tabletop photocopier that filled a niche for small companies. Xerox's customers could not use the little tabletop machines, but in a different application, the inferior machine moved up-market, and so customers consumed infinitely more photocopying as a result of that industry having been disrupted.

Interdependent and Modular Architectures

In the early 1980s, IBM had better operating-system technology than did Microsoft and better microprocessor technology than did Intel. Yet, IBM chose to outsource the two technologies to those respective companies and in the process put into business the two companies that subsequently dominated the computer industry. IBM designed and assembled computers, where subsequently no money was made. At the time IBM made those decisions, it was generally regarded as managerially astute. This experience brings up the question as to when to work in-house and when to outsource.

A review of the history of most industries reveals that during the early period when the performance of the product or service is not good enough for what customers in the mainstream need, the architecture of the product tends to be proprietary and interdependent in character. When the product is not good enough, competition is dictated through the development of better products. A proprietary architecture in an interdependent architecture can come closest to the frontier of what is technologically possible.

For example, in the early years of the mainframe computer industry, an independent supplier of operating systems, core memory, or logic circuitry could not have existed because the design of each of those subsystems depended upon the design of each of the other subsystems. They had to be interdependently architectured. An independent-contract manufacturer of mainframe computers could not exist because the way mainframe computers were made depended upon the way that they were designed, and the way that they were designed depended upon the way they were going to be made.

Integrated companies initially tend to dominate their industries, such as IBM and Digital, General Motors and Ford, Alcoa and U.S. Steel, and so on. In the most demanding tiers of these markets, the architectures are still proprietary and interdependent.

Once the technology is more than good enough, how do companies compete for the business of customers in the now-overserved tiers of the market? The answer is that they must compete differently. They must compete in speed-to-market, in the ability to respond quickly to changes in customers' needs, and in the ability to customize the features and functions of products to the needs of smaller and smaller niches of the market. To be fast, flexible, and responsive, the architecture of the product has to evolve from an interdependent to a modular architecture—ultimately one that is built around industry standards. When that happens, the industry tends to disintegrate as specialized companies provide one piece of value added. The population of special-

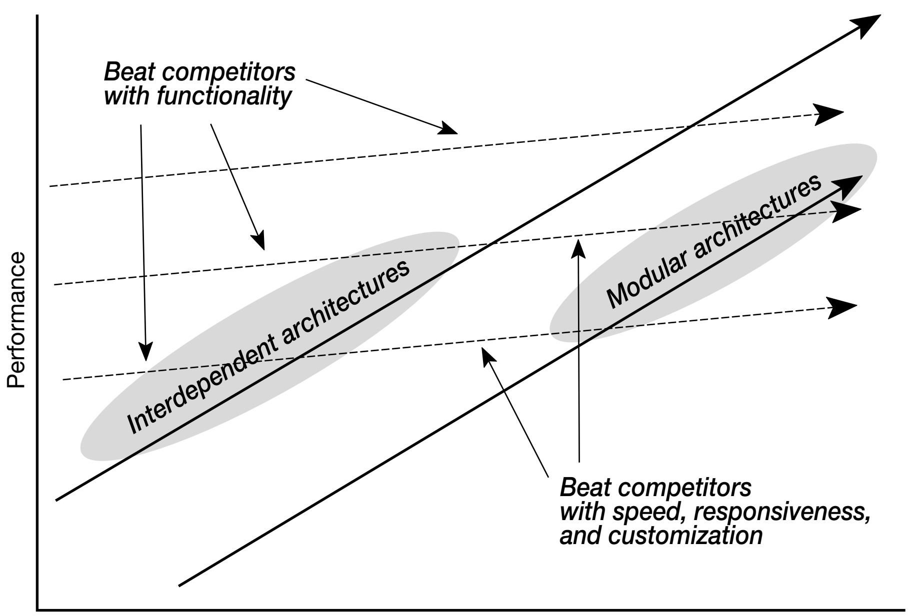

Figure 2. Model of integration and disintegration. Integrated competitors have the advantage when products are not good enough. Focused firms overtake over-served markets.

https://doi.org/10.1557/mrs2002.81 Published online by Cambridge University Press

ized companies then tends to replace what were the dominant integrated companies in an earlier era. In fact, the dominant integrated companies have to disintegrate in order to compete effectively (see Figure 2).

In an interdependent architecture, the companies that make the money tend to be the ones that do the design and assembly of the product that is not good enough, that is, the end-use product. Companies that were the suppliers tended to live a miserable profit-free existence year after year. So, for example, IBM, with 70% market share, made 95% of the industry's profit. The companies that supplied IBM were wrung out every quarter. General Motors had a 55% market share; it made 80% of the industry's profit. A supplier to General Motors made no money. In a modular architecture, the company that makes the money tends to flip; those that design and assemble the products that the customers use tend to get knocked out and the ones that make money are the ones who supply the subsystems that are not yet good enough. To illustrate, if an engineer at Compaq receives instructions to design a better computer than Dell, the engineer could put in a faster microprocessor, a higher-capacity disk drive, or more dynamic random-access memory. However, anything the engineer does, the competitor can instantly copy. So the ability to make money flips back to Intel, who provides the performance-generating subsystem that the assembler has to use in order to keep it at their leading edge.

I recently published an article in *Harvard Business Review* called "Skate to Where the Money Will Be." The title refers to what made Wayne Gretzky a great hockey player: He skated not to where the puck was, but to where it would be. It suggests that the place in the value-added chain where the money is made at one point is not likely where it will be made in the future. This is happening now in the car industry. Automobiles have become more than good enough for what most consumers need. The industry used to take six years to design and deliver a new car. Now the design cycle has converged on two years, and customers can order a car custom-assembled to their specifications, deliverable in five days. Car companies accomplish this by shifting the architecture of these cars from an interdependent to a modular architecture. They still procure subsystems from tier 1 suppliers. Inside the subsystems, just like an Intel chip set, the architecture becomes more proprietary and interdependent even while the external interfaces of these subsystems become more standardized. True to form, the integrated car companies are having to disintegrate to keep pace with the accelerating cycles. However, car companies like GM sold off the pieces of their business where the money was subsequently likely to be made, as did IBM. In other words, GM and IBM skated to where the money used to be rather than to where the money would be in the future.

Modularity in Action: Advanced Materials

In many ways, the purpose of advanced materials is to be a sustaining technology; that is, to improve the functionality of products at the leading edge. When we started our company, Ceramics Process Systems, in the early 1980s, we introduced components made of ceramics to the automobile-engine world. One of the leading European car companies that was a major investor in our company wanted us to start making silicon nitride piston pins because the calculations revealed that if we could reduce the mass of a reciprocating component in the engine, it would greatly improve its efficiency. With the specifications they gave us, we molded the piston pins. When the piston pins were tested, however, the test engine vibrated uncontrollably. It turned out that over the years, the car company had distributed mass elsewhere in the engine to compensate for the reciprocating mass of the steel. There simply was not a plug-compatible modular opportunity for the advanced materials at the edge of performance. Everything depended upon everything else. We succeeded because we finally found applications where the technology was almost plug-compatible, which is far away from the leading edge.

How Disruptive Growth Begins

The way disruptive technology makes it to the obvious existing market is by finding simple applications and then gradually improving. Figure 3 shows that the transistor was a disruptive technology relative to the vacuum tube because it could not handle the power in the early 1950s that would be required to be used in the market as it existed at the time. Every one of the vacuum-tube players saw the transistor coming. They all worked on the technology in their laboratories, but framed it as a technological deficiency instead of an opportunity. In other words, companies were challenged with ways of making transistors good enough to be used in the vacuum-tube market. But transistors took root in a completely new market. The first application was for a small hearing aid in 1953. This application valued the transistor for the very attributes that made it useless in the mainstream market. In 1955, Sony introduced its pocket transistor radio, which was such a low-quality product that it did not compete in the mainstream market. It only appealed to low-end customers, such as teenagers, because it enabled them to do something that had not been possible before—such as listen to Elvis Presley out of earshot of their parents. Sony moved up into portable televisions in about 1959, where a huge

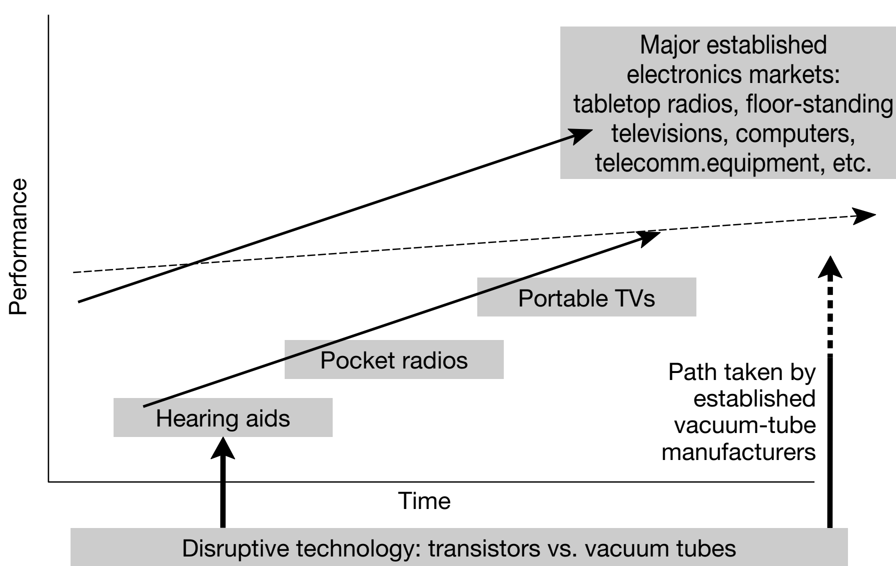

Figure 3. Successful disruptors target smaller, "green space" markets instead of stretching toward existing, larger markets.

https://doi.org/10.1557/mrs2002.81 Published online by Cambridge University Press

new market emerged. Other companies felt no pain because the market served different kinds of customers and different kinds of applications than their mainstream ones. By about the mid-1960s, solid-state electronic components became good enough for the mainstream market to use and they replaced vacuum tubes. This is a typical pattern. A disruptive technology first takes root away from the core of the mainstream market, where users value it despite its limitations.

In the last chapter of *The Innovator's Dilemma*, I discuss whether the electric vehicle and photovoltaics are disruptive technologies. In its concern over air pollution, California passed legislation in the early 1990s mandating that every car company had to sell 2% of its unit volume in California as zero-emission or electric vehicles by 1998, or it would be prohibited from selling gas-powered cars. The car companies worldwide mounted huge efforts to develop battery technology for electric vehicles. By about 1996, it was clear that they would not accomplish this goal. As an example, customers told Chrysler that they wanted a car that went about 250 miles between refueling stops and could accelerate from 0 to 60 miles per hour in about nine seconds. To get its electric minivan to accelerate fast and cruise far, Chrysler loaded 1600 pounds of lead-acid batteries in the back of it. The minivan, however, only cruised 80 miles and accelerated from 0 to 60 mph in 19 seconds. One unfortunate side effect was that it took four times longer to stop this minivan than the gas-powered one because of the balance in the back. It retailed for $100,000. During consumer tests, Chrysler found no buyer. The government has since pushed the target date to 2003.

In studying this issue, one of my students determined that the industry is coming up with the right answers to the wrong questions. The right question is locating the market in which a customer wants a car that does not cruise far or accelerate quickly, and is cheap. He suggested that parents of teenagers and people in the retirement community would prefer this type of car. He suggested that people in places like Bangkok would find a use for such a car since they cannot travel quickly on the streets and would actually prefer the car to shut down as they waited in traffic.

Another limitation to forcing electric vehicles into existing applications is that they do not fit with existing infrastructures. If plugged into a 100-V outlet, they take an unacceptable eight hours to recharge. Even if every filling station was outfitted with 480 V of service, the car would take 20 minutes to refuel; consumers will not want to wait that long. On the other hand, if the electric vehicle were a neighborhood vehicle, designed not to drive on a freeway, but for teenagers to cruise around town to high school and see their friends, or for retirement communities, it would be convenient for the users to plug in the vehicle overnight. By targeting such a market, the car companies avoid the infrastructural interdependencies that happen when they stretch to use the new technology in the most demanding applications. Over and over again, though, leading companies try to force the technology into demanding, obvious, and large applications.

One of the most personally rewarding things about *The Innovator's Dilemma* has been the number of people who have told me that our theory has helped them frame a problem that they are facing in their own lives. I hope that you too will be able to use our frameworks to better understand some of the issues that your company may be confronting.

Recommended Reading

1. C.M. Christensen, *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail* (Harvard Business School Press, Boston, 1997).
2. C.M. Christensen, M. Raynor, and M. Verlinden, "Skate to Where the Money Will Be," *Harvard Business Review* (November 2001), p. 72.
3. C.M. Christensen, T. Craig, and S. Hart, "The Great Disruption," *Foreign Affairs* **80** (2) (March/April 2001), p. 80.
4. C.M. Christensen and M. Overdorf, "Meeting the Challenge of Disruptive Change," *Harvard Business Review* **78** (2) (March–April 2000), p. 66.

***Clayton M. Christensen**, the Robert and Jane Cizik Professor of Business Administration at the Harvard Business School (HBS), is the author of* The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail, *a book that won the Global Business Book Award for the best business book published in 1997. In 1984, Christensen co-founded Ceramic Process Systems, a developer of ceramic products, for which he served as chair and president. Christensen holds a BA degree in economics from Brigham Young University and an MPhil degree in economics from Oxford University, where he studied as a Rhodes Scholar. He received his MBA degree from HBS in 1979, graduating as a George F. Baker Scholar, and was awarded his DBA degree from HBS in 1992. Christensen was a White House Fellow, serving the secretaries of transportation.*