

From Smart Pricing to Smarter Pricing: Using the Whole Pricing Picture

by Scott Borg and David Pyke

October, 2002

Abstract

Businesses have recently recognized that smarter pricing is often the fastest and easiest way to increase profits. Taking better account of their customers' price sensitivities can allow firms to set more advantageous prices and adjust them more dynamically, sometimes even on a customer-by-customer basis. To do this effectively, however, it is vital to take account of the whole pricing picture, not just part of it. This means paying equal attention to two curves: the Willingness-to-Pay curve of the customers and Opportunity Cost curve of the supply chain. Both of these curves have interesting, complex shapes. The goal in pricing is to capture as much as possible of the area between the two.

From Smart Pricing to Smarter Pricing: Using the Whole Pricing Picture

The fastest and easiest way to increase profits is often to improve pricing policies. Whether the product is an automobile, a software application, an amusement park admission, or a hamburger, raising or lowering the price by the right amount can sometimes change the entire financial profile of a company, even when everything else is left pretty much the same. Many CEO's have attained celebrity status and received huge bonuses largely because they got their pricing right. Companies have regularly beaten the rest of their industry simply by pricing better than their competitors.

In fact, the power of pricing is so great, it can sometimes conceal the fact that a company may not be doing other things right. This was what happened recently at Ford. They were managing their pricing so much better than their competitors that it seemed as though they must be doing everything else better too. The recent problems at Ford don't demonstrate the limitations of effective pricing nearly as much as they demonstrate its power. Yet even Ford could have made greater gains by taking its pricing innovations further.

The viability of a new product, such as the Segway, will often depend even more on pricing decisions. The Segway, unveiled with great fanfare in December 2001, is essentially a gyroscopically controlled scooter, designed to carry a single, standing passenger on sidewalks and over moderately uneven terrain. If the company producing it prices the Segway too high, they may get too few orders to survive. If they price it too low, they may get swamped with orders they can't fill. If they change their price or delivery time too often or too much, potential customers may end up so uncertain about the product that they hold off on ordering it. Any of these mistakes could cause losses to mount so rapidly that they could drive the company into bankruptcy.

The amazing thing, given the importance of pricing, is how few conceptual tools most executives have for dealing with the complete pricing picture. Managers making pricing decisions will often be very sophisticated in the way they deal with one part of the pricing problem. But the impact of their decisions on company profits is often disappointing, because they will allow another, equally vital part of the pricing picture to remain out of focus. To set prices for maximum profit, for example, it is critical to understand the effect of price changes

on manufacturing and the supply chain. Yet, most pricing decisions are made by marketing managers who have little understanding of supply chain dynamics. And most supply chain decisions are made by operations managers who have little exposure to pricing decisions. The habit of treating pricing as a specialist, functional matter has prevented managers from noticing how closely each aspect of a company's activities is connected to pricing. A vision that puts the pricing pieces together is rarely articulated.


The problem is growing more urgent as an increasing range of industries attempts to adopt the sort of dynamic pricing techniques pioneered by airlines and hotels. Using new software, point of sale (POS) data, and the internet, businesses are increasingly able to segment their markets, sometimes down to individual customers, and to adjust their prices to suit each segment. But this complicates the pricing problem more than it solves it. Decisions about pricing need to be made with a broader awareness of the way those choices will limit, expand, or segment business operations and a deeper awareness of how they will affect profits.

Fortunately, a state-of-the-art picture of the pricing problem is not, in its essentials, very complicated. To derive the maximum benefits from pricing policies, there are ultimately two things—and only two things—that you need to take into account. One is the Willingness-to-Pay curve of your customers. The other is the Opportunity Cost curve of your supply chain. The problem is that you need to have as clear as possible an idea of: a) how these two curves are shaped, b) how they line up against each other, and c) how they are likely to change over time. When these two curves, taken together, fall into certain configurations, a relatively small change in price can have a huge effect on profits.

The Willingness-to-Pay Curve

The Willingness-to-Pay curve deals with the part of the pricing picture that tends to receive the most attention. Each year companies spend large sums of money investigating and measuring “customer demand.” Yet even here, insufficient attention is sometimes given to the one aspect of customer demand that matters most—the Willingness-to-Pay curve itself. The Willingness-to-Pay curve is the most direct way of describing whether customers will purchase the product and how much of the product they will purchase. Each potential customer is represented by a horizontal segment of the Willingness-to-Pay curve. The monetary level of that segment represents the customer's “indifference point.” If the payment required to buy the product is below that amount, the customer will generally be willing to make the purchase. If the payment required is above that amount, the customer will walk away. If the customer is a business, the Willingness-to-Pay is the “break even point,” determining whether or not the deal would be profitable for that business. The horizontal dimension of the Willingness-to-Pay curve corresponds to the quantity of the product the customers would be willing to buy at that per-unit amount. This means that the Willingness-to-Pay curve provides a way to tell at a glance how changes in price will affect product sales.

The most important thing to realize about Willingness-to-Pay curves is that they usually don't look like the smooth, straight lines that were used to represent the graph of "demand" in your undergraduate economics class. Instead, a Willingness-to-Pay curve is more likely to resemble a series of steps or plateaus. (See Figure 1.) The step-like shape of Willingness-to-Pay curves is especially conspicuous when the products are more innovative, more distinctive, or constitute a larger or less frequent purchase. It is due to the fact that different groups of customers will usually be adopting the product in different contexts, to accomplish different things, in comparison with different alternatives. Customers will have a very different Willingness-to-Pay for a Segway if they are adopting it as an alternative to a delivery van than they will if they are adopting it as an alternative to a skateboard. Customers who are considering installing a microchip in the desktop computers they are manufacturing will have a very different Willingness-to-Pay than customers who are considering installing the same chip in cell phones. As the price of a product becomes lower, it will typically be used for completely different purposes than it was used for when its price was higher. Each new use will tend to create a new step or plateau in the Willingness-to-Pay curve.



The Willingness-to-Pay curve will usually reveal situations where a relatively large change in a set price will have little effect on the quantity sold. This is the case whenever there is a large difference between the levels of successive steps or plateaus. The manufacturers of Segways may find themselves in this sort of situation after they have already begun to sell well to businesses, public services, and wealthy enthusiasts. At this point, it is possible that

lowering the price of a Segway by several hundred dollars will still leave it well above the Willingness-to-Pay of the next large category of customers. Hence, that particular price change might affect sales hardly at all.

The Willingness-to-Pay curve will also usually reveal situations where a relatively small change in a set price will have a huge effect on the quantity sold. This is the case whenever the current price is near one of the plateaus in the Willingness-to-Pay curve. Often the plateau will be defined by a competing product. In the case of the Segway, for example, one of the competing products will probably be a kind of motorized mail cart that is already on the market. As the price of a Segway moves above or below the price of this kind of mail cart, the quantity of Segways sold will probably undergo a sudden decrease or increase. The relevant Willingness-to-Pay plateau for the Segway will not be exactly the same as the price of the competing mail cart, because the benefits of each product will not be exactly the same. But there will still be a Willingness-to-Pay plateau for the Segway in that vicinity, and its effect on Segway sales will be considerable.

The things that can sometimes make the complexity of the contours in a Willingness-to-Pay curve less obvious do not make these contours less important. The step-like shape of Willingness-to-Pay curves becomes less conspicuous when the products are less innovative, less distinctive, or constitute a smaller or more frequent purchase. This is because, in these cases, there will usually be more gradations of alternatives and more gradations of uses. This is often the situation when a product has been around for a while, has gone through a large drop in price, is available in numerous variations, and is now used constantly in a wide range of contexts. Mass-produced consumer goods, whether desktop computers or toothbrushes, tend to fall in this category. But these are the very situations where *small* variations in the contours of a Willingness-to-Pay curve matter most. Volumes in these cases are typically high, and margins are typically low. Fine tuning the pricing decisions in these cases can have a huge impact. Not surprisingly, a large portion of the recent smart pricing activity has been in the grocery and mass retail industries.

There are two ways to determine the levels and contours of a Willingness-to-Pay curve. The first way is to look inside the customer's operations to see what the alternatives are and how they would affect the customer's balance sheet. For business customers, the result will generally be a set of hard numbers. It should be possible, for example, to calculate with considerable accuracy how much each Segway is worth to the U.S. Postal Service if the device is going to be used for deliveries in an urban area. It should also be possible to calculate how many Segways the Postal Service would need for that application. Surveying other possible customers and doing similar calculations for them should result in a clear picture of the Willingness-to-Pay curve for the higher end Segway customers. The Willingness-to-Pay for private customers isn't such a hard set of numbers, because it will depend more on subjective judgments. But it can still be estimated by looking at the alternative expenditures to which a Segway purchase will be compared.¹

The second way to determine the levels and contours of a Willingness-to-Pay curve is to do the experiment of offering the product at various prices. Every time a product is offered to a different group of customers at a given price, the acceptance rate indicates what portion of those customers has a Willingness-to-Pay higher than the designated price.² Bargaining situations, while more time consuming, can be used to locate the Willingness-to-Pay of a single customer with even greater precision. In practice, this sort of experimental test of the Willingness-to-Pay curve is generally necessary even when it is possible to take an inside look at the customer's operations.

Both these ways of determining the levels and contours of the Willingness-to-Pay curve are greatly facilitated by the internet. The internet makes it much easier to collect information on the operations of potential customers and to determine what alternatives are available to them. The internet also makes it much easier to conduct pricing experiments, offering special prices to select populations or engaging in bargaining processes with single customers.³ Over the last four or five years, this has helped convert Willingness-to-Pay curves from theoretical constructs to practical pricing tools.

Every observation about customers needs to be translatable into an observation about the level and shape of the Willingness-to-Pay curve in order to be directly relevant to pricing decisions. When it comes to pricing, "early adopters," for example, are generally customers with a very high Willingness-to-Pay. If there is a "gap" or "chasm" between the early adopters and the next group of customers, this essentially means there is a big difference between the first step or plateau on the Willingness-to-Pay curve for that product and the next step or plateau.


The Opportunity Cost Curve

The Opportunity Cost curve is the part of the pricing picture that tends to receive the least attention. Managers refer casually to "the unit cost" without paying enough attention to the way that cost will vary with the quantity of product supplied. Managers are also regularly tempted to use a cost estimate that doesn't include all the resources the company employed in production. Equipment that has already been depreciated, personnel budgeted to other departments, and a host of other assets are regularly left off production cost estimates. If these resources could otherwise have been put to some other use, they are "sacrificed opportunities," and it is important to charge them to the production cost of the product.


To appreciate the full implications of the Opportunity Cost curve in pricing decisions, it is necessary for companies to discuss it more explicitly. Like Willingness-to-Pay, the Opportunity Cost is an "indifference point" or "break even point," although this time, for the supplier. In fact, to be more symmetrical, we could call the Opportunity Cost the supplier's "Willingness-to-Sell." But the Opportunity Cost curve that is relevant to pricing decisions doesn't represent a whole population the way the Willingness-to-Pay curve does. It only represents the supply costs associated with one company.

The most important thing to realize about Opportunity Cost curves is that they are just as interestingly shaped as Willingness-to-Pay curves. In fact, an Opportunity Cost curve is even *less* like the graph of “supply” drawn in undergraduate economics classes than the Willingness-to-Pay curve is like the graph of “demand.” Instead of running flat or tilting in the opposite direction from the Willingness-to-Pay curve, an Opportunity Cost curve is more likely to resemble a series of deepening wave troughs with sharp peaks in between.

This is because the expenditures necessary to produce increasing quantities of a produce do not increase in a steady way. After relatively high initial investments in new production facilities and worker training, the cost per unit of product will tend to fall with increasing economies of scale and with the effects of organizational learning. (The Economies of Scale curve.) This is what will almost certainly happen with the manufacture of Segways, for example, as the first large-scale factory is set up and begins running increasingly smoothly.



Economies of Scale

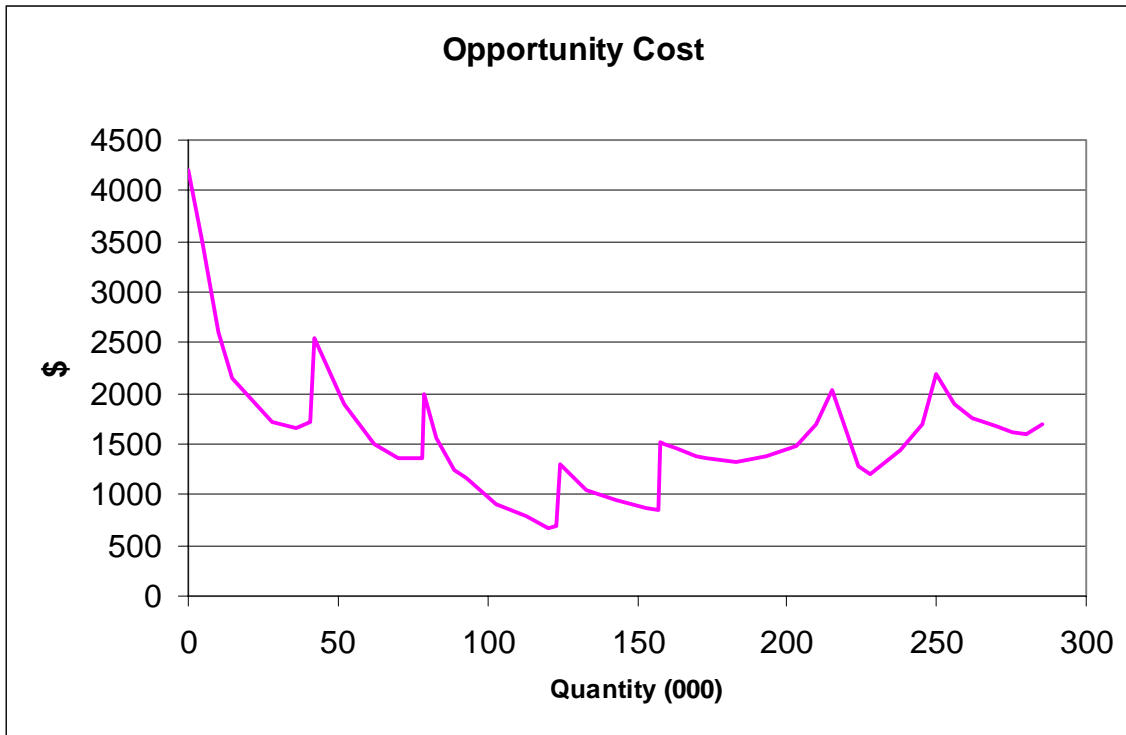


Increasing Utilization

But as soon as the utilization of the factory and associated facilities reaches a certain level—often about 80% capacity—cost per unit will once again begin to climb. (The Increasing Utilization curve.) The biggest reason for these rising costs is that the total supply chain, including the factory, suppliers, and distribution apparatus, will become less and less able to cope with variations in the level of production. Even with extensive sharing of information, small fluctuations in the inputs or outputs at one point in the system will tend to produce large fluctuations elsewhere, a phenomenon known as the *Bullwhip Effect*.⁴ Queue lengths will grow, lead times will increase, and managers will find themselves constantly firefighting and expediting. Additional expenditures will then be necessary to make up the shortages, to discount the surpluses, and to manage the growing lead times.

Eventually, the whole system will become overloaded, and new production lines or facilities will need to be built and staffed. This will produce a sharp upward spike in costs, starting the pattern all over again. Hence, the upward spikes in the complete Opportunity Cost curve represent the purchase of new equipment or factories, followed by their more efficient use as yet more customers are added. Eventually, the entire curve begins to turn upward as

management systems, supply chain partners, logistics providers, and other resources face excessively high utilization levels.



Opportunity Cost curves can only be reliably determined by actually adding up the necessary expenditures for supplying customers under different circumstances and with different lead times. This kind of calculation can incorporate many hypothetical numbers and still remain useful. But it is not possible to get an accurate idea of how Opportunity Costs will change with changes in quantity by simply taking the current costs and extrapolating from them.⁵

Like Willingness-to-Pay curves, Opportunity Cost curves have become a much more useful tool since the advent of the internet. The closer relationships with suppliers and greater sharing of information that the internet facilitates have made it possible to estimate Opportunity Costs under different circumstances with much greater precision.

The Way Effective Pricing Depends on Both Curves

The profits a company can make by producing a product depend *equally* on the Willingness-to-Pay curve and the Opportunity Cost curve. In fact, if a company could charge each customer the maximum that customer would be willing to pay, the profit would be that customer's Willingness-to-Pay minus the company's Opportunity Cost for supplying that customer. But in that situation, the customer would be capturing no value. Most of the time,

the value created is divided between the supplier and the customer. The price can be anywhere between the customer's Willingness-to-Pay and the supplier's Opportunity Cost without anyone losing on the deal. The exact level at which the price is set determines what portion of the value being created is captured by the supplier and what portion is captured by the customer.⁶

It is not in a company's interest to capture as much as possible of the value created. Instead, the better policy is to let the customers capture enough value so that there are more of them. This allows more value to be created, so that a smaller share of it results in larger total profits. It also makes it harder for a competitor to lure the customers away.


When the Willingness-to-Pay curve and the Opportunity Cost curve are put together on the same graph, they provide an easy way to visualize the effect of any possible price. For any given set price, there is a simple, easy-to-visualize rectangle that represents the resulting profits.

Setting a price for a product draws a horizontal line across the graph at that level. If the company is going to offer the same price to all customers, it can expect to sell its product to all the customers whose Willingness-to-Pay is above that level. Hence, the set price determines what portion of the possible customers will be supplied.

Because the company won't be selling to any customers with a lower Willingness-to-Pay, setting a price also draws a vertical line through the point where the price level and the Willingness-to-Pay curve intersect. This vertical line marks the total quantity that will be sold.

The Opportunity Cost for supplying the customers is indicated by the level of the Opportunity Cost curve at that quantity. (I.e., where the vertical line marking the total quantity that will be sold intersects the Opportunity Cost curve.) This Opportunity Cost can be represented by another horizontal line.

The profits from supplying the customers will be the set price, which those customers will be paying, minus the Opportunity Cost of supplying them. Hence, the total profits will be represented by the rectangular area bounded by the price level above, the Opportunity Cost level below, the zero-quantity axis on the left, and the total quantity line on the right. This might sound complicated, when stated in words or in mathematical formulas, but it is easy to grasp visually by looking at the relevant area on the graph.



The power of pricing is easy to demonstrate by simply looking at what happens to profits as the set price level is raised or lowered. Most managers would expect profits to increase gradually as the set price approaches an optimum level and then to decrease gradually as the set price moves away from this level. But much of the time, this is not the relationship between set prices and profits. Instead, as the set price changes, the profits will abruptly increase or decrease by varying amounts. Furthermore, there will typically be more than one peak to the profit curve.


If the set price level is just above a plateau on the Willingness-to-Pay curve, lowering the set price by a small amount will dramatically increase the quantity sold. This was what happened when Ford lowered the prices on its Explorers and other higher end vehicles. It is what will happen to Segways each time a price drop makes them a viable alternative for a new category of use.

But this increase in quantity sold won't usually translate into increased profits unless the Opportunity Cost curve is dropping significantly at the same time the Willingness-to-Pay curve is extending to a new plateau. Ford benefited from its increased sales of higher end vehicles, because it had the extra manufacturing capacity that allowed it to produce those vehicles at a decreasing cost. Segway's manufacturer will benefit from its increasing sales whenever the production increase will allow it to gain economies of scale.

This kind of relationship between the two curves is *not* something that can be taken for granted. Often the Opportunity Cost curve will be rising toward a peak just as the Willingness-

to-Pay curve is extending to a plateau. Lowering the set price slightly in this situation could dramatically increase sales—while *decreasing* profits.

Sometimes, when the set price level is near the bottom of a step in the Willingness-to-Pay curve that is relatively steep and fairly high, the best move is to *increase* the set price by a substantial amount. The company producing Segways may find itself in this position if the difference in Willingness-to-Pay between industrial and private customers is greater than anticipated. If this happens and lowering the set price for Segways does not attract the expected private customers, it may be better to raise the set price for a time (providing improvements to make the increase more palatable) in order to capture more of the value the Segway will have for its commercial and government users. In this kind of situation, the quantity sold might go down slightly when the set price is increased, but the profits should increase substantially.



Using This Pricing Picture as a Guide to Segmenting Markets

Apart from picking a more favorable set price, the obvious way to increase the value captured is to segment the market in some way. From a pricing standpoint, segmenting the market is equivalent to picking more than one price level. This is what companies are accomplishing by producing variations on the same basic product that can be priced differently. Offering discounts, coupons, and other special prices for different customer groups is another way to segment the market. But these segmenting tactics still leave the problem of how many

price levels of product to introduce and how to price each product variation for the relevant customer segment.

Here the combination of a Willingness-to-Pay curve with the corresponding Opportunity Cost curve is once again a powerful tool. The Opportunity Cost curve has to be carefully recalculated for each segmentation tactic that being considered. But once this is done, the market segmenting problem is roughly the same as the pricing problem, except that multiple prices can be chosen to produce a better result. In practice, a company will want to choose the fewest prices that will still capture most of the areas where the Willingness-to-Pay curve and the Opportunity Cost curve are farthest apart. If the locations and shapes of the two curves are fairly clear, this is not an especially difficult problem.

How the Two Curves Are Likely to Change Over Time

Because the particular shapes of both the Willingness-to-Pay curve and the Opportunity Cost curve are so important, it is vital to be aware of how the shapes of each curve are likely to change over time. Some changes in each curve can occur very quickly. Many of these changes will result from the pricing decisions themselves. Hence, the relevant curves in any pricing decision are not the ones that exist *before* the price is changed, but the ones that will exist *after*.

Competitors will tend to match a lower price whenever their Opportunity Costs permit. This will often have the effect of lowering a Willingness-to-Pay plateau just as a new price is introduced to take advantage of it. Hence, any assessment of the opportunities provided by the two curves will need to take account of projected moves by competitors.

Raising or lowering a price too abruptly can cause Opportunity Costs to rise sharply. This is especially conspicuous when companies introduce a temporary price drop in the form of a “special offer.” The sudden increase in the output required from the supply chain will typically make special expenditures necessary to fill the orders, or greatly increase the time necessary to fill them, or both. A supply chain with several stages will tend to multiply these special expenditures and increased lead times, as each stage over- or under-compensates, after finding itself with shortages to make up or with surpluses to warehouse. If the orders then drop off as a result of the demand being temporarily satisfied by the burst of extra orders, the oscillation in the supply chain could increase further. Thus, supply chain mechanisms can make a small change unexpectedly costly if the change is sudden enough.

Equally important, changing the price of one product will often affect the customers’ Willingnesses-to-Pay for related products. This is because products are used as parts of product systems, rather than as isolated tools. Lowering the price on Segways, for example, might increase people’s Willingnesses-to-Pay for the sorts of doorways and ramps that can easily accommodate Segways. This could lead to a temporary increase in the prices for those items, which, in turn, could temporarily reduce people’s Willingness-to-Pay for Segways. In addition to its implications for companies producing more than one related product, this sort of mechanism can affect the shape of the original Willingness-to-Pay curve considerably.

Applying More Dynamic Pricing Systems

Dynamic pricing systems can be a powerful tool for managing the pricing problem, but they don't change its essentials. There is still a population of potential customers with varying Willingnesses-to-Pay, and there is still a supply chain with Opportunity Costs that vary with the quantity of product. What's different is that instead of using a *handful* of set prices that are left constant for weeks at a time, dynamic pricing uses *many* prices that are changed from day to day, hour to hour, or even sale to sale. This means that the market is divided into many more segments, and that prices can be adjusted more frequently to take account of changes in the two relevant curves. But it's important to recognize that dynamic pricing doesn't do more than this.

To benefit from dynamic pricing, it is necessary to keep the customers with higher Willingnesses-to-Pay from going for the lowest price offered across many markets and over a longer stretch of time, thus making it into a *de facto* set price. One technique for accomplishing this is to establish different channels for various price levels and to vary the mix between them. Airlines apply this technique with their different channels for selling tickets, and groceries do it with their different channels for distributing coupons. Another technique is to customize the product on a customer-by-customer basis. This is the technique employed by build-to-order businesses like Dell Computer. Both of these strategies reduce the customer's ability to access the prices intended for other, lower Willingness-to-Pay customers.

But customers will *always* find ways to access the lower prices to some extent. In other words, no matter how much channels are customized, customers will find ways to reassemble themselves into 'market segments,' because it's in their interest to do so. If all else fails, customers will have other customers make their purchases for them, or they will form buyers' clubs to bargain for a collective, lower price. Customized products can only counteract this tendency to a limited degree, because customers will eventually find ways to group products into a limited number of categories for purposes of comparison. This means that businesses will always be dealing with a market that has organized itself into segments, even if the business has not organized the market into segments itself. Any segmented market, even if it has a great many segments that are adjusted rapidly, still presents a variation of the essential pricing picture.

Prospects for the Automated Management of Pricing

The analysis of the pricing problem presented here suggests that managing it effectively in practice will require both automated systems and close human oversight. Automated systems are necessary for two major reasons. One reason is that there are simply too many Willingness-to-Pay curves and Opportunity Cost curves for managers to keep track of. For every region where a product can be offered at different prices and has different costs, there are potentially different Willingness-to-Pay curves and Opportunity Cost curves. The other reason is that the

effects of one pricing decision on other pricing decisions can rapidly become too complex to estimate without automated systems.

Meanwhile, close human oversight will remain indispensable for a long time to come. This is because there are so many parts of the pricing problem where intuitive judgments by human beings need to be taken into account and where human actions can creatively change what is going on. Intuitive judgments and creative actions on the part of humans are not things that automated systems have so far been very successful in anticipating or even recognizing. Hence, until there is some dramatic new breakthrough, verifying the “reasonableness” of pricing recommendations from automated systems will require the sort of overall understanding outlined in this article.

¹ For more on Willingness-to-Pay curves and shopper behavior, see W. W. Moe and P. S. Fader, “Uncovering Patterns in Cybershopping,” *California Management Review* 43, no.4 (2001), 106-117.

² Special procedures are necessary for analyzing this data. The common statistical practice of summarizing “demand data” by means of a simple linear regression is extremely ill advised, because it obscures the very thing that is most valuable to know about a Willingness-to-Pay curve: its step-like structure.

³ Walter Baker, Mike Marn, and Craig Zawada, “Price Smarter on the Net,” *Harvard Business Review* (February 2001). Also P. K. Kannan and P. K. Kopalle, “Dynamic Pricing on the Internet: Importance and Implications for Consumer Behavior,” *International Journal of Electronic Commerce* 5, no. 3 (2001), 63-83.

⁴ For further discussion of the Bullwhip Effect, see H. Lee, P. Padmanabhan, and S. Whang., “The Bullwhip Effect in Supply Chains,” *Sloan Management Review* 38, no. 3 (1997), 93-102.

⁵ Fortunately, new software is available to help, such as the factory modeling software from companies like i2 Technologies and from Manugistics.

⁶ Scott Borg, Adam Brandenburger, and Harborne Stuart, “An Introduction to Business-Centered Economics” (Foundations Course Guide, distributed by Harvard Business School), 1996. Also Harborne Stuart and Adam Brandenburger, “Value-Based Business Strategy,” *Journal of Economics and Management Strategy* 5, no. 1 (1996).