11 CUSTOMIZED PRICING

11.1 BACKGROUND AND BUSINESS SETTING

All of the pricing tactics we have studied so far assume that buyers and sellers interact in the same basic fashion. The seller offers a stock of goods to different customer segments through different channels. Potential customers arrive, observe a price, and decide whether or not to purchase. The seller needs to decide what price to offer for each product to each customer group through each channel. The seller monitors sales of his goods and periodically updates the prices and/or availabilities he is offering. This is the *list-pricing* model. All of the pricing settings we have covered so far, including peak load pricing, revenue management, markdown management, and dynamic list pricing, are varieties of list pricing.

List pricing may seem quite general. However, there is another important class of pricing—*customized pricing*—that is of equal or greater importance in many industries. In customized pricing, potential customers approach the seller, one by one, and the seller quotes each one a price. Often (but not always) each potential buyer wants something different—either a different bundle of products or services or a different quantity or a different variation on a basic product. In customized pricing, the seller can quote a different price for each request. The price can be based on knowledge of the individual buyer, the product(s) requested by the buyer, and other factors, such as general market conditions and competitive offerings. We now describe some examples of customized pricing.

11.1.2 Unsecured Consumer Loans

In Great Britain a common way for people to borrow money to renovate their houses, buy a used car, go on vacation, or consolidate credit card debt is by obtaining an unsecured consumer loan. Typical amounts for unsecured consumer loans range from £500 to £20,000, and typical terms are one to seven years, with larger loans typically having longer terms. Unsecured loans are offered not only by traditional banks, such as Barclays and NatWest, but also by retailers, such as Sainsbury's (a major grocery chain), and specialized Internet lenders, such as egg.co.uk. The market has many competitors, with more than 20 major nationwide providers in addition to many regional lenders. Typical annual percentage rates (APRs) and corresponding monthly payments for a five-year loan of £3,000 as advertised by different lenders are shown in Table 11.2.

TABLE 11.2

Interest rates (APRs), monthly payments, and total interest payments for a 5-year unsecured personal loan of £3,000 as advertised by various lenders in June 2003

Provider	APR (%)	Monthly payment (£)	Total interest (£)
Northern Rock	6.9	92.49	329.79
Nationwide	7.9	93.87	379.35
AA Loan	10.9	98.07	530.67
Egg	10.9	98.07	530.67
Sainsbury's Bank	11.5	98.93	561.41
American Express Bank	11.7	99.21	571.69
MBNA Express Loan	11.9	99.50	581.99
Lombard Direct	11.9	99.50	581.99
Woolwich	11.9	99.50	581.99
Virgin	11.9	99.50	581.99
Alliance and Leicester	12.5	100.36	612.99
Barclays	12.9	100.94	633.75
Halifax	13.9	102.39	685.94
HSBC	14.9	103.85	738.57
NatWest Bank	15.9	105.32	791.63
Lloyd's TSB	17.7	108.01	888.22

SOURCE: *Money Facts*, July 2003.

The sales process for the Bank of Albion, a typical (albeit mythical) provider of unsecured loans is shown in Figure 11.1. Potential customers enter the process through one of four channels:

- *Branch:* entering a branch and inquiring about a loan
- *Direct:* calling an 800 number
- *Internet:* Inquiring about a loan through the bank's Web site
- *Agent:* referred by a third party, such as a real estate broker

Regardless of channel, a potential customer is asked how much she wants to borrow over what term and a brief series of questions to establish name, age, and time at current address. The bank then quotes an APR and monthly payment. About 15% of the customers drop out at this point ("lost quotes"), while the remainder go on to fill out a more extensive

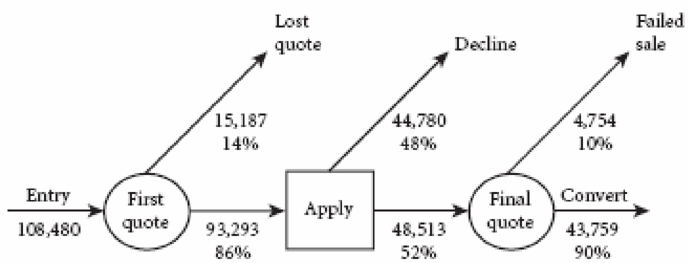

Figure 11.1 Consumer loan sales and pricing process at the Bank of Albion.

application that includes additional questions about such items as employer, income, expenses, and other loans outstanding. This information enables the bank to assign a credit score to each applicant. About 50% of applications are rejected by the bank as too risky. For the accepted applications, the bank has the opportunity to adjust the APR—for example, it can offer a higher APR to high-risk customers if it chooses. About 95% of customers whose applications are accepted go on to take a loan.

The bank faces two pricing decisions: what APR to offer in the initial quote, and what adjusted APR to offer to customers whose applications have been accepted.[1] This is a customized-pricing problem because the bank can choose different APRs for different customers based on type of loan, channel, and market segment.

11.1.3 Other Customized-Pricing Services

In addition to the foregoing two examples of customized pricing—one business to business and one business to consumer—there are some others.

- Businesses purchasing trucking services usually solicit bids from a number of competing carriers, such as Roadway and Schneider. Each carrier will submit a bid detailing a price schedule at which it will carry freight for the customer for some period of time, usually a year. Each carrier faces the problem of determining what price schedule to bid.
- A metropolitan police department wishes to purchase 50 police cars. It solicits bids from Ford, General Motors, and Daimler-Chrysler. The bid solicitation describes the performance and options required. Each of the three manufacturers needs to determine what price to bid.
- A hospital has decided to purchase a magnetic resonance imaging (MRI) machine to improve its diagnostic capabilities. It sends copies of a request for proposal (RFP) to the five leading MRI manufacturing companies. The RFP specifies the required capability as well as the criteria the hospital will choose to determine the winning bidder. Each of the manufacturing companies needs to determine what price to bid.
- A customer applies to his local bank for a mortgage. Based on the customer's credit history, the amount of the loan, and the characteristics of the house he is planning to purchase, the bank must determine whether or not to offer him a mortgage and, if so, at what rate.

In each of these customized pricing situations the potential customer has approached the seller. Since the customer has already signaled his willingness to purchase, demand induction is usually not the major issue. Instead, the major issue in customized pricing is usually not *if* the customer will purchase but *who* the customer will purchase from. This is one important difference between customized pricing and list pricing. Here are two others.

1. *In customized pricing, the seller can quote a different price to each customer*. The customized price can (and should) reflect the best information the seller has at that time. Of course, the seller will usually not have perfect information about the buyer's preferences or about the competition. But, as we will see, he can use statistical reasoning to improve his expected profitability.
2. *In customized pricing, the seller can track lost business*. In a customized-pricing situation, the seller will either win—i.e., get the business—or lose. In many (but not all) situations, a losing seller can find out which of his competitors won the business. However, even when this information is not available, the buyer knows that a potential customer has inquired about his product and decided not to purchase. This is in contrast to list pricing, in which the seller can only observe how many units he sold—not how many customers considered buying his product but decided against it. This *lost bid* information is important because it forms the basis for statistical estimation of customer bid-response functions.

The existence of a *bid* or a *price quote* is a common element in customized pricing. In business-to-business settings, the quote is often in response to a request for proposal (RFP) or a request for quote (RFQ) that has been issued by a buyer. In business-to-consumer situations, such as unsecured consumer loans and mortgages, the price quote often occurs after a customer has made a phone call or filled out an application. In any case, the price quote occurs only after the seller knows something about the buyer and what she wants. This contrasts with list pricing, in which potential customers find prices by accessing a price list (online or in a catalog) or reading a label in a store. The prices need to be set before the seller knows the identity of any individual potential customer.

List pricing and customized pricing often coexist, in the sense that a seller will use a list price to set a "ceiling," while most customers are quoted a price at a discount from the list. For example, the *manufacturer's suggested retail price* (MSRP) is usually an upper limit on the price for, say, a new car, with the selling price ultimately determined by some combination of the customer's ability to negotiate and the dealer's willingness to sell at a lower price.[2] In a similar vein, shipping companies maintain voluminous tariff schedules that specify list prices for shipping different types of freight between every zip code pair in the United States. However, only a small fraction of sales (usually less than 5%) takes place at list price. The rest moves at discounted rates negotiated with individual customers.

In the remainder of this chapter, we formulate the customized-pricing decision as an optimization problem and show how that problem can be solved to maximize the expected contribution margin from each bid. A key element in the customized pricing-problem is the *bid-response curve*, which specifies the seller's expectation on how each customer will respond to his bid price. We show how price-response curves can be estimated for different customers seeking to purchase different products. We then show how the customized-pricing model can be enhanced to deal with other pricing settings as well as objectives other than maximizing expected contribution. Finally we discuss how these computations fit into a general process for customized pricing.

11.2 CALCULATING OPTIMAL CUSTOMIZED PRICES

To motivate the formulation of the customized-pricing problem, we put ourself in the place of a seller who has been asked to bid on a potential deal. By *deal* we mean a particular piece of business, which may be a single order or a contract to provide future products and services. By *bid* we mean the price we offer.[3] This is a simplified setting since it assumes the buyer has dictated all elements of the deal and that the only decision on the part of the bidder is what price to bid. In the real world, a bidder may also have to decide on other elements in his bid, such as delivery date, contract terms and conditions, and even product functionality. Furthermore, we assume initially that the bidder is submitting a single take-it-or-leave-it bid. Of course, in many situations, not only price but every element of the deal may well be on the table.

The first thing we need to determine is our goal for this bid. That is, what is our objective function? For now, we assume that we want to maximize expected contribution from the bid. For a given price, p, expected contribution can be expressed as

$$\text{Expected contribution at price } p = (\text{Deal contribution at } p) \times (\text{Probability of winning bid at } p) \qquad (11.1)$$

The first term on the right-hand side of Equation 11.1 is the *deal contribution*—the margin we will realize if we win the bid. If the bid is to purchase d units of an item with unit cost of c, then our deal contribution would be $(p - c)d$. The calculation of deal contribution would be more complicated if we were bidding to supply a complex bundle of products and services, each of which has a different unit cost, or if we were negotiating a contract for future services in which the volume and mix of services to be supplied were uncertain.

The second term on the right-hand side of Equation 11.1 is the probability that we will win the deal at a given price. Our uncertainty about winning at a given price derives from two sources.

- *Competitive uncertainty*. We do not know what our competitors will bid. In fact, in many cases we may not know the identity or even the number of competitors we face.
- *Preference uncertainty*. We do not know exactly what criteria (both conscious and unconscious) the buyer will use to evaluate competitive bids. Nor do we know for sure how the buyer values our brand and our product/service offerings relative to the competition.

In most situations both preference uncertainty and competitive uncertainty will be present: We won't know exactly what our competitors are going to bid, nor will we know exactly how the buyer will choose among competing bids. Of course, if we knew what our competitors were bidding and how the buyer would choose among bids, our problem would be relatively simple—we would submit the most profitable bid that would win. However, in the real world, bidders rarely, if ever, have access to that level of information.[4]

There is, however, one important situation in which preference uncertainty is absent—when we know that the buyer is going to choose the lowest bid. This is the case, for example, when a government agency is required by law to choose the lowest-cost bidder. In these situations, the agency often issues extremely detailed specifications for the goods to be purchased so that price is the only remaining issue on which selection is made. Preference uncertainty can also be absent in supplier auctions, either online or offline. In a supplier auction the buyer initially selects a set of acceptable suppliers and commits to purchase from the lowest bidder. In most seller auctions each supplier has the opportunity to observe the bids of other suppliers and then to submit a lower bid if he desires to do so. This ability to rebid based on the bids of other suppliers injects additional complexity into the pricing process that we will not address here.[5]

11.2.1 Single-Competitor Model

Imagine you are an auto manufacturer bidding against a single competitor to sell 50 pickup trucks to a county park district. The park district will buy all 50 trucks from a single supplier and is committed by law to pick the lowest bidder. Each supplier has been asked to submit a single sealed bid. The bids are final, and the lower of the two bids will win. Your production cost per truck is \$10,000. Based on past experience, your belief about what your competitor will bid can be described by a uniform distribution between \$9,000 and \$14,000, as shown in Figure 11.2. What should you bid to maximize expected profitability?

Call your bid p and the competing bid q. You will win if your bid is less than the competing bid, that is, if $p < q$. (For simplicity, we ignore the possibility of a tie.) Let $\rho(p)$ be the probability that you would win if you bid a price p. Then $\rho(p) = 1 - F(p)$, where F is the cumulative distribution function of the competing bid—that is, $F(x)$ is the probability that

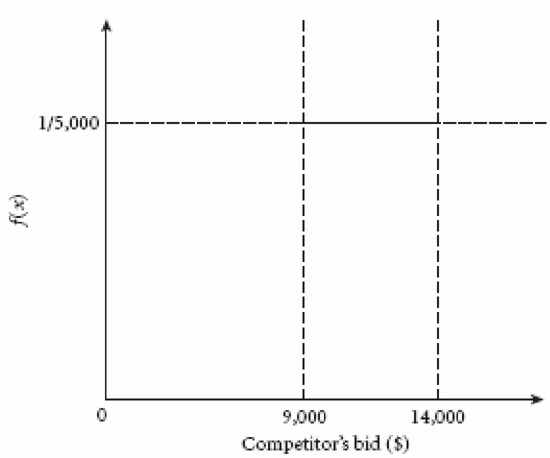

Figure 11.2 Uniform probability distribution on a competitor's bid.

the competing bid will be less than x. Your probability of winning the bid as a function of price is then as shown in Figure 11.3. If you bid below \$9,000, then you will win for sure, since you know that the competitor will bid more than \$9,000. Your probability of winning the bid decreases linearly between \$9,000 and \$14,000 and is zero if you bid above \$14,000.

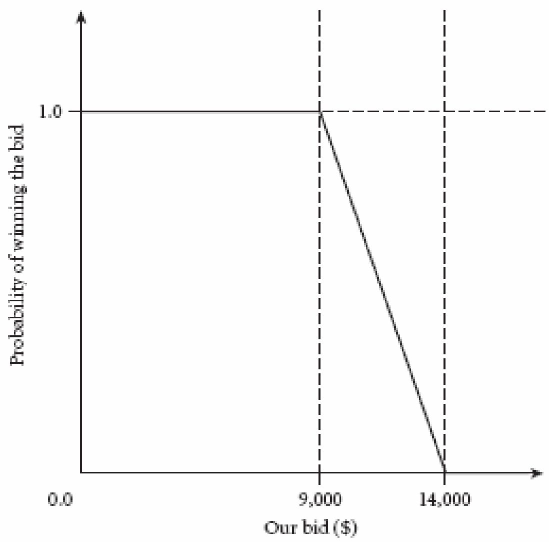

Figure 11.3 Probability of winning the bid as a function of price when our belief about the competitor's bid follows the uniform distribution in Figure 11.2.

We call $\rho(p)$ the *bid-response function* for this deal. For each deal, the bid-response function specifies the probability of winning the deal as a function of price bid. It is the customized-pricing analog of the price-response curve introduced in Chapter 2. Where the price-response curve specifies total expected demand as a function of list price, the bid-response curve specifies the probability of winning an individual bid as a function of bid price. Like the price-response curve, the bid-response function is downward sloping; that is, the probability of winning the bid decreases as the bid price increases. The bid-response probability is less than or equal to 1 for all prices and decreases to 0 at high bids. For the example, the bid-response curve is given by

$$\rho(p) = 2.8 - p/5{,}000 \tag{11.2}$$

and the price that maximizes expected profitability can be found by solving

$$\text{Maximize} \quad 50(2.8 - p/5{,}000)(p - 10{,}000) \tag{11.3}$$

Using the techniques from Section 4.4, we find that the objective function in Equation 11.3 is maximized at $p^* = \$12{,}000$. At this price, your probability of winning the bid is $2.8 - 12{,}000/5{,}000 = 0.4$. The margin per unit if the deal is won is $\$12{,}000 - \$10{,}000 = \$2{,}000$ and expected total margin is $40\% \times 50 \times \$2{,}000 = \$40{,}000$.

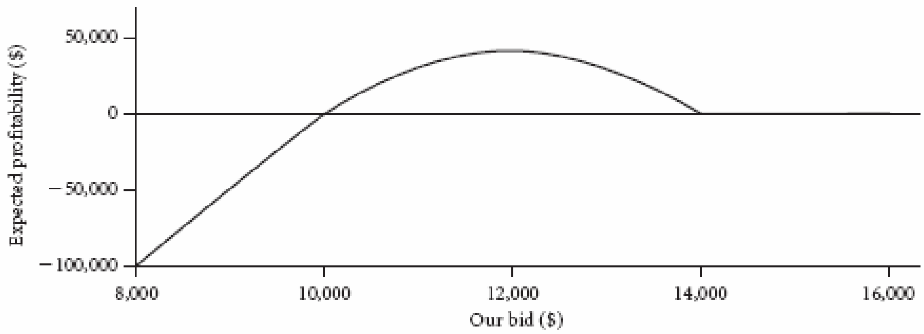

Figure 11.4 Expected profitability as a function of price bid in example.

Figure 11.4 shows how expected profit varies as a function of your bid. The expected profit function has the familiar hill shape. At a bid price below $9,000, you are certain to win the deal, but at a loss. Above $9,000, your probability of winning the bid decreases. However, you begin to gain some positive expected profit at prices above $10,000. Above $10,000, unit margin increases with increasing price, but this increase is counterbalanced by the decreasing chance of winning the deal. At $12,000, these two effects balance, and expected contribution is maximized. Above $12,000, decreasing chances of winning the deal overwhelm the increased margin. At any bid above $14,000, you are certain to lose the deal, so expected contribution is $0.

The optimization problem in Equation 11.2 is identical to the problem confronting a seller facing a deterministic price-response curve of $50(2.8 - p/5{,}000)^+$ and a unit cost of $10,000. In both cases, the optimal price is $12,000. However, there is an important difference between the two situations. A seller facing a deterministic price-response curve of $d(p) = 50(2.3 - p/5{,}000)^+$ and charging $12,000 per unit will sell exactly 20 units and will make a profit of $40,000. But in the bid-response case, by bidding $12,000 per unit, you will either win the bid and sell 50 units, for a total contribution of $100,000, or you will lose the bid and realize nothing. Changing price does not change the number of units you sell; it changes your probability of selling all the units. At the optimal price of $12,000, you have a 40% chance of winning the deal and making $100,000 and a 60% chance of losing the deal and making nothing.

11.2.2 The General Problem

We can generalize the example in a very straightforward fashion. Specifically, the *customized-pricing problem* for a particular deal is

$$\text{Maximize} \quad \tau(p) = \rho(p)m(p) \tag{11.4}$$

where $\tau(p)$ is expected contribution at price p, $\rho(p)$ is the bid-response function, and $m(p)$ is deal contribution at price p. For our purposes, $\rho(p)$ will be a continuous, downward-sloping function of p, and $m(p)$ will be a continuous, upward-sloping function of p. This means that $\tau(p)$ will be hill shaped, as shown in Figure 11.4, and there will be an optimal price p^* that maximizes expected contribution.

Assume, as in the example, that we are bidding for an order of d units and that our unit cost is c. In this case, $m(p) = d(p - c)$ and $\tau(p) = d\rho(p)(p - c)$. The value of p^* that maximizes $\tau(p)$ can be found by setting the derivative of $\tau(p)$ equal to 0. That is,

$$\tau'(p^*) = d\rho'(p^*)(p^* - c) + d\rho(p^*) = 0$$

which, after a bit of algebra, gives

$$-\frac{\rho'(p^*)p^*}{\rho(p^*)} = \frac{p^*}{p^* - c} \tag{11.5}$$

The expression on the left-hand side of Equation 11.5 is the *bid-response elasticity*, that is, the percentage change in probability of winning the bid that will result from a 1% change in price.[6] A bid-price elasticity of 2 means that a 1% *increase* in price will lead to a 2% *decrease* in the probability of winning the deal. The expression on the right-hand side of Equation 11.5 is the inverse of the contribution margin ratio. In other words, at the optimal customized price, the bid-response elasticity is equal to the inverse of the contribution margin ratio:

$$\text{Bid-response elasticity at optimal price} = 1/(\text{Contribution margin ratio})$$

This is the customized-pricing analog of the condition for optimal deterministic prices in Equation 3.21.

Note that using probabilities to represent our chances of winning a deal does not imply that either the buyer or our competition is making decisions by flipping a coin or using any other type of random process. Indeed, we would expect the competition to be setting the prices they will bid based on rational analysis of the opportunity—including his best information on what he believes *we* will bid. The uncertainty in the process arises entirely from our lack of information on the behavior of the other players in the bidding game: the buyer and the competition.

11.2.3 Multiple Competitors

We would naturally expect the addition of more competitors to reduce our probability of winning a deal. Otherwise, we would be willing to pay additional companies to compete against us—behavior that is rarely seen in real life. We can demonstrate the effect of multiple competitors using the simple example of the county park district. What if we faced two competitors instead of one bidding for the order of 50 pickup trucks? To simplify matters, let's assume we have identical beliefs about how both competitors will bid. That is, we believe that both competitor 1 and competitor 2 will price according to the uniform distribution shown in Figure 11.2.

We also assume that the two competing bids will be independent—that is, that information about one competitor's bid would not change our assessment of the bidding behavior of the other competitor. While the independence assumption may generally be reasonable, there are situations where it is not appropriate. One case would be when we believe that our two competitors are colluding to coordinate their bids. In that case, knowledge of how one of them is going to bid would obviously tell us quite a bit about how the other competitor planned to bid. The second case would be when the two competitors are not necessarily colluding but are basing their bids on joint information that we do not share. If both of the competitors share a common cost (say, a common supplier) that is unknown to us, learning what one of them planned to bid might enable us to estimate the common cost and therefore to come up with a better estimate of what the other competitor would bid.

We will win the deal in the example if and only if our bid is below the lower of the bids submitted by competitor 1 *and* competitor 2. For a given p, the probability that our bid will be lower than the bid from competitor 1 is $1 - F(p) = 2.8 - p/5{,}000$, which, by assumption, is also the probability that our bid will be lower than the bid from competitor 2. Since competitors 1 and 2 are submitting independent bids, the probability that p will be lower than both competing bids is $\rho(p) = (2.8 - p/5{,}000)^2$. This bid-response curve is shown in Figure 11.5. You can see that the addition of the second competitor decreases our chance of winning the bid for all realistic prices.

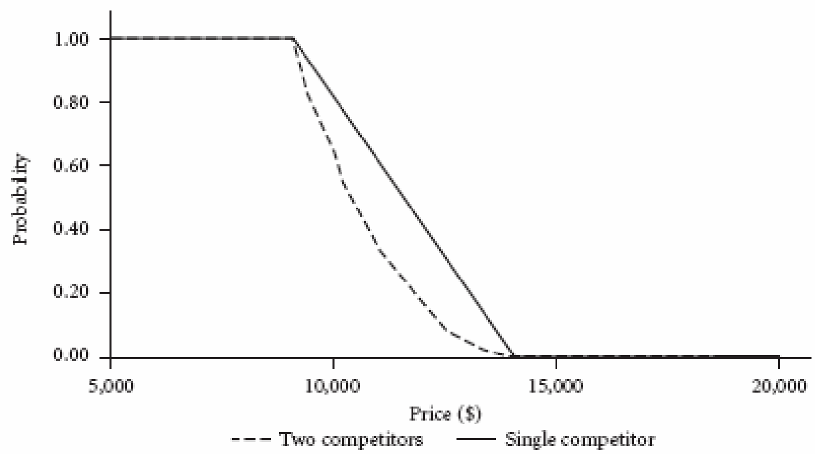

Figure 11.5 Our probability of winning the bid as a function of price with one competitor and with two identical competitors.

With two competitors, our expected contribution is given by

$$\tau(p) = 50(2.8 - p/5{,}000)^2(p - 10{,}000)$$

In this case, the price that maximizes expected revenue is $p^* = \$11{,}333.33$ per unit, with a corresponding probability of winning the bid of 28.4% and an expected total profit of \$18,963. Note that the additional competition creates pain for us in two ways: Our optimal bid is lower, so total contribution is lower if the deal won, *and* our probability of winning the deal is lower, even at our lower bid. The overall effect is dramatic: reducing our expected contribution from the deal by more than 50% from \$40,000 to less than \$19,000. This is a graphic illustration of the power of competition to reduce both the price and the profitability of a deal. Obviously, there is a strong incentive for the buyer to increase the number of bidders—or at least make the bidders believe there are more competitors!

As a matter of reference, the probability that a bid p is the lowest when there are n competitors submitting independent bids is given by

$$g(p) = [1 - F_1(p)] \times [1 - F_2(p)] \times \cdots \times [1 - F_n(p)] \tag{11.6}$$

where $F_i(x)$ is the cumulative distribution function on the price bid by competitor i—that is, the probability that competitor i will bid less that x.

Equation 11.6 gives the probability that our bid will be the lowest in any situation where competitors are submitting independent bids. If the buyer is choosing a winner based strictly on lowest bid, then this also gives our probability of winning at any price; that is, $\rho(p) = g(p)$. However, if the buyer is taking factors other than price into account when choosing a winner, then our probability of winning the bid will not be the same as our probability that we have submitted the lowest bid. We will treat this situation next.