



**COLUMBIA UNIVERSITY**

Center for Pricing and Revenue Management

WORKING PAPER SERIES: NO. 2010-1

---

## Customized Pricing

Robert Phillips,  
Columbia University and Nomis Solutions

2010

---

<http://www.cprm.columbia.edu>

# Customized Pricing

**Robert Phillips**

Columbia University Graduate School of Business

Nomis Solutions

[rp2051@columbia.edu](mailto:rp2051@columbia.edu)

Columbia Center for Pricing and Revenue Management

Working Paper 2010-1

October, 2010

## 1 Background and Introduction

Consider the following three pricing situations:

- A global telecommunications company sells its services in North America. The company sells more than 40 individual services including local, long-distance, data communications, and corporate network. Different customers request different levels and combinations of these services. The company sells its services on a contract basis, with the typical term of a contract being a year. An important segment of the company's business consists of customers with annual telecommunications bills of between \$10,000 and \$10 million. Customers in this segment are primarily large and medium-sized businesses, educational institutions, and various state and local government agencies. Prospective customers in this segment typically submit a Request for Quote (RFQ) to two or more telecommunications companies annually. Based on the bids received, they determine which company will be their service provider for the next year. Bids by competing telecommunications companies are usually specified as discounts from their standard list prices. The problem facing the company is what discount to quote in response to each RFQ (Phillips, 2005, pp. 264 - 265.).
- The Soundco Radio Company sells aftermarket radios and CD players for automobiles. Its customers are regional and national electronics retailers, automotive stereo catalog sellers,

and electronics wholesalers. The dealer base price for one of Soundco's most popular models, the CDR-2000 is \$109. However, the average price actually paid by customers in 2007 was \$67.50. The difference of \$41.50 was due to a variety of discounts, promotions, and concessions that Soundco made on different deals. For example, while the average price paid by customers for the CDR-2000 was \$67.50, there was a wide variation in actual price paid: about 18% of sales were at \$55 or less while 17.5% were at \$85 or more. Soundco wants to know if it is giving the right discounts to the right customers to maximize profitability and, if not, how it should change its discounting (Marn, et. al. ,2007, pp. 31-33.).

- A medical testing device company sells a gas chromatograph refill cartridge with a standard list price of \$11.85. The refill cartridge is typically ordered in batches. Orders for fewer than 200 units are handled through the company's website or through a reseller with no discount. Large orders (more than 1,000 units) are negotiated by a national account manager, usually as part of a large, bundled sale. Orders for 200 - 1,000 units are handled by a regional sales staff that has considerable freedom to set discounts. Analysis of historical data showed that discounts ranged from 0% to more than 30% for some deals. The company wants to determine the "right" levels of discount to offer for various deals (Agrawal and Ferguson 2007, pg. 220.).

The three cases described above are examples of *customized pricing*. Customized pricing is most common in business-to-business settings, although it is also found in consumer lending and insurance. Customized pricing is defined by three characteristics:

1. Customers approach a seller. Each buyer describes a product, or set of products that she would like to purchase. Often this is specified in a Request for Proposal (RFP) or Request for Quote (RFQ) or, in the case of insurance or credit, an application. We call such an inquiry a *customer request*.
2. The seller must decide how to respond to each customer request. One possibility is not to respond – that is, pass up the opportunity for the business. If a seller does respond, he needs to determine what price to quote for each customer request. At the point the price is quoted, the seller knows the product(s) requested by the buyer, the channel through which the buyer approached, and at least some information about the buyer. Often in business-to-business settings, price is quoted in terms of a discount from a list price or tariff.

3. The seller has some freedom to quote different prices for different customer requests based on the products in the request, the sales channel through which the request is received, and the characteristics of the buyer.

Each customer request can be characterized by the characteristics of the order (e.g. the number and configuration of products ordered), the channel through which the request was received (e.g. Internet versus direct sales), and the characteristics of the customer making the request (e.g. large manufacturer in the Northeast versus small distributor in the Southwest). The problem facing the seller at any time is what price to quote for the request currently under consideration. In the extreme, the seller could quote a different price for each request – this would approach the ideal of “one-to-one” pricing. More typically, the seller differentiates prices by a finite set of combinations of order characteristics, customer characteristics, and channel. For example a seller might define five order-size tiers, with increasing levels of discount by order-size. We use the term *pricing segment* to denote each combination of order characteristic, customer characteristic, and channel for which a seller can set a price.

This chapter focuses on the use of mathematical analysis by a seller to set and update the prices for a finite set of pricing segments. The idea of using mathematical analysis to optimize prices in this fashion is relatively new. Several authors such as Friedman (1956), Gates (1967) and Morin and Clough (1969) describe models in which price is the only criterion used to select a supplier. Talus Solutions was the first company to develop a software system to optimize customized prices (Boyd, et. al., 2005). The first detailed treatment of an optimization approach to customized pricing for segmented markets is Chapter 11 in Phillips (2005). Agrawal and Ferguson (2007) apply a similar analytical approach to two examples of what they call *customised-pricing bid-response models (CPBRMs)* and compare the performance of segmented and unsegmented approaches. Phillips (2010) describes the application of optimal customized pricing in the specific context of consumer credit pricing.

## 2 Business Application of Customized Pricing

A stylized view of the customized pricing sales process is shown in Figure 1. Prospective customers approach a seller one-by-one through different channels. Each customer describes her desired purchase to the seller. Based on the combination of the desired purchase, the channel, and the

(Figure 1 about here.)

Figure 1: A stylized view of the customized pricing sales process.

characteristics of the customer; the seller chooses the price to offer. Once the price is known, the customer decides whether or not to purchase from the seller. The seller knows the characteristics not only of the seller's own customers and their purchases, but also the characteristics of those potential customers who did not purchase from him – i.e., *lost sales*. Typically the majority of a seller's lost sales went on to purchase from a competitor although some portion of the lost sales may have decided not to purchase at all.

In many business-to-business markets, a purchase is often initiated when a buyer sends a Request for Proposal (RFP) or Request for a Quote (RFQ) to one or more competing suppliers describing the buyer's needs for goods or services. Each supplier receiving the RFP or RFQ needs to decide whether or not to bid for the business. A supplier may decide not to bid if they feel that they cannot satisfy the requirements of the buyer, or if they feel that their probability of making the sale is not high enough to justify bidding, or for other reasons. If a supplier does bid, he needs to choose a price. At the highest level, the tradeoff in determining the price is straightforward – if the seller sets the price high, he will make a high profit if he wins, but has a low probability of winning; if he sets the price low, he will make less profit if he wins, but he has a higher probability of winning. The price that maximizes his expected profit optimally balances these two effects.

Customized pricing is commonly used for products that are highly configured. Heavy trucks are a good example. The 2008 Model 388 Day Cab heavy truck sold by Peterbilt allows the customer to choose his desired options for 34 different components of the truck ranging from the rear-wheel mudflap hangers which can be straight, coiled or tubular to 15 different choices for rear axles to seven different transmissions to the choice between an aluminum or a steel battery box.<sup>1</sup> Multiplying the number of options in the Peterbilt list would imply that they could combined in  $75 \times 10^{16}$  possible ways – each corresponding to a different truck. In practice, not all combinations are feasible and the actual number of possible trucks is smaller – probably by two or three orders of magnitude. Each of the options has a list price and the sum of the list prices for the chosen options plus a base price is the “list price” for the configured truck. However, in the vast majority of cases, heavy trucks are sold at a discount from this list price. The customized pricing problem

---

<sup>1</sup>From the Peterbilt website at <http://www.peterbilt.com/trad388dc.4.aspx> as accessed in November, 2008.

is what discount to apply for each quote.

Highly configurable products such as heavy trucks lend themselves to customized pricing for a number of reasons. First of all, configurability usually implies the existence of a large number of potential products – as in the case of Peterbilt. One heavy truck manufacturer claimed in conversation that no two customers had ever independently ordered exactly the same truck. This means that there is a lot of potential to segment based on product dimensions. Secondly, configurable products and services are usually “big ticket” items that often involve considerable interaction between the buyer and the seller during the sales process. This provides an opportunity for customer-based price differentiation. Finally, options often vary widely in their profitability to the seller. This variation gives an additional motivation for price differentiation. In the case of Peterbilt, a truck built of highly profitable options provides much more scope for possible discounting than a less profitable configuration.

The stylization of customized pricing in Figure 1 should not be interpreted to imply that sellers are passively waiting for customer requests to arrive. On the contrary, the vast majority of sellers will be actively engaged in marketing their products and services through advertising and brand support as well as soliciting business through direct and indirect sales channels. These sales and marketing activities are often critical in generating business, however they largely (if not entirely) represent “sunk” costs by the time the final price is quoted. Sales and marketing activities influence the volume of customer requests that a seller receives and the willingness-to-pay of potential buyers. A company with strong sales and marketing will have more opportunities to bid and will be able to bid successfully at higher prices than a competitor with weaker sales and marketing. The goal of a customized pricing process should be to enable a supplier to determine the best price to bid for each customer request given the supplier’s strengths and weaknesses relative to the competition.

Customized pricing is most commonly associated with business-to-business markets: most consumer products and services are priced using list prices combined with various promotions. However, customized pricing is also common in consumer and small business loans and insurance. Both lenders and insurance companies typically require applications from prospective customers describing both the desired product (e.g. size and term of loan) as well as information about the prospective borrower themselves. Table 1 lists the information required for an on-line application

Name	Zip Code	Address
Currently insured?	Insured with GEICO in last 6 months?	Date of Birth
Gender	Marital Status	Student?
Employment Status	Age drivers license first obtained	Education level
Industry	Veteran?	Field of Occupation
Accidents/Violations/Thefts/Suspensions ?	Current Insurance Company	Years with Current Company

Table 1: Information required for an on-line auto insurance application with GEICO. Source: www.geico.com, August 2008.

for auto-insurance from the insurance company GEICO's website – www.geico.com.<sup>2</sup> GEICO uses the information received with a customer's application to determine whether or not they will offer insurance to that customer.<sup>3</sup> GEICO can use the information on these applications not only to determine which customers to accept, but the rate that they want to quote to each customer. Mortgages, home equity loans, student loans, and auto loans all use a similar process of customer approach followed by an application with significant disclosure. Loan prices can depend both on characteristics of the lending product such as term and size of loan and on characteristics of the customer – particularly the lender's estimation of the customer's creditworthiness. Thus, customers with better credit history are likely to be offered lower rates for auto loans than those with poor credit. For a fuller discussion of the application of customized pricing to consumer credit markets see Caufield (this volume) and Phillips (2010).

### 3 Formulating and Solving the Customized Pricing Problem

The level and type of pricing segmentation that a seller uses is a major determinant of the effectiveness of customized pricing. If a seller differentiates price by five order size tiers, six regions, and three channels, then the seller has  $5 \times 6 \times 3 = 90$  pricing segments. If the lender decided to increase the number of order size tiers to 10, then the number of pricing segments would increase to 180. If, in addition, he decided to differentiate pricing between existing and new customers, then the number of pricing segments would double to 360. For many sellers, particularly those with highly configurable or bundled products, the number of pricing segments

<sup>2</sup>Accessed on November, 4, 2008.

<sup>3</sup>In the words of Stiglitz and Weiss (1981), both insurance and consumer credit are *rationed*. There are customers to whom it is unprofitable to extend credit at any price due to adverse selection in the face of private information held by the customers. The same holds true of insurance markets as noted by Akerlof (1970).

(Figure 2 about here.)

Figure 2: Steps in the customized pricing process. (Figure adapted from Phillips [2005]).

can be very large. One mortgage lender in the United States manages over two million prices at any one time. The number of pricing segments is a measure of the amount of price differentiation being employed by a lender – the more pricing segments, the greater the level of differentiation.

To formalize the customized pricing problem, let  $N$  be the number of pricing segments. For each segment  $i$ , the customized pricing problem is to determine a price  $p_i$  for  $i = 1, 2, \dots, N$ . Define  $D_i$  as the number of customer requests that will be received in segment  $i$  during some future period. For the moment, we assume that  $D_i$  is independent of price,  $p_i$ . Define  $\rho_i(p_i)$  as the bid-response function for segment  $i$  – that is, for a segment  $i$ ,  $\rho_i(p_i)$  is the fraction of customers in segment that will purchase if the seller bids at price  $p_i$ . Define  $f_i(p_i)$  be the incremental profit function for segment  $i$  as a function of price. Then, the Unconstrained Customized Pricing Problem is:

$$\max_p \sum_{i=1}^N D_i \rho_i(p_i) f_i(p_i) \quad (1)$$

$$\text{s.t. } p > 0 \quad (2)$$

where  $p = (p_1, p_2, \dots, p_N)$ . In (1), the objective is to maximize the total expected profit from all customer segments. We will discuss later the situation in which a seller wishes to pursue other objectives for some or all segments.

We note that this formulation of the problem is quite general with respect to the level of segmentation. Setting  $N = 1$  corresponds to a single price for all customers. At the other extreme, a seller could define a separate segment for every possible combination of customer characteristics, channel, and product. For example, a heavy truck manufacturer might wish to set a different discount for an order of 39 trucks than for one of 40 trucks or 38 trucks. By defining segments ever more finely, the seller can approach the limit of “market-of-one” pricing. As an example, each piece of information shown in Table 1 is a potential pricing dimension for GEICO. If GEICO used all possible values of all of these dimensions in setting their prices, then it would come close to “market-of-one” pricing.

Effective customized pricing involves solving the problem in (1) for every possible bid request. As shown in 2, this requires five steps:

1. *Segmenting the Market.* This involves establishing a set of *pricing segments*, each of which,

in theory, could be charged a different price.

2. *Estimating Bid-Response Functions.* For each pricing segment, a bid-response curve needs to be estimated that predicts the probability that an order in that segment will accept a bid as a function of price.
3. *Calculating Incremental Profit.* For each bid request, the incremental profit that the seller expects to realize if he wins the bid needs to be determined as a function of price.
4. *Optimizing.* Given the bid-response function and incremental profit function for a given bid, the bidder needs to determine the price that is most likely to help him achieve his goals. Often – but not always – the goal is to maximize expected profitability.
5. *Monitoring and Updating.* As bid results are received, the seller needs to monitor the results relative to expectation. Based on the results, he may need to update the market segmentation and/or the parameters of the bid-response functions.

We discuss each of these steps in more detail in the following subsections.

### 3.1 Segmenting the Market

In theory, a seller would want to segment his market as finely as possible in order to maximize his profitability. In practice, several factors limit the amount of market segmentation that a particular market can support. The most important of these factors are:

- *Informational Limits.* A seller can only differentiate prices based on information available at the time of a bid. Even if a seller believes that, say, a customer's annual income has a strong influence on her response to his pricing, he cannot use that information in pricing if he does not know it when he is bidding.
- *Infrastructure Limitations.* It is not uncommon that the IT infrastructure supporting pricing limits the level of differentiation that a seller can support. For example, a seller's pricing execution system may only support up to three pricing dimensions with five values per dimension. In this case, the seller cannot support more than  $5^3 = 125$  pricing segments. Infrastructural limits are very real barriers to pricing differentiation for many companies. In some cases, they can be relaxed by investing in improved IT systems, however, changes to

price distribution and bid evaluation systems can be expensive and time-consuming and the benefits from better segmentation need to be weighed against the cost and time of changing systems.

- *Legal and Regulatory Limits.* Laws and regulations may limit the types and levels of price differentiation that a seller can employ. For example, in the United States, the Robinson-Patman Act of 1936 prohibits manufacturers and resellers from charging different prices to different retailers under certain circumstances (Marn, Roegner, and Zawada, 2004 pp. 257-258). Laws and regulations on pricing differentiation can vary widely from country to country. The Fair Lending Act in the United States prohibits lenders from discriminating among prospective borrowers strictly on the basis of age (Ross and Yinger, 2002). In contrast, setting loan APR's on the basis of age is both legal and commonplace in the United Kingdom.
- *Simplicity and Transparency.* A seller may refrain from extensive price differentiation because he believes that there is value in maintaining a simple pricing structure that is fully disclosed to all customers, particularly when the product is sold through intermediaries. Resellers and intermediaries often express a desire for "simple" pricing, although evidence is often lacking that they are willing to accept higher prices for a simpler structure.<sup>4</sup>
- *Arbitrage.* The ability of a seller to differentiate pricing among different customers may be limited due to the potential for arbitrage – customers offered a low price could purchase more than they need and resell the surplus to other customer segments, undercutting the seller. The threat of arbitrage is particularly great for standardized products with low transportation costs. It often limits the extent to which companies selling easily transportable products can charge different prices through different channels, to different customer segments, or to different countries. Services and highly customized products are not as subject to arbitrage and often support higher levels of customer segmentation.
- *Fairness Concerns.* A seller may be reluctant to differentiate prices along a certain dimension due to concerns about fairness. These concerns can be of two types: the seller may believe

---

<sup>4</sup>The passenger airlines present a case of an industry in which customers have often complained about the complexity of the fare structure, but there is no evidence that any group of customers systematically chooses higher priced flights offered by airlines with simpler fare structures. The perception that customers – especially intermediaries – will reject complex or non-transparent pricing remains quite high in many industries, often with little or no concrete evidence.

(Figure 3 about here.)

Figure 3: An example bid-response curve.

that differentiating price along a certain dimension is unfair. Alternatively, the seller may be concerned that differentiating along a certain dimension may be perceived by customers as unfair. This could lead to resentment on the part of the customers and, ultimately, lower demand. See Maxwell (2008), Chapter 13 in Phillips (2005), and Özer and Zheng (this volume) for some of the “perception of fairness” issues encountered in pricing.

Due to these limitations most sellers operate in a world of finite customer segmentation well below the theoretical limit of one-to-one pricing.

### 3.2 Estimating the Bid-Response Function

A key element in customized pricing is the supplier’s uncertainty about how a prospective buyer will respond to a bid price. Presumably a supplier who bids on a piece of business believes that he has a non-zero probability of winning – otherwise he would not waste the time and effort to bid.<sup>5</sup> It is also reasonable to assume that the seller’s estimate of the probability that his bid will win should be a decreasing function of price – that is, the higher the price that he bids, the lower his probability of winning. For each bid, the *bid-response function* specifies the seller’s probability of winning as a function of the price that he bids. A typical bid-response function is shown in Figure 3. In this figure, the horizontal axis is price, the vertical axis is the probability of winning the bid and the bid-response function is a decreasing function of price.

The bid-response function shown in Figure 3 is analogous to the more familiar “price-response” or “demand” curves found in many discussions of price theory such as Phillips (2005) and van Ryzin (this volume). The bid-response function represents the seller’s probability that he will win a particular bid as a function of price while a demand curve represents the total demand that a seller would expect to receive if he posted a fixed price in some market for some period of time. In a customized pricing setting, the seller will be responding one-at-a-time to different bids with the

---

<sup>5</sup>This is not strictly true. A supplier might decide to make a bid that he knows will not win simply to signal his willingness to do business with the buyer in order to be included in future bid opportunities. Or he might decide to bid to “keep the competition honest” – that is, to provide the buyer with some leverage that the buyer could use to force a competitor to lower his price.

freedom to set a different price for each bid if he so desires. Since each bid is likely to be different (different bundles of products and services on order, different customers, different competitors, different channels), the supplier could, in theory, have a different bid-response function for each bid.

In general, the bid-response curve for a particular bid must incorporate two types of uncertainty:

1. *Competitive Uncertainty.* Typically a seller will not know details of competing bids – including prices. In many cases a seller will not know who is bidding against him or how many competing bids he is facing.
2. *Uncertainty on Buyer Preferences.* Even if a seller knew with certainty both the identity of his competitors on a particular bid *and* the prices that all of the competitors were bidding, he still may not be able to predict with certainty whether or not his bid would win at a given price. The seller will usually not know the preferences or the exact process by which the buyer will choose a winning bid. In most cases, the buyer is not certain to choose the lowest-price bid.<sup>6</sup>

We can conceive of many different ways in which a seller might estimate the bid-response curve for a particular customer request. For a particularly important bid, a seller might invest considerable time and effort in preparing its bid. For example, a large airline seeking to purchase aircraft for its fleet is likely to solicit proposals from both Boeing and Airbus. If the order is sufficiently large, both manufacturers will devote substantial amounts of time and thought to all elements of their proposals – including price. Typically each manufacturer will convene an internal group of people who have knowledge of the particular deal, experience with the customer, understanding of the competition, and experience with similar bidding situations in order to derive the best possible understanding of how the potential customer is likely to respond to the price associated with the deal. When a sale involving hundreds of millions or billions of dollars is at stake, it is not unusual for companies to devote many man-months – even man years – to developing their proposals. The ultimate price offered in the proposal will be determined as the result of long discussions, competitive role playing, and complex calculations. The final price for the transaction may be only determined after many rounds of negotiation.

---

<sup>6</sup>As an exception, some government procurements are required by law to select the lowest-price bid. As a result, government RFP's tend to be exceptionally detailed in order to minimize non-price differences among bids.

While each very large customer request can be treated as unique and subjected to in-depth analysis, there are many situations in which the following three conditions hold:

1. The seller makes many relatively small quotes – in many cases thousands or tens of thousands – during the course of a year.
2. The seller retains historic “win/loss” data. That is, the seller retains full information on the details of each quote including the product requested, the customer, the channel, the price quoted, and the outcome of the bid – i.e. if the business was won or lost.
3. The seller offers a menu of standardized products. This is in contrast to fully customized services such as architecture, construction or management consulting where each job is unique and it is difficult to establish comparability.

Under these three conditions, a seller can use statistical regression to estimate bid-response functions that can be used to determine the optimal prices for all pricing segments.

When the rate of incoming customer requests is high, it is typically too expensive or difficult for the seller to devote substantial amounts of time or effort to analyzing each deal. Auto F&I (finance and insurance) executives requesting a quote for an auto-loan typically require a response within a few seconds. Given that an auto-lender may be receiving hundreds or thousands of such requests daily it is infeasible for them to convene a corporate task force to determine the rate to quote for each one. They need a more rapid and automated approach. Fortunately, a high rate of incoming quotes usually implies a large volume of historical data that can be used to estimate stable and predictive bid-response functions that can be applied to future bids.

If sufficient historic win/loss data is available, a bid-response function can be estimated based on historical win/loss data using techniques of *binary regression*. The *target* data for binary regression consist of ones and zeroes. In estimating the coefficients of a bid-response function, the target data is the history of “wins” and “losses” that the seller has experienced in the past. In this case, a one can be used to indicate a win and a zero a loss. The *covariates* of the model include all of the other information available with the past bids including the price, the characteristics of the product or products ordered, the characteristics of the customer, and which channel the request was received through. Table 3.2 illustrates the bid history data available for a manufacturer of printer ink cartridges. This manufacturer sells three grades of cartridge: silver, gold, and platinum, to four

			Request Characteristics		Customer Characteristics		
No.	W/L	Price	Grade	Size	New?	Type	BusLevel
1	W	\$7.95	S	1,259	N	Res	0
2	L	\$9.50	P	655	E	Gov	\$153,467
3	L	\$9.75	P	790	E	Gov	\$1,402,888
4	W	\$8.50	G	840	N	Ret	0
5	L	\$8.47	G	833	E	Ret	\$452,988
6	L	\$9.22	P	540	E	Edu	\$55,422
7	W	\$8.61	G	850	N	Ret	0
...	...	...	...	...	...	...	...

Table 2: Sample historic bid data available to a printer ink cartridge seller. “No.” is an arbitrary bid number. W indicates that the bid was won, L that it was lost. “Price” is the unit price bid. “Grade” indicates the grade of the product ordered: S = Silver, G = Gold, P = Platinum. “Size” is the number of units in the order. N indicates a new customer, E an existing customer. The seller services four types of customer: Res. = Reseller, Gov = Government, Ret = Retailer, Edu = Educational Institution. “Business Level” is the total amount of business sold to this customer over all business lines in the previous 12 months.

types of customers: resellers, government agencies, educational institutions, and retailers. The manufacturer has kept track of the outcome of each bid and the amount of business that it has done with each customer over the previous year. This information has been stored in a database and can be retrieved in a format similar to that shown in Table 3.2.

The challenge facing the cartridge manufacturer is how to use the data in Table 3.2 to estimate bid-response functions. This typically requires four steps: 1) segmenting the market, 2) determining the model structure, 3) estimating the coefficients of the model based on historic data, and 4) measuring the quality of the model. This is a classical problem of model specification and estimation – not different in principle from similar problems faced in promotion response estimation as discussed in Blattberg and Briesch (this volume). We will give a broad introduction to some of the issues involved in developing bid-response models, more detailed discussions of statistical modeling can be found in any standard text on statistical modeling.

Model specification, and model estimation typically proceed iteratively – that is, market segments and model structure will often be sequentially “tweaked” until the model fit meets some criteria or until additional tweaking fails to make improvements. To illustrate this process, consider the example shown in Table 3.2 and define the following notation:

- GRADES = 1 if Grade is Silver, 0 otherwise; GRADEG = 1 if Grade is Gold, 0 otherwise;

GRADEP = 1 if Grade is Platinum, 0 otherwise;

- $PRICE$  = price;
- $SIZE$  = order size;
- $NEW$  = 1 if customer is new,  $NEW$  = 0 if customer is existing;
- $RES$  = 1 if the customer is a reseller, 0 otherwise;  $GOV$  = 1 if the customer is a government agency, 0 otherwise;  $RET$  = 1, 0 otherwise; if the customer is a retailer, and  $EDU$  = 1 if the customer is an educational institution, 0 otherwise<sup>7</sup>;
- $BUSLEVEL$  = total amount of business done in the last 12 months.

A bid-response model using this data would specify the probability of winning a bid as a function of the price  $p$  and all of the other information available with the order. For the printer cartridge manufacturer, this information is grade, size, customer status, customer type, and level of business sold to the customer. Denote all of these non-price characteristics as a vector  $x$ . One extremely popular function for bid-response modeling (and, in fact for binary response models in general) is the *logit* function, which is given by:

$$\rho(p, x) = 1/(1 + e^{g(p,x)}) \quad (3)$$

where  $g(p, x)$  is an affine function of price and the (possibly transformed) non-price characteristics of a bid. The logit function is popular both because it is easily tractable and because statistical packages such as SAS and R include extensive support for it. The logit is a member of a larger category of statistical models known as *Generalized Linear Models* (GLMs). Other commonly used GLM forms include the probit and the linear.

---

<sup>7</sup>The variables indicating which grade and type the order falls into are called *categorical variables*. Since, by assumption, the order can only be for one grade of cartridge, than exactly one of GRADES, GRADEG, and GRADEP can be 1 with the other two 0. A similar property holds for the variables indicating type. Experienced modelers will recognize that, when there are  $n \geq 2$  categorical variables, that only  $n - 1$  need to be included in the model since the value of the missing variable can be inferred from the values of the others. For example, in the example, if GRADES and GRADEG are both equal to zero, than GRADEP must equal 1. If either GRADES or GRADEG are equal to one, then GRADEP must equal 0. This means that  $GRADEP = 1 - GRADES - GRADEG$ . In other words, GRADEP is *co-linear* with GRADES and GRADEG and does not need to be included as an explanatory variable. In what follows, we will ignore this and continue to include all of the variables in the model formulation.

Once the choice has been made to use the logit, the next step is to estimate the coefficients of the variables in the function  $g(p, x)$ . One obvious choice is simply to include all of the available variables in their raw form. In the case of the printer cartridge manufacturer, this would result in a specification of the form:

$$\begin{aligned}
 g(p, x) = & \beta_0 \\
 & + \beta_1 \times PRICE \\
 & + \beta_2 \times GRADES + \beta_3 \times GRADEG + \beta_4 \times GRADEP \\
 & + \beta_5 \times SIZE \\
 & + \beta_6 \times NEW \\
 & + \beta_7 \times RES + \beta_8 \times GOV + \beta_9 \times RET + \beta_{10} \times EDU \\
 & + \beta_{11} \times BUSLEVEL.
 \end{aligned} \tag{4}$$

Equations 3 and 4 specify a *statistical model* for bid response. Once a model has been specified, the next task is to determine the values of the coefficients – in this case, the values of  $\beta_0$  through  $\beta_{11}$  – that best fit the historic data. This is a standard problem of binary regression which we will not address in detail here except to note that most common statistical packages such as SAS and R include routines for estimating the coefficients of a logit model. They will also calculate various measures of statistical fit such as Concordance and the Akaike Information Criterion (AIC) that estimate how well the chosen model actually fits the data. How binary regression can be used to estimate the coefficients for a logit bid-response model is discussed in more detail in Phillips (2005), pages 284-287. Some good additional references on the properties and estimation of GLMs are McCullagh and Nelder (1989), Lindsey (2000) and Dobson and Barnett (2008).

Equation (4) is not the only possible formulation of a bid-response model using the data from Table 3.2. For example, (4) includes the term  $\beta_5 \times SIZE$ . This term specifies the effect of order size on the probability that the seller will win a bid, all else being equal. In many situations, order size has a strong influence on the probability that a bid will be won at a particular price – typically, customers placing larger orders tend to be more price-sensitive. The model specified in Equation (4) represents this dependence as linear in the size of the order. In many cases – particularly when order size can span a very large range – the strength of bid-response may be more closely correlated with the *logarithm* of order size. This would suggest a model in which the term  $\beta_5 \times SIZE$  is replaced with  $\beta_5 \times \log(SIZE)$ . The supplier might also want to consider a

model in which the product of price and order size influences the probability of winning a bid, in which case he could add an additional term of the form  $\beta_{12} \times PRICE \times SIZE$ . Determination of the best statistical model for bid-response is partly art and partly science. The process usually proceeds by sequentially trying different models and keeping the one that best fits the data. By sequentially comparing alternative models and choosing winners, an experienced modeler can usually develop a predictive and stable model.

We note that the statistical procedure described here is only one possible approach to estimating the bid-response curve and, indeed, is only feasible when the seller has retained historic win/lose information *and* there is sufficient historic data to support the estimation. Our experience has been that when these two conditions are satisfied, standard binary regression approaches such as maximum likelihood estimation can deliver stable, significant and highly predictive estimates of the bid-response curve. However, there are a number of factors that can confound the estimation, particularly the presence of *endogeneity*. Endogeneity occurs when the price offered to a customer is influenced by variables correlated with his price sensitivity that are not included in the data. For example, a car salesperson may use how well a customer is dressed as an indication of the willingness of the customer to accept a higher price. To the extent the salesperson is correct, the price offered to customers is not independent of their willingness-to-pay and, as a result binary regression will tend to underestimate price sensitivity. The effect of endogeneity can be significant<sup>8</sup>. To the extent that the magnitude of endogeneity is understood, the regression can be adjusted to account for it. Alternatively, random price tests can be used to generate observations that are free from any potential taint of endogeneity.

There are also cases in which sufficient historic data may not be available (as in the introduction of a new product) or in which the seller has not preserved a record of wins and losses. In these cases, alternative approaches such as judgmental estimation of the bid-response curves must be used. In any case, the initial estimation of the bid-response curve should always be updated over time as customer response to new bids is observed.

---

<sup>8</sup>In a meta-analysis of price-estimation studies, Bijmolt et. al. (2005) found that the treatment of endogeneity had a major effect on the estimates of price elasticity across industries.

### 3.3 Calculating Incremental Profit

The objective function in the Unconstrained Customized Pricing Problem in (1) specifies that the seller is seeking to maximize total expected profit, which is calculated as the product of the probability of winning a bid and the incremental profitability if the bid is won. *Incremental profitability* is calculated as the total expected profitability of the seller if the bid is won minus the expected profitability if the bid is lost. In many cases, the incremental profitability of a transaction is simply the price charged minus the unit cost, that is,  $f_i(d, p) = d(p - c_i)$  where  $p$  is the price charged,  $c_i$  is the unit (incremental) cost per sale in segment  $i$ , and  $d$  is the order size. This definition of incremental profitability assumes that unit cost is fixed and independent of both the selling price and the number of units sold. In this case, the unit cost  $c_i$  should be calculated so that the total cost of the order  $dc_i$  is equal to the difference between the total cost that the company will incur if it makes the sale minus the total cost that the company will incur if it doesn't make the sale.

There are cases where the simple linear relationship  $f_i(d, p) = d(p - c_i)$  does not apply. In some cases, a seller is bidding to provide an unknown level of products or services to a buyer for future business. For example, UPS competes with FedEx in the package express business. Typically a potential customer will request bids from both companies. A typical bid by UPS or FedEx would be to serve all of the package express business generated by that customer for the coming year. At the time of the bid, neither UPS nor FedEx nor the customer can perfectly forecast its shipping needs for the next year. In this and similar cases, incremental profitability is a random variable and optimization is over expected incremental profitability.

The level and composition of products or services that will be demanded from a supplier under a contract may not only be uncertain, they may depend on the price. For example, many manufacturers contract with less-than-truckload (LTL) trucking companies such as Roadway or Yellow Freight to transport their products to distributors or retailers. A common practice is to choose two or three LTL companies as *preferred suppliers* from ten or more who bid on an annual contract. Under this arrangement, the shipper commits to use one of the preferred suppliers for every shipment. The shipper may also guarantee a minimum level of business to each preferred supplier. However, when it comes time to move a particular shipment, the shipper is more likely to choose the supplier who has bid the lowest price for that type of shipment. Even if a trucking company wins the right to be a preferred supplier, the amount of business that it will receive will

(Figure 4 about here.)

Figure 4: Calculating expected profitability as a function of price. Expected profitability is the product of bid-response and incremental profit. The profit-maximizing price is shown as  $p^*$ .

be a decreasing function of price. More detail on the LTL trucking industry can be found in Kintz (this volume).

Finally, as noted by Phillips (2010), in consumer credit markets, incremental cost is typically an *increasing* function of price due to *adverse selection*. This means that, as the prices offered by a lender increase, the loss rates for the loans that it funds will also increase. This will occur even if the lender does not change its underwriting guidelines – that is, it does not change the criteria by which it chooses which applicants it will lend to. As the prices offered by a lender rise relative to the competition, customers with a lower probability of default will defect to other lenders at a higher rate than those with a higher probability of default – who have fewer alternatives. A similar phenomenon occurs in insurance markets. In these cases, the unit cost of serving segment  $i$ ,  $c_i$  cannot be treated as a constant, but must be represented as a function of price – that is, as  $c_i(p)$ . More detail on the effect of adverse selection on pricing in consumer lending can be found in Phillips and Raffard (2010).

### 3.4 Optimization

Once the bid-response function and the incremental profit function have been determined for each pricing segment, the next step is to determine the optimal price for each segment. If the seller is seeking to maximize expected profitability and does not wish to apply any constraints, then this is equivalent to finding the set of prices that solves the optimization problem in (1). Since this problem is separable, the optimal price for each segment can be determined independently. The optimization problem for a single segment is illustrated graphically in Figure 4. In this figure, the dashed downward-sloping curve is the bid-response function and the upward sloping solid line is the incremental profit function. The product of bid-response and incremental profit is expected profit which is the hill-shaped curve in Figure 4. In the absence of any constraints, a seller who seeks to maximize expected profit would choose the price at the “top of the hill”, labeled  $p^*$  in Figure 4.

For most realistic bid-response curves and incremental profit functions, expected profitability is

well-behaved in the sense that it is a smooth function of price with a single peak.<sup>9</sup> This is the case in Figure 4. This means that, in the absence of constraints, that the optimal price can be calculated using standard “hill-climbing” approaches such as gradient ascent.

### 3.4.1 Optimality Condition

While the unconstrained customized pricing problem in (1) can be easily solved using numerical techniques, it is useful to observe that the optimal price for a segment obeys a standard price-optimality condition. For each segment, define the *bid price elasticity* as the percentage change in the probability of winning a bid divided by the percentage change of price:

$$\epsilon_i(p) = |\rho'_i(p_i)p_i/\rho_i(p_i)|, \quad (5)$$

$\epsilon_i(p_i)$  in Equation 5 is the analog of the familiar concept of *own-price elasticity* which is defined as the percentage reduction in demand resulting from a 1% increase in price. It should be noted that, in the unconstrained case, Equation 5 implies that the optimal price for segment  $i$  must satisfy:

$$e_i(p_i^*) = \frac{f'_i(p_i^*)p_i^*}{f_i(p_i^*)}. \quad (6)$$

Note that if  $f_i(p_i) = p_i - c_i$  – that is, incremental profit is equal to price minus unit cost – then condition 6 reduces to

$$e_i(p_i^*) = p_i^*/(p_i^* - c_i).$$

This is the “elasticity equals the reciprocal of unit margin” condition for price optimality (See Phillips [2005], pp 64-65 and van Ryzin, this volume). Equation (6) is the extension of this well-known condition to the case in which incremental profit is a more complex function of price.

### 3.4.2 Constrained Problems

The optimization problem in (1) is unconstrained. In most business applications, the seller will wish to set constraints on prices. Examples of typical constraints include:

---

<sup>9</sup>Most standard bid-response functions such as probit, logit, or linear satisfy a property known as the Increasing Failure Rate (IFR) property (Barlow and Proschan, 1965). If the bid-response functions  $\rho_i(p)$  demonstrate the *Increasing Failure Rate* (IFR) property *and* all the incremental profit functions  $f_i(p)$  are increasing, continuous, and concave, it can be shown that a unique optimal price will exist for each segment. This uniqueness property was apparently first identified by the econometrician Theil (1948) in his Master’s Thesis.

- *Bounds*: Typically, user-specified upper and lower bounds,  $p_i^+$  and  $p_i^-$  are applied on each price by specifying constraints of the form  $p_i^+ \geq p_i^*$  and  $p_i^* \geq p_i^-$  for each segment  $i$ . Price bounds can be applied for a number of different reasons. They may be used in order to maintain some level of price stability by making sure that the new price does not deviate too much from a previous one. For example, an auto manufacturer may want to make sure that the per-unit price quoted to its fleet customers for an order is never more than 10% higher than the last price quoted previously to the same customer. Bounds are also applied to ensure that recommended prices are within the region of statistical reliability of the bid-response curve calculation. In other cases, regulations may require a cap on the maximum rate that can be quoted to a particular pricing segment. For example, usury laws in some states specify a maximum interest rate that can be charged for consumer loans.
- *Monotonicity Constraints*: In many cases, sellers want to ensure that prices consistently increase or decrease along certain dimensions. A seller might want to ensure that, for otherwise identical customer requests, the unit price bid for a larger order should never be higher than the unit price bid for a smaller order – otherwise customers could get a lower price by breaking a large order into several smaller orders. As another example, sellers often require that an order from an existing customer should never be priced higher than an identical order from a new customer.
- *Business Performance Constraints*: A seller may wish to maintain certain minimum levels of total sales or revenue, even at the expense of profitability. Management might give a directive such as; “we want to maximize contribution, but we can’t allow sales to fall below \$10 million for our flagship product during the next quarter or analysts are likely to downgrade our stock.” This can be imposed by adding a constraint of the form:

$$\sum_{i \in I} D_i \rho_i(p_i) p_i \geq \$10,000,000$$

where  $I$  is the set of segments that include the flagship product.

- *Price banding*. A seller may want prices in a particular region or through a particular channel to maintain some relationship to prices in other regions or channels. For international companies, such price bands are often necessary to prevent arbitrage. A semiconductor manufacturer may need to specify that chips cannot be sold in Brazil for less than two cents per unit less than they are sold in the United States, otherwise it would be profitable for

arbitrageurs to purchase chips at the lower price in Brazil and resell them in the United States.

- *Channel constraints.* Sellers often wish to maintain relationships among the prices charged through different channels. For example, a seller might want to ensure that the price quoted through the Internet for a customer request should never be higher than the price quoted for the same request received through a call center.

Each of the conditions described above can be imposed by adding one or more constraints to the optimization problem specified in (1). From a technical point-of-view, adding constraints usually makes the pricing optimization problem harder to solve. However, as long as the constraints define a convex feasible region, standard solution approaches can be used to solve for optimal prices. From a business point-of-view, adding and managing constraints can be more problematic. In particular, care needs to be taken that users do not over-constrain the problem.

### 3.5 Monitoring and Updating

As shown in Figure 2, an effective customized pricing process needs to include a mechanism for monitoring the market and updating prices over time. No matter how carefully crafted, a set of prices cannot be optimal forever. Prices need to be adjusted in response to changes in the macroeconomic environment, changes in costs, shifting customer preferences, and competitive actions. Depending upon the market and velocity of transactions, prices might need to be updated daily, weekly or monthly. Most companies using an analytical approach to customized pricing update prices on a fixed periodic basis with interim *ad hoc* changes in response to external events.

Not only prices, but model coefficients also need to be updated. That is, the values of  $\beta_0, \beta_1, \dots, \beta_{11}$  estimated by the printer cartridge manufacturer for the model in (4) will need to be monitored periodically and refined. Typically, model coefficients are updated much less frequently than the prices themselves. Prices need to be changed whenever market conditions change, costs change, or the business goals and constraints change. Depending upon the market, coefficients might be updated monthly, bi-monthly or even semi-annually. A seller should periodically monitor the performance of the statistical model relative to actual results. If the predictions of his model begin to deviate significantly from reality, then it is a good idea to re-estimate the coefficients.

One way to update the coefficients is to append the most recent win/loss observations to the data

file and re-run the regression. This is typically done with weighting the most recent observations more heavily than the historic data. Alternatively, various *Bayesian approaches* can be used. With Bayesian updating, new observations are used directly to update the values of the parameters<sup>10</sup> Updating is particularly important early in the adoption of an analytic approach in order to ensure that the initial set of coefficients has been estimated accurately. It is also important in markets with rapidly changing costs.

## 4 Enhancements and Extensions

The previous sections have described the application of an analytical approach to customized price optimization in the “plain vanilla” case in which the seller seeks to maximize expected profitability, each customer request contains only one product, the prices offered by competitors are not available, and the bid price is not negotiated. In most real-world applications, one or more of these conditions will not hold. We now briefly discuss the effects of relaxing these conditions.

### 4.1 Alternative Objective Functions

The customized pricing problem as specified in Equation (1) maximizes expected contribution. This is consistent with a corporate goal of maximizing expected short-run profitability. However, it is often the case that a seller might wish to maximize expected *revenue* rather than expected profitability for one or more segments. For certain segments, a seller may wish to hold or increase its market share for strategic reasons and is willing to give up some profitability to do so. In that case, the objective function in (1) would be replaced with:

$$\max_p \left[ \sum_{i \in I_1} D_i \rho_i(p_i) f_i(p_i) + \sum_{i \in I_2} D_i \rho_i(p_i) p_i \right].$$

where  $I_1$  is the set of pricing segments for which expected contribution is to be maximized and  $I_2$  is the set for which revenue is to be maximized.

Maximizing revenue always results in lower prices than maximizing profitability. Intuitively, this occurs because the “profit maximizing price” is an increasing function of the unit cost and the revenue-maximizing price is the same as the profit-maximizing price with a unit cost of 0. When unit cost is greater than 0, the additional revenue generated by the revenue-maximizing price

---

<sup>10</sup>Gill (2008) provides an introduction to the use of Bayesian statistics.

(Figure 5 about here.)

Figure 5: An efficient frontier. The frontier represents all points at which profit is maximized subject to a minimum revenue requirement or, equivalently, revenue is maximized subject to a minimum profit constraint. The firm is currently operating within the frontier at point A. Point B maximizes profit at current revenue and Point C maximizes revenue at current profit. Any point on the frontier between B and C achieves both higher profit and higher revenue than point A. Point D maximizes profit, but at lower revenue than Point A.

comes at the expense of profitability – this is often described as “buying market share”. One way to visualize the tradeoff between contribution and revenue (or market share) is through the use of an *efficient frontier* as illustrated in Figure 5. The efficient frontier shows all of the combinations of prices at which profit is maximized subject to achieving at least a certain level of revenue. Points inside the frontier can be achieved by changing prices; points outside the frontier cannot be achieved by changing prices. It could be argued that a firm would always wish to be operating at a point on the efficient frontier. If it is operating inside the efficient frontier, it could increase both revenue and profitability simply by changing prices. Thus, a firm currently operating at Point A in Figure 5 could achieve the same level of profitability but higher revenue by moving to Point B. Alternatively, it could maximize profitability at the same level of revenue by moving to Point C. Points between B and C on the efficient frontier all represent points at which the firm could increase both revenue and profitability relative to operating at Point A.

The efficient frontier enables the seller to calculate how much profit it would lose by meeting any particular revenue target. This can be a valuable insight since it allows management to set revenues with a full understanding of the implications for reduced profit.

## 4.2 Incorporating Competitive Information

So far, we have ignored competition. At first glance, this might appear to be a major lapse. After all, most companies would likely nominate “what the competitor is charging” as a major determinant of whether or not they will win a bid. In fact, there is no theoretical difficulty in incorporating competitive prices into the bid-response function. If competitive prices are known at the time bids are made, they can (and should) be incorporated into the bid-response function. To see how this is done, assume that the printer cartridge seller has two competitors. Let  $p_A$  denote the price bid by the first competitor and  $p_B$  the price bid by the second competitor. Then, the

model specified in 4 could be supplemented with the two additional terms  $\beta_{12} \times p_A$  and  $\beta_{13} \times p_B$ . The values of  $\beta_{12}$  and  $\beta_{13}$  can be determined using regression in the same fashion as the other coefficients and used as predictors for future bid-response.

The rub is that in the majority of customized pricing settings, competing prices are typically not known when the price for a bid must be determined. A seller responding to an RFP or an on-line order inquiry will typically not know the identity or even the number of competitors he is facing – much less what prices they are bidding. In fact, in most cases, the seller will not have full information about competitive pricing even if he wins a bid. Thus, there will be little or no competitive price information in the historical win/loss data. In this case, the best option open to the seller may be to use a bid-response model that does not explicitly include competitive pricing terms.

Note that excluding competitive pricing terms from the bid-response curve is not the same as “ignoring competition”. Rather, it is equivalent to assuming that *competitors will set prices in the future in the same way as they have done in the past*. Even though they do not contain explicit competitive price terms, the bid-response curves implicitly incorporate within them the effect of past competitive pricing. The results of running a regression may be to show that a particular segment seems to be highly price-sensitive. This could be because customers in that segment are intrinsically more sensitive to price. However, it could also indicate that one or more competitors typically price “aggressively” (i.e. low) in this segment. As long as competitive pricing in this segment does not change, the predicted level of price sensitivity should be stable and predictive of the future. If a competitor changes its pricing strategy for a segment – say by raising price, a supplier will find that it is winning more deals in that segment than it forecasted. This should be a trigger to re-estimate the model coefficients. The updated coefficients will reflect the new competitive strategy.

### 4.3 Segment Selection

The formulation of the Unconstrained Customized Pricing Problem in (1) implicitly assumes that some price will be quoted to all customer segments. There are situations in which this is not the case, that is, the seller does not want to sell to every segment. The most obvious case is lending and insurance – as noted before, due to risk, there are some customers who are unprofitable at any price and therefore cannot obtain credit at all. However, there are other cases in which a seller

might not wish to sell to every potential customer segment. It may be unprofitable to sell certain items through certain channels or to certain regions. It is not uncommon to set minimum order size restrictions on small, inexpensive items and so on. This suggests that the formulation of the Unconstrained Customized Pricing Problem in (1) should be expanded to enable simultaneous pricing and segment selection:

$$\begin{aligned} \max_{p,x} \quad & \sum_{i=1}^N D_i x_i \rho_i(p_i) f_i(p_i) \\ \text{s.t.} \quad & p > 0 \\ & x \in \{0, 1\} \end{aligned} \tag{7}$$

where  $x = (x_1, x_2, \dots, x_n)$  is a vector such that  $x_i = 1$  means that the seller should sell into customer segment  $i$  and  $x_i = 0$  means that he should not sell into that segment. The unconstrained version of the problem shown in (7) is not difficult to solve: simply solve the Unconstrained Customized Pricing Problem in (1) and set  $x_i = 1$  for all segments whose maximum profit is greater than 0 and set  $x_i = 0$  for all segments with maximum profit less than 0. A similar approach can be used when the only constraints are price bounds. However, if there are many constraints applied, the problem can become quite difficult because the constraints may force some segments to be served at unprofitable prices in order to meet the constraints.

#### 4.4 Multi-dimensional Prices and Bundled Products

In many cases, there is more than one “price” required for a bid. In the United States, lenders typically charge a number of up-front fees to obtain a mortgage. In addition, the mortgage may have “points” associated with it, where the points are a fee paid by the borrower expressed as a percentage of a loan. Points can enable a borrower to obtain a lower APR by paying a fee expressed as a percentage of the amount they wish to borrow. For example, a \$100,000 mortgage might be available at 7.9% with \$1,000 in fees and .5 points. The .5 points would be an additional fee of  $.5\% \times \$100,000 = \$500$ . Since fees and points are usually rolled into the initial balance, the borrower is actually borrowing \$101,500 at a 7.9% APR and her monthly payment is computed accordingly. All three components – APR, fixed fees, and points – are part of the “price” of the mortgage. All three can influence the buyer’s decision whether or not to accept the mortgage as well as the seller’s profitability if the mortgage is chosen.

As another example, many business-to-business service providers such as telecommunications

companies price based on an “ $n$ -part” tariff in which the total cost to the buyer is based on a periodic base price plus a usage cost that depends on the level of service used – the higher the level of service, the lower the per-unit usage cost. Typically the per-unit cost is a step-function of usage, thus the charge might be a base cost of \$100 per month plus \$.02 per minute for the first 300 minutes per month, \$.015 per minute for the next 100 minutes, and \$.01 per minute for any usage above 400 minutes. A bid must specify the base cost as well as all the per-unit costs as well as the breakpoints. In theory, a supplier could specify a personalized  $n$ -part tariff for each bid. In practice, most sellers tend to maintain a few standard tariff structures. When they bid, they specify the tariff structure and a discount to be applied. The problem facing the bidder is to determine which structure and what level of discount level to bid for each request.

Another example of multi-dimensional pricing is provided by business-to-business software licenses that often specify timed payments. For example, a software license might specify a payment to be made upon signing, a further payment to be made upon installation, and additional support and maintenance fees to be paid annually for five years. Each of these payment amounts is a component of the overall “price”.

There are several approaches to optimizing multi-dimensional prices. One is to compute an “aggregate price” which ideally reflects the metric that buyers are using to choose among bids. For example, it is not unreasonable to assume that borrowers use monthly payment as the “pricing metric” that they use to compare alternative mortgages. A lender might then consider monthly payment to be a logical choice for the “aggregate price” of a mortgage. A software company might consider the Net Present Value of payments to be the aggregate price of its software. Let  $p = (p_1, p_2, \dots, p_m)$  be the elements of price and let  $q_i(p)$  be the aggregate price for segment  $i$ . Then, the customized pricing problem with aggregate prices can be written:

$$\max_p \sum_{i=1}^N D_i \rho_i(q_i(p)) S_i(p)$$

where  $S_i(p) : \mathbb{R}^m \rightarrow \mathbb{R}^1$  is a function that specifies incremental profitability for segment  $i$  as a function of the various price elements. The aggregate price can be used as an explanatory variable within binary regression in order to estimate the bid-response functions  $\rho_i(q_i(p_i))$ .

While the aggregate price approach has the advantage of being simple, it assumes that buyers in all segments are indifferent among options with the same aggregate price. However, in most cases,

it is likely that different buyers weight price elements differently and that no single aggregate price measure holds for all buyers. For companies purchasing software licenses, some may face short-run budget constraints that make them very sensitive to the up-front cash outlay while other buyers may be quite willing to accept a higher up-front cash payment in return for a lower overall cost of ownership. This suggests the alternative approach of including each pricing dimension independently in the regression and allowing the regression to determine their weights. For the software license example,  $p^1$  might be the up-front cash payment,  $p^2$  the total license fee, and  $p^3$  the annual support and maintenance. By including  $p^1$ ,  $p^2$ , and  $p^3$  as covariates in the binary regression, the seller could, in theory, determine how different segments weight each of the pricing elements in their pricing decisions and determine the corresponding bid-response curves as a function of all three elements.

As usual in regression, the best approach will depend upon the situation. The aggregate price approach is more parsimonious because it collapses all of the pricing elements into a single aggregate price. However, as discussed, the aggregate price chosen may not accurately reflect how buyers actually compare alternatives. Several different aggregate prices may need to be considered. While the approach of incorporating all of the pricing elements in the regression is more flexible, it also has drawbacks. It requires much more data for stable estimation since a coefficient needs to be estimated for each pricing element. Choosing among approaches may require several rounds of trial-and-error and comparison of the results of different approaches.

## 4.5 Bundled Products

A problem similar to multi-dimensional pricing is faced by sellers who are selling bundled products or highly configured products. In many cases, a price (possibly quoted as a discount from list price) must be quoted individually for every element of the order. For example, automotive fleet RFP's often include several different types of vehicles – for example, 10 cars, 5 pick-up trucks, and 5 panel vans. If the bid is indivisible – that is, the buyer firmly commits to purchase all of the vehicles from the same seller – then the best strategy is to determine the optimal price to quote for the entire bid. If, on the other hand, the bid is divisible in the sense that the buyer may choose to purchase the cars from one supplier, the trucks from another, and the panel vans from a third; then the optimal strategy is to treat the bid as three independent bids – that is, determine the optimal bid for the cars, the optimal bid for the pick-up trucks, and the optimal bid

(Figure 6 about here.)

Figure 6: Setting bounds for guide price negotiation.  $p_L$  and  $p_U$  are the prices at which expected profit is \$13,500: 10% lower than the optimal expected profit of \$15,000 which is achieved at  $p^* = \$150$ . Any price between  $p_L$  and  $p_U$  will result in expected profit within 10% of the optimal.

for the panel vans. While these two extreme cases are straightforward, many bidding situations fall between these extremes. In particular, a buyer may express a preference for purchasing from a single supplier but reserve the right to purchase from two or more. In this case, the allocation of the total price among the vehicle categories becomes important – the price allocated to each vehicle category may need to be chosen to be “competitive” – that is not too far out-of-line with expected competitive bids. More discussion of bundled pricing can be found in Chapter 11 of Phillips (2005) as well as in Oren (this volume) and Gallego and Stefanescu (this volume).

#### 4.6 Negotiated Deals

The discussion so far has assumed a “take it or leave it” pricing situation – that is, the buyer describes her needs, sellers quote their prices, and the buyer then chooses which seller (if any) from which to purchase. The price quoted by each seller is final – the only decision facing the buyer is which bid (if any) to accept. While this reasonably characterizes many customized pricing situations, there are cases in which the final price is the product of one or more rounds of negotiation between the buyer and the seller. The prevalence of negotiation in customized pricing differs widely from industry to industry and even from segment-to-segment within an industry. For example, in the US auto-lending market, negotiation is relatively rare in prime markets but almost universal in sub-prime markets. Negotiation is generally more common for larger deals. It is also more common in face-to-face selling situations (whether direct or indirect) than telesales or Internet channels. For those situations in which negotiation is likely, an effective approach should provide more than simply a single optimal price – it should also provide some guidance regarding an acceptable range of prices and, ideally, some idea of the tradeoffs between price and other aspects of the deal.

A common approach to support negotiated pricing is to use the expected profit function to help define a range of acceptable prices. The basic idea is illustrated in Figure 6. Here, the unit price that maximizes expected profit is \$150, at which price the expected profitability from the

bid under consideration is \$15,000. A lower price and an upper price have been set so that the expected profitability is \$13,500 at both  $p_L$  and  $p_U$ . At any price between  $p_L$  and  $p_U$ , the expected profit will be within 10% of the optimum. This range along with the target price can be used by the seller to guide sequential rounds of negotiation and ensure that the final outcome is within a desired range.

The approach illustrated in Figure 6 is often used to support a combined centralized/decentralized approach to negotiated pricing. The target price  $p^*$  and the upper and lower bounds  $p_L$  and  $p_U$  can be calculated centrally consistent with corporate goals and business constraints. The local sales person has the freedom to negotiate the best price he can within the specified bounds. This approach can deliver the best of both worlds: all of the data available to the corporation is used to calculate the bid-response curves and set pricing ranges for each segment based on overall corporate strategy while local knowledge of individual customers can be used to negotiate the best deal within the specified range. In many cases, the local sales person will have specific customer or competitive knowledge of the deal under consideration that can lead to better overall performance than using the same optimal price for all deals within a segment.

Of course, in many negotiated settings, price is not the only element of the deal in play. An auto lender might be willing to offer a lower APR if the borrower is willing to put more money into the down payment. A heavy truck manufacturer might be willing to lower the price if the buyer agrees to purchase an extended warranty. An enterprise software vendor selling a multi-year license deal might be willing to lower the price if the buyer agrees to pay more cash up front. Complex negotiations can involve multiple rounds of give and take and a full discussion of the “art and science of negotiation” is well beyond the scope of this section – Raiffa, et. al. (2003) provides an introduction to the topic. However, the analytic concepts behind price optimization can also provide invaluable guidance to negotiation. In particular, they give insight into the central questions: If we change this non-price characteristic of the deal, what will be the effect on incremental profitability? How, then, should we change price in order to maximize expected profitability given the new profitability level? A number of customized pricing optimization software systems enable users to perform these computations in order to help support the negotiation process.

## 5 The Future of Customized Pricing

The primary focus of this article has been on the technical and mathematical aspects of pricing optimization. The problem is of sufficient complexity that, for sellers selling multiple products to many different customer segments, automated software systems are often used to set customized prices and update them over time. Whether or not the size and complexity of the customized pricing problem facing a seller is sufficiently great to justify implementing a software system, many companies have the opportunity to improve their pricing by implementing a consistent process of learning and improvement such as the one illustrated in Figure 2.

Of course, the primary motivation for adopting an analytical approach to customized pricing is usually to increase profitability. The package shipper UPS reported a profitability increase of more than \$100 million per year in North America through use of a software system to optimize customized prices (Boyd, et. al. 2005). The sub-prime auto lender AmeriCredit saw a \$4 million increase in profitability within three months of implementing a customized pricing optimization system (Phillips, 2010.). Results such as these provide strong motivation for companies utilizing customized pricing to invest in improving both the processes and the systems that they use to set prices. There is at least anecdotal evidence that an increasing number of firms using customized pricing are applying analytic approaches to setting their prices. The *Wall Street Journal* described the benefits resulting from introduction of more analytical approaches to customized pricing at industrial parts manufacturer, Parker Hannifin (Aepfel, 2007). A number of companies such as Nomis Solutions and Zilliant offer automated pricing optimization solutions to companies facing customized pricing situations.

Customized pricing also offers a number of promising research opportunities. Relative to its importance in the economy, it has been relatively less studied than other pricing mechanisms such as auctions or list pricing. There at least two promising areas for future research. One is the estimation of customer price-sensitivity information from historical data, especially the elimination of endogeneity effects as described in Section 3.2. The identification and use of instrumental variables to address endogeneity as in Berry et. al. (1995) holds promise but its use has not been documented in a customized pricing setting. The other area for additional research is in the area of optimal segmentation – when the ability to segment is limited, how should a seller choose and price to segments in order to maximize expected profitability? Both of these

represent areas where research breakthroughs could lead to substantial improvements in practice.

## 6 Acknowledgments

I would like to thank Joe Nipko, Özalp Özer, Robin Raffard, and an anonymous reviewer for their thoughtful comments that led to improvements in this chapter.

## 7 References

Aeppel, T. (2007). “Seeking Perfect Prices, CEO Tears up the Rules,” *The Wall Street Journal*, March 27, 2007;

Agrawal, V. and Ferguson, M. (2007). “Bid-Response Models for Customised Pricing”, *Journal of Revenue and Pricing Management*, **6** 3, 212-228;

Akerlof, G. A. (1970). “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism.” *Quarterly Journal of Economics*, **84** 3, pp. 488-500;

Barlow, R. E., and Proschan, F. (1965). *Mathematical Theory of Reliability*. New York, NY: John Wiley and Sons, Inc.;

Berry, S., Levinsohn, J., and Pakes, A. (1995). “Automobile Prices in Market Equilibrium,” *Econometrica* **63** (4) 841-890;

Bijmolt, H., van Heerde, H. J., and Pieters, R. G. M. (2005). “New Empirical Generalizations of the Determinants of Price Elasticity.” *Journal of Marketing Research*, **42** 141-156;

Boyd, D., Gordon, M., Anderson, J., Tai, C., Feng, Y., Kolamala, A., Cook, G., Guardino, T., Purang, M., Krishnamurthy, P., Cooke, M., Nandiwada, R., Monteiro, B., Haas, S. (2005).

*Manugistics Target Pricing System* US Patent No: 6963854;

Caufield, S. (2009). "Pricing of Consumer Credit," *The Handbook of Pricing Management*. Oxford, UK: Oxford University Press, Oxford, UK;

Dobson, A. J. and Barnett, A. (2008). *An Introduction to Generalized Linear Models*. Third Edition. Boca Raton, FL: Chapman and Hall/CRC;

Gill, J. (2008). *Bayesian Methods: A Social and Behavioral Sciences Approach*. Second Edition. Boca Raton, FL.: Chapman & Hall/CRC;

Lindsey, J. K. (2000). *Applying Generalized Linear Models*. New York, NY: Springer;

Makridakis, S., Wheelwright, S. C., and Hyndman, R.J. (1998). *Forecasting: Methods and Application*. 3rd Edition. Boston, MA.: John Wiley and Sons;

McCullagh, P., Nedler, J. A. (1989). *Generalized Linear Models* 2nd Edition. Boca Raton, FL.: Chapman and Hall/CRC;

Marn, M. V., Roegner, E. V., and Zawada, C. C. (2004). *The Price Advantage*. Hoboken, NJ.: John Wiley and Sons;

Maxwell, S. (2008). *The Price is Wrong: Understanding what makes a Price Seem Fair and the True Cost of Unfair Pricing*. Hoboken, NJ.: John Wiley & Sons;

Phillips, R.L. (2005). *Pricing and Revenue Optimization*. Stanford, CA: Stanford University Press;

Phillips, R. L. (2010). “An Introduction to Price Optimization for Consumer Credit”. *Credit Control*. Forthcoming;

Phillips, R. L. and Raffard, R. (2010). “Adverse Selection in Consumer Lending,”. Working paper. Available at [www.nomissolutions.com](http://www.nomissolutions.com);

Raiffa, H., Richardson, J., and Metcalfe, D. (2003). *Negotiation Analysis: The Science and Art of Collaborative Decision Making*. Harvard University Press, Cambridge, MA;

Ross, S. L. and Yinger, J. (2002). *The Color of Credit: Mortgage Discrimination, Research Methodology, and Fair-Lending Enforcement*. The MIT Press. Cambridge, MA.: The MIT Press;

Stiglitz, J. E. and Weiss, A. (1981). “Credit Rationing in Markets with Imperfect Information,” *American Economic Review*, **71** 3, pp. 393 - 410;

Theil, H. (1948). “A Static Theory of Entrepreneurial Behavior,” Master’s Thesis, University of Amsterdam, Amsterdam (in Dutch).