
Service businesses struggle with a reality that is foreign to manufacturers: Customers “interfere” with their operations. To deliver consistent quality at sustainable cost, companies must learn to manage that involvement.

Breaking the Trade-Off Between Efficiency and Service

by Frances X. Frei

Service businesses struggle with a reality that is foreign to manufacturers: Customers “interfere” with their operations. To deliver consistent quality at sustainable cost, companies must learn to manage that involvement.

Breaking the Trade-Off Between Efficiency and Service

by Frances X. Frei

What if a manufacturer had to deal with customers waltzing around its shop floor? What if they showed up, intermittently and unannounced, and proceeded to muck up the manufacturer’s carefully designed processes left and right? For most service businesses, that’s business as usual. In a restaurant or a rental car agency or most of the other service companies that make up the bulk of mature economies today, customers aren’t simply the open wallets at the end of an efficient supply chain. They’re directly involved in ongoing operations. The fact that they introduce tremendous variability—but complain about any lack of consistency—is an everyday reality.

Dealing with that variability is a central challenge in making a service offering profitable. But little in managers’ conventional training or tool kits equips them to deal with it effectively. Operations management theory, rooted in the manufacturing context, typically has only one thing to say about variability: It must be eliminated. Any educated manager learns to recognize it as the enemy of quality.

In the service context, the challenge is far more subtle. First, it wouldn’t be wise to drive out all variability; customers judge the quality of their experience in large part by how much of the variability they introduce is accommodated, not how sternly it is denied. Second, it wouldn’t be possible to do so. While manufacturers have virtually complete control over the cost and quality of their production inputs, service companies face this one, huge exception: Their customers are themselves key inputs to the production process. That form of input is, by its nature, capricious, emotional, and adamantly disinterested in the company’s profit agenda.

My research over the past several years has been aimed at helping service organizations overcome the challenge of customer-introduced variability. I’ve studied a wide variety of service companies, some of which prospered while others experienced escalating costs in the face of eroding customer satisfaction. The framework that has emerged from that study can help managers make better decisions about

how and how much to reduce or accommodate the variability customers introduce. As the stories in the following article make clear, there are multiple ways to combat the effects of any type of variability, and the best solution is not always immediately apparent. But by using a systematic process to diagnose problems and design and fine-tune interventions, managers can reduce the impact of variability and enhance the competitiveness of their service.

Five Types of Variability

The first step in managing the variability introduced by customers is to understand the forms it can take. Customers introduce variability to operations in no fewer than five ways, so it is critical to sort out which type is causing mischief before designing interventions.

Arrival variability. The first type of variability that creates challenges for service companies is an obvious one: Customers do not all want service at the same time or at times necessarily convenient for the company. Many a grocery store manager has bemoaned shoppers' inability to space their transactions such that checkout clerks remain busy and lines do not form at the registers. The classic way to address arrival variability is to require appointments or reservations, but that makes sense only in certain situations. In many service environments, such as retail stores, call centers, or emergency rooms, the customers themselves cannot foresee or delay their needs. The resulting inefficiencies have inspired a large body of work in what's known as queuing theory and many solutions (including those described by W. Earl Sasser in "Match Supply and Demand in Service Industries," HBR November–December 1976).

Request variability. Film buffs will recall the diner scene in the movie *Five Easy Pieces*, in which actor Jack Nicholson asks for a side order of wheat toast. The rule the waitress invokes—no substitutions—is a time-honored way to limit request variability, or the range of what customers ask for in a service environment. While it's hard to imagine operations grinding to a halt over an order of toast, the fact that customers' desires don't emerge along standard lines poses real challenges for virtually every kind of service business. At an advertising agency, each client is executing a unique marketing strategy. At a resort, vacationers want dif-

ferent amenities. Even at a single-service business like Jiffy Lube, customers show up with different makes and models of automobiles.

Capability variability. Perhaps less obviously, service businesses must also work with customers whose own capabilities differ. Whether because of greater knowledge, skill, physical abilities, or resources, some customers perform tasks easily and others require hand-holding. This capability variability clearly becomes more important when customers are active participants in the production and delivery of a service. A cleaning service may arrive, do its work, and leave, having had no real interaction with the customer. The customer's particular capabilities make little difference to how well the crew does its job. In a medical setting, by contrast, a patient may be more or less able to describe his symptoms, and that will affect the quality of the health care he receives.

Effort variability. When customers must perform a role in a service interaction, it's up to them how much effort they apply to the task. An internal accountant may or may not take care to hand over well-organized files to her company's independent auditor. A shopper at a warehouse club may or may not have the remaining energy to return his massive shopping cart to one of the corrals in the parking lot. Such effort variability has an impact on service quality and cost, either directly for the engagement at hand or indirectly for other patrons.

Subjective preference variability. Customers also vary in their opinions about what it means to be treated well in a service environment. One diner appreciates the warmth of a waiter's first-name introduction; another resents his presumption of intimacy. When a top partner in a law firm lavishes attention on engagements, some clients will be gratified by the proof of their cases' importance. Others will think those expensive billable hours could be doled out more judiciously. These are personal preferences, but they introduce as much unpredictability as any other variable and make it that much harder to serve a broad base of customers.

It's possible to think of these five forms of variability sequentially because they reflect the process by which many service transactions unfold. The customer arrives, makes a request, plays a part in the process requiring some level

Frances X. Frei (frei@hbs.edu) is an associate professor of business administration in the Technology and Operations Management unit at Harvard Business School in Boston.

of capability and effort, and assesses the experience according to personal preferences. At any of these points, life is easier for a service provider if it is dealing with a narrow band of variability. Where the band is wide, service quality and efficiency are at risk.

The taxonomy above is important because operational issues in a service business can often be traced to problems created by customer-introduced variability. But the right strategies to manage, say, effort variability (often involving incentives) can be completely different from the strategies for dealing with capability variability (typically some sort of training). Before managers can draft an appropriate response, they must diagnose which variability is at issue.

A Classic Trade-Off

Wherever customer-introduced variability creates operational issues for a company, managers face a choice: Do they want to accommodate that variability or reduce it? Generally, companies that emphasize the service experience tend toward accommodation, and those that emphasize operational simplicity—usually as a means to keep costs low—tend toward reduction. The two approaches are in constant tension.

Consider a classic illustration of a reduction strategy: the restaurant menu. Menus, by their nature, are a way to constrain request variabil-

ity. They put a limit on what would otherwise be an infinite number of potential orders and therefore make it possible for a restaurant to offer meals of consistent quality at a reasonable cost. But customers chafe under too many constraints (again, recall Jack Nicholson’s rage in *Five Easy Pieces*). For them, the ability to request variations in preparation, ingredients, and side dishes—or to order off the menu entirely—is part of a premier dining experience. When restaurants do not accommodate special orders, they reduce the complexity of the operating environment but also may diminish service quality. Companies that use reduction strategies tend to attract price-conscious customers who are willing to trade off an excellent service experience for low prices. People who choose discount airlines, bulk retailers, movie matinees, and off-peak travel options essentially reduce their collective variability by conforming to a company’s operational needs, even at the risk of an inferior service experience.

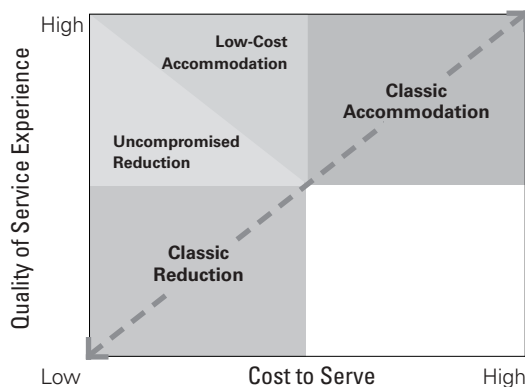
Accommodation strategies take different forms, depending on the business and type of customer-introduced variability. Very often, accommodation involves asking experienced employees to compensate for the variations among customers. For example, in a business where customers have divergent views of how service should be delivered (a business, that is, with high subjective-preference variability), a veteran employee learns to diagnose customer types. By making on-the-fly adaptations to suit their preferences, he essentially “protects” the customers from having to make many adjustments of their own.

It costs more, of course, to hire, train, and keep employees who can compensate for customers. Like most accommodation strategies, this one forces the company to bear the brunt of the variability. Therefore, the success of an accommodation strategy usually hinges on a company’s ability to persuade customers to pay more to cover the added expense. Generally, only companies at the high end of their competitive landscape can command such a premium. Those at the low end must rely on strategies to reduce variability.

But managing customer-introduced variability does not have to come down to a stark trade-off between cost and quality. Some companies have met the challenge without damaging either the service experiences they provide or their operating environments. In a matrix

Overcoming the Trade-Off

Managers in service businesses often assume that they face a tough choice: either accommodate customers’ various desires and behaviors at high cost or refuse to accommodate variability and risk customer defection. But other options exist—those above the diagonal of the matrix—which let companies offer a high level of accommodation at low cost or reduce variability without damaging the service experience.



representing the classic trade-off as a linear function of cost to serve versus the quality of the service experience, these companies have gone “above the diagonal.” (See the exhibit “Overcoming the Trade-Off.”) The matrix shows possibilities beyond *classic reduction* and *classic accommodation* strategies: the potential for what can be termed *uncompromised reduction* and *low-cost accommodation*.

Here’s an example of an uncompromised reduction approach. A company can greatly re-

duce the impact of variability on its operating environment without compromising the service experience by targeting customers on the basis of variability type. If, for example, a college fears that admitting students of varying intellectual capabilities will complicate its operations, it can choose only students whose standardized test scores fall within a narrow band. The students get the benefit of a tailored curriculum without the school’s having to support more than one. Likewise, a company

Strategies for Managing Customer-Introduced Variability

Once a company has determined which type of customer-introduced variability is creating operational difficulties, it must choose which of four basic strategies to pursue. The chart outlines tactics that have proven to be effective in each category.

	Classic Accommodation	Low-Cost Accommodation	Classic Reduction	Uncompromised Reduction
Arrival	<ul style="list-style-type: none"> • Make sure plenty of employees are on hand 	<ul style="list-style-type: none"> • Hire lower-cost labor • Automate tasks • Outsource customer contact • Create self-service options 	<ul style="list-style-type: none"> • Require reservations • Provide off-peak pricing • Limit service availability 	<ul style="list-style-type: none"> • Create complementary demand to smooth arrivals without requiring customers to change their behavior
Request	<ul style="list-style-type: none"> • Make sure many employees with specialized skills are on hand • Train employees to handle many kinds of requests 	<ul style="list-style-type: none"> • Hire lower-cost specialized labor • Automate tasks • Create self-service options 	<ul style="list-style-type: none"> • Require customers to make reservations for specific types of service • Persuade customers to compromise their requests • Limit service breadth 	<ul style="list-style-type: none"> • Limit service breadth • Target customers on the basis of their requests
Capability	<ul style="list-style-type: none"> • Make sure employees are on hand who can adapt to customers’ varied skill levels • Do work for customers 	<ul style="list-style-type: none"> • Hire lower-cost labor • Create self-service options that require no special skills 	<ul style="list-style-type: none"> • Require customers to increase their level of capability before they use the service 	<ul style="list-style-type: none"> • Target customers on the basis of their capability
Effort	<ul style="list-style-type: none"> • Make sure employees are on hand who can compensate for customers’ lack of effort • Do work for customers 	<ul style="list-style-type: none"> • Hire lower-cost labor • Create self-service options with extensive automation 	<ul style="list-style-type: none"> • Use rewards and penalties to get customers to increase their effort 	<ul style="list-style-type: none"> • Target customers on the basis of motivation • Use a normative approach to get customers to increase their effort
Subjective Preference	<ul style="list-style-type: none"> • Make sure employees are on hand who can diagnose differences in expectations and adapt accordingly 	<ul style="list-style-type: none"> • Create self-service options that permit customization 	<ul style="list-style-type: none"> • Persuade customers to adjust their expectations to match the value proposition 	<ul style="list-style-type: none"> • Target customers on the basis of their subjective preferences

Wherever customer-introduced variability creates operational issues for a company, managers face a choice: Do they want to accommodate that variability or reduce it?

faced with subjective preference variability can target customers who are predisposed to want service to be delivered the same way. It isn't always easy to know where customers fall on the relevant spectrum of variability, and there isn't always sufficient demand within a given band of customers to sustain a business. However, companies that find such a niche can benefit from reduced variability without requiring customers to adjust.

Companies that achieve low-cost accommodation most often do it by persuading customers to serve themselves. This strategy is very effective for high arrival or request variability, both of which complicate labor scheduling. Obviously, when the customer is responsible for much of the labor, the right labor is provided at the right moment. Further, by having customers serve themselves, companies are allowing the service experience to vary with customers' capability and effort (accommodating capability and effort variability) and giving customers control of the service environment (accommodating subjective preference variability). The online auction house eBay shows how far this model can be taken: Virtually all the labor of selling and buying on the site is performed by customers, not by eBay employees.

The problem is that many companies, unlike eBay, have established precedents whereby employees perform certain tasks for customers. For those companies to succeed with a low-cost accommodation approach, they must persuade customers to do the work. This "persuasion" is typically achieved through some redefinition of the customer value proposition. That is, customers need to feel compensated in some way—whether through lower prices, greater customization, or other benefits of being in control—in order to feel good about doing work they think the company should be doing.

Solutions in Practice

Once a management team understands the types of variability customers introduce, and the possibilities for reducing or accommodating variability, the challenge of managing service operations becomes more tractable. Let's revisit the four strategic responses discussed above: classic accommodation, classic reduction, low-cost accommodation, and uncompromised reduction. (The exhibit "Strategies for Managing Customer-Introduced Variabil-

ity" gives examples for each.) The history of successful service companies reveals that they've used every one of these strategies at one time or another.

In the late 1990s, for example, Dell faced the challenge of high arrival and request variability in its customer service operations as the company considered adding large servers to its product array. It knew that these high-end servers, and the corporate customers who bought them, would create significant new demands for responsive service. Given the competitive context, Dell would have to be prepared to satisfy these demands around the clock and across a broad spectrum of possible malfunctions. As a new entrant in the market, lacking scale in its service operations, the company faced a trade-off between maintaining an underutilized and expensive service operation (accommodation of variability) and achieving higher predictability and utilization by, for instance, asking customers to schedule appointments (reduction of variability). Dell understood that, from its customers' perspective, accommodation was the only alternative, so the company set out to find a way to insulate itself from the effects of variability without compromising customers' service experiences.

Dell's solution was to outsource on-site customer service to third-party providers, who served more than one client and thus were less disrupted by the variability imposed by Dell's customers than Dell would have had it acted alone. The move posed some risk: By giving up this customer contact in exchange for lower costs, Dell could have lost control of its customer relationships. The company prevented that through strict vigilance, staying in close touch with customers to discuss their needs and to assess their experiences with the third-party providers. By maintaining this contact, Dell effectively made the providers' role less prominent. In the end, the company achieved a low-cost accommodation of the variability its customers brought to the service relationship.

Starbucks provides an excellent example of the deft handling of capability variability. The coffee shop chain allows customers to choose among many permutations of sizes, flavors, and preparation techniques in its beverages. In the interests of filling orders accurately and efficiently, Starbucks trains its counter clerks to call out orders to beverage makers in a particu-

Managing customer-introduced variability does not have to come down to a stark trade-off between cost and quality.

lar sequence. It is all the better when customers themselves can do so. Therefore, Starbucks attempts to teach customers its ordering protocol in at least two ways. It produces a “guide to ordering” pamphlet for customers to peruse, and it instructs clerks to repeat the order to the customer not in the way it was presented but in the correct way. The tone is not one of rebuke, but nevertheless most customers learn to avoid the implied correction by stating their order in the way that helps Starbucks’s operations—with no hit to the service experience. Indeed, for some customers, getting the order right is an aspiration, a small victory on the way to the office. It’s a clever solution, achieving an uncompromised reduction of variability.

Companies facing issues relating to effort variability often resort to the classic accommodation approach: They simply require employees to do the work for the lazier customers, with an obvious impact on operating costs. Some companies, however, try to compel those customers to work a little harder. As decades of research on employee motivation have emphasized, there are two ways to change behavior: *instrumental means* and *normative means*. Instrumental means are formal rewards and penalties for specified behaviors—the basic carrots and sticks of discipline. Normative means rely more subtly but often more effectively on shame, blame, and pride. In the case of Zipcar, an auto-sharing service, motivating customers to make the effort asked of them is particularly important because their actions influence not only themselves but also other customers. A car returned to its parking space late by one user spells real inconvenience for the next. While late fees are a common instrumental control for this type of situation, they risk being perceived by the customer as a license to be late. Indeed, late fees often help compensate a business for customers’ costly choices, but they are not always effective in changing their behavior.

Normative controls, which make customers *want* to behave, can be far more successful, but these incentives are difficult to craft. Why would one customer necessarily care about the inconvenience suffered by another? To use normative controls effectively, companies need to create an environment in which customers care about the impact of their behavior on others. Such an environment exists on eBay, where customers serve one another with great

care, in large part because of the customer-to-customer commitment the company has built through tools such as feedback stars, which publicize buyers’ and sellers’ past behavior. Normative controls can be particularly important when instrumental incentives have failed. (As Steven Levitt and Stephen Dubner relate in *Freakonomics*, when a day care center instituted late fees for parents who were not on time to pick up their children, lateness got worse. The fee reduced the parents’ guilt, which had been a powerful normative incentive.) Companies like Zipcar must not only determine how they need customers to behave but also come up with effective ways to promote that behavior.

The best strategy for changing customers’ behavior is not always obvious, nor is the best strategy for managing a specific type of variability. Tiffany & Company, the luxury jeweler, suffered from missteps in 2001, when it failed to anticipate how customers would react to what seemed like a logical solution. Its problem was one that many retailers would like to have: The brand’s popularity was soaring among the so-called mass affluent segment—a fast-growing market of moneyed consumers. Consider that Tiffany’s hallmark had long been the graciousness of its service. As customers began crowding into its stores, this traditional service experience was being undermined. In particular, management noticed, with so many people milling around the floor it was hard for employees to uphold the first-come-first-served norm.

Tiffany dealt with this arrival variability with a tried-and-true device: the beeper. Upon arrival in the shop, customers were given a beeper and told they would be buzzed as soon as a service person was available. Unfortunately, the reaction of the customers Tiffany most wanted to protect—its most wealthy and loyal ones—was outrage. Management had failed to recognize that a more problematic form of variability—subjective preference variability—had disrupted the business. While the mass-market customer arriving in the store was well acquainted with beepers, and even felt well served by them, the more demanding luxury customer found them to be inconsistent with Tiffany’s historic commitment to white-glove service. Only after the company saw a dramatic plunge in satisfaction among the latter group did it confront its fundamental man-

agerial challenge: whether (and how) to serve two distinct segments of customers through a single retail channel. Tiffany's challenge was complicated by the fact that the less expensive silver jewelry was popular with both segments, which made it difficult to come up with a solution that segmented service on the basis of product type. (Subjective preference variability is also the focal point of Southwest Airlines' current dilemma. See the sidebar "Should Southwest Airlines Be More Accommodating?" for details.)

Gateway's attempt to manage customer variability failed for different reasons. Since its inception, the personal-computer maker had sold its products solely through direct channels. But faced with eroding market share, management decided to address the capability variability common in high-tech markets. It knew that it would be able to sell more PCs if it provided more hand-holding to consumers who lacked technical knowledge and confidence. This meant entering the retail market—and more, it meant creating exceptional retail

environments that enabled deep customer learning. When Gateway's new stores opened in 1996, they were undeniably impressive. Employees were experienced, helpful, and abundant (the employee-to-customer ratio was unusually high). Excellent educational materials were on hand, and the stores were conveniently located to ensure heavy foot traffic. Gateway succeeded spectacularly at bringing customers with all levels of expertise through the doors.

Fast-forward to April 2004, when the company was shuttering the last of more than 300 storefronts. How could this have happened? It wasn't that the strategy was ludicrous. The company had accurately diagnosed a problematic form of customer variability, and it had devised a way to manage its impact. Unfortunately, that way was expensive, and Gateway hadn't guaranteed that the people receiving the benefits of all that prepurchase accommodation would also bear the costs. Far too often, customers took their newly acquired understanding of what they needed and how it

Should Southwest Airlines Be More Accommodating?

As this article was being prepared for publication, a controversy erupted among the customers of 35-year-old Southwest Airlines.

Throughout its history, the company has employed an unusual boarding policy: no assigned seats. Instead, customers choose where to sit as they board the plane. The resulting cattle call has always produced some grumbling, particularly among the frail members of the herd. More troubling to the analysts who follow Southwest's fortunes, the boarding policy alienates some business travelers, who tend to arrive at airports just minutes before departure—and are generally willing to pay more for choice seats.

What the policy clearly has in its favor is efficiency. Southwest's average turnaround—the time that elapses between a plane's pulling up to a gate and its pulling away for the next take-off—is 40% faster than competitors'. The net effect is a level of fleet productivity that keeps costs and fares low. For some customers, the no-assigned-seating policy has another thing in its favor: egalitarianism. Passengers who travel infrequently, using their own money—

and who want to enjoy the experience—are not necessarily well served by airlines that favor frequent business fliers.

So in June 2006, when Southwest announced an experiment in which flights out of San Diego would have assigned seats, the resistance from longtime customers was dramatic. The experiment was designed to discover how boarding time would be affected by a less charged but also less chaotic process. Apparently, management had not anticipated the effect on the customer experience.

What's going on here, and how should Southwest proceed? It's useful to see the situation in terms of customer-introduced variability. Most airlines cater to a wide range of subjective preferences among their customers. But Southwest, in support of its low-cost model, opted for a strategy of reducing, rather than accommodating, variability. People who like the airline's approach self-select into its customer base. Others adjust their preferences to avail themselves of Southwest's low fares.

It's important to recognize that part of the reason so many customers are willing to do

this is that Southwest's offering is so obviously egalitarian. Indeed, the airline seems to have realized that any departure from this ethos could be problematic. In the past, it has denied requests for various services from some frequent fliers even when they would have cost nothing to implement and would not have affected critical operating metrics, such as turnaround time. Amid the credible threat of customer defection, the company held fast, knowing that if passengers saw some receiving special perks, their subjective preference for equal treatment would be violated.

Now that Southwest Airlines has a customer base accustomed to open seating, it faces an uphill battle of behavior change if it decides to make a switch. Management of customer-introduced variability will come to the fore as a key to competitiveness and profitability. For the moment, the airline seems ill equipped to take on that challenge. "I am averaging easily 100 letters a day," Southwest president Colleen Barrett told the *Baltimore Business Journal*. "I am just literally in a state of shock."

worked and then placed an order with one of Gateway's low-price competitors.

Managing the Operational Behavior of Customers

It's clear from the examples above that the effective management of variability in service operations often requires a company to influence customers' behavior. That can be a hard goal to achieve, given that a company's operational concerns are not usually foremost in its customers' minds. Managers attempting this kind of intervention should plan their actions carefully in a three-step process.

Diagnose the problem. The operational problems caused by customers' discretionary behavior can range from the seemingly minor—some customers are late to their appointments—to issues that can have a large impact on profitability. As a first step, managers must understand the root causes of problematic customer behavior. Unless the behavioral problem is accurately diagnosed, no subsequent action to correct it will be effective.

The experience of retail bank First Union in the late 1990s makes this point dramatically. Because the bank misdiagnosed the type of customer variability it faced, it took actions that were inappropriate to the situation. First Union had created many self-service options for customers—primarily through ATMs, voice response units, and Web pages—and hoped that the cost of the innovations would be more than recouped by lower costs in branch operations. However, when customers continued to visit the branches to transact business in person with tellers, the investment in self-service technology failed to meet expectations. Management concluded that the problem was, in essence, one of capability variability: Not all customers had learned what the technology could do and how to use it. To address this problem, First Union stationed greeters at the doors of its branches to ask customers the nature of their business with the bank that day. If the transaction could easily be accomplished through an ATM (as was usually the case), the greeter would recommend the self-service technology and offer guidance on how to use it. Within months of this management intervention, First Union had lost roughly 20% of its most recently acquired accounts. Not long after, First Union merged with Wachovia and dropped its name.

The cause for the loss was not hard to trace: It came down to a misunderstanding of why the self-service options had not caught on among all customers. The variability that was actually at issue was not capability variability but effort variability. Customers with time on their hands preferred to wait in line to have the teller do all the work.

Managers can avoid that kind of misdiagnosis by conducting a thorough analysis guided by some straightforward questions:

- What is problematic about customers' current behavior? What is the danger of leaving the behavior unchanged?
- What are the hypotheses of the cause of the behavior? In determining the hypotheses, consider the role of the five types of customer-introduced variability and state hypotheses for each as the cause.
- Which hypotheses make the most sense? Which are less plausible? Is management invested in a particular outcome? What assumptions is the company making about what customers value?
- How will these hypotheses be tested? Who will be responsible for the data they produce? If the outcome has significant implications for strategy or operations, who will lead the change process?

Had First Union (or Tiffany, drawing on an earlier example) gone through this kind of exercise, the ineffectiveness of the solution would have been identified well before it was rolled out in a full-scale, live operating environment. First Union hypothesized that customers' resistance to self-service technologies reflected a gap in their capabilities, so the bank jumped directly to training them (using greeters) without sufficiently testing the hypothesis. Acting on untested hypotheses is a common mistake when the logic of what is (presumably) good for customers is widely accepted. First Union reasoned that if customers only knew how much better off they would be using ATMs, they would surely choose to serve themselves. Had the bank tested this assumption—by, say, asking customers why they used particular channels and what they thought of alternative channels—it would have exposed the flaws in its thinking. Managers often confuse capability and effort variability because their symptoms can be identical.

At Tiffany, the company observed overcrowding, hypothesized that arrival variability

was the issue, and designed a store-level solution. Had the company been more thorough in exploring the problem—particularly in analyzing the differences in subjective preferences between customer segments—it could have learned about the potential incompatibility of the two segments and designed a company-level solution.

Design a mutually beneficial operating role for customers. With the appropriate diagnosis, companies can design an operating role for customers that creates explicit value for both parties. As in step one, a set of questions can guide the creation of this mutually beneficial role:

- What do customers gain from their new role? Are they better off than before? Are they still better off than they would be in the hands of competitors?
- What does the company gain from customers' new role? What is the intended impact of their new behavior on the company's performance?
- Is it realistic that customers will behave the way the company wants them to? What assumptions are managers making about human motivation?

The difficulty in creating value for customers often comes from untested assumptions about their behavior and perceptions, like the ones made by management at First Union. Usually there are many ways to create value for customers—but one of them is not to make customers feel they are worse off than they were before the change.

The difficulty in creating value for service companies is that revenue and cost are often not tightly linked in such businesses. This isn't the case in product-based businesses, where each transaction can be evaluated according to the clear associated revenue minus the cost of production. Service businesses often use a model more akin to buffet pricing: Customers, having paid a fee, can conduct as many transactions as they desire. This makes it difficult to understand the value being created at different points in the relationship and allows such mistakes as Gateway's foray into high-touch retailing. Indeed, the free riding the company suffered is a major risk for any business in which customers need expensive prepurchase service and rivals offer easy substitutions.

Test and improve the solution. Because of the inherently complicated nature of cus-

tomers behavior, it is useful to test approaches to influencing behavior before rolling them out on a broad scale. However, while pilot tests can reveal critical system flaws at a limited cost, such tests are often executed incorrectly. The three most common mistakes are as follows:

- *Creating testing environments that are substantially different from the real environment.* Sometimes pilots take place in a better climate than customers will actually experience. The most common differences in a testing environment are more experienced employees, artificially ample resources, and limited exposure to variability.
- *Creating incentives—whether implicit or explicit—for the test to have a positive outcome.* This often comes in the form of a promise that the test manager will be responsible for the full-scale rollout if the test has a positive outcome (regardless of whether the company learned anything).
- *Designing a test that has no controls.* If customers change their behavior following a test, it is difficult to know whether the change should be attributed to the test or to other external factors if the test had no controls.

One way to overcome the last mistake is to use what Wells Fargo refers to as the “challenger-champion” model. For every new initiative, the company selects a sample to test the new initiative (the challenger sample) and a similar, matched sample (the champion sample). After the initiative is tested on only the challenger sample, the company tracks differences in behavior between the two samples.

More generally, we have found that pilot tests are effective when managers can affirmatively answer the following questions:

- Is the pilot program being tested under typical circumstances? Are the employees, customers, and resources consistent with the company's real operating environment?
- Is the goal of the pilot to learn as much as possible (rather than to demonstrate the value of the new system)? Is this goal clear to both employees and managers?
- Is it clear that managers' performance is not based on a positive outcome of the pilot?
- Are customers and frontline employees involved in evaluating the circumstances of the test and in assessing results?
- Can managers articulate the explicit changes made as a result of the pilot test? (If

relatively few changes are made, that should be a red flag that the primary motivation of the test was proof-of-concept, not learning.)

Throwing a Customer in the Works

Profitably managing the variability implicit in customer heterogeneity, and developing effective levers to influence it, is a central challenge for service businesses. By extension, it is also a central challenge for developed economies. In the typical mature economy, service providers conduct more than 70% of commerce—yet the frameworks and tools for managing these businesses lag significantly behind those developed for manufacturing environments.

Understanding the workings of service businesses more thoroughly begins with identifying the things that make them different from manufacturers. Chief among these is the presence of the customer in operations. Customers perform roles that are either well or poorly designed for them and engage in behaviors that either benefit or harm the company. They make it nearly impossible to manage production in isolation from consumption. Companies that learn to manage the variability customers bring to the works will find that customers are the key to competitive advantage.

Netflix is an example of a company that capitalized on incumbents' mishandling of customer variability. When customers rent DVDs,

late returns are a major source of tension for both rental companies and customers. Companies have charged late fees—which customers often perceive to be draconian—in order to encourage people to return movies on time. But late fees have not only failed to change customers' behavior but also have been a significant source of customer dissatisfaction. Enter Netflix and its subscription model, which makes late fees obsolete by allowing people to keep movies for as long as they want. The customer's incentive to return a movie is being able to get the next movie on her request list.

Netflix saw an opportunity in the tension over late fees. The company knew from its research what its competitors didn't: Some customers value having control over how long they keep movies, but not at the high cost (and anxiety) of late fees. This left room for a middle ground, a premium subscription service that guarantees revenues while accommodating variability in usage time. While incumbents were trying to strong-arm their customers into "behaving," Netflix built a winning business model based on a deeper understanding of the true drivers of customer behavior.

Reprint [R0611E](#)

Harvard Business Review OnPoint [1498](#)

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

[Breaking the Trade-Off Between Efficiency and Service](#) is also part of the *Harvard Business Review* OnPoint collection [Serve Your Customers—Efficiently and Profitably](#), Product no. 1500, which includes these additional articles:

[Lean Consumption](#)

James P. Womack and Daniel T. Jones
Harvard Business Review
March 2005
Product no. 9432

[Putting the Service-Profit Chain to Work](#)

James L. Heskett, Thomas O. Jones,
Gary Loveman, W. Earl Sasser, Jr., and
Leonard A. Schlesinger
Harvard Business Review
July–August 2000
Product no. 4460

Harvard Business Review

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 Rich Gravelin at
617-783-7626,
or e-mail him at
rgravelin@hbsp.harvard.edu