Optimal Timing of Price-Quote Revisions*

Miguel Sousa Lobo†

September 7, 2007

Abstract

We propose a model for the timing of individual discounts, where a seller makes sequential offers to a buyer with private valuation for a good. After receiving a price quote below the valuation, the buyer has a random delay until purchase. A buyer for whom the quoted price is above the valuation has to search for an alternative and is not immediately lost to the seller. By timing the price-quote revision, the seller achieves partial third-degree price discrimination. Significant gains can be obtained by the seller, even when the demand function is unknown. The optimal policy is characterized, and conditions on the random-time distributions are provided under which it is unique. The model predicts that, in less competitive markets, price negotiation is more prevalent, with a higher dispersion of transacted prices and more protracted negotiations. We discuss practical estimation procedures for the model parameters and numerical examples.

*First draft: July 5, 2007

†Duke University, Fuqua School of Business, email: `mlobo@duke.edu`

1 Introduction

The practice of offering individual discounts to buyers seems to be more common in some industries than others, and some sellers quicker than others to offer a discount as an incentive for a specific buyer to complete a transaction. How can the seller determine the optimal timing of individual discounts, and how deep should the eventual discount be? While there exists an extensive literature that relates to this problem, particularly in the fields of industrial organization and bargaining, it has received little attention in the revenue management literature which has mostly focused on posted-price transactions.

While it is difficult to obtain a precise estimate for the number and total value of transactions in which the price is negotiated, this seems to be the dominant mode in business-to-business transactions (see, *e.g.*, Boyd (2007b) and Boyd (2007a)). Such transactions from one firm to the next along the value-chain account, by value, for the bulk of transactions in the economy. Hancock and Humphrey (1997) cite a range of estimates from 50 to 83% for the fraction of transactions in the US economy that are made in cash, but note that these transactions account for less than 1% of payment expenditures, and that 79.2 billion non-cash payments totaling more than 750 trillion dollars were made in 1994. By comparison, based on Bureau of Economic Analysis statistics, personal consumption expenditures in the US in 1994 amounted to less than 5 trillion dollars.

There is strong support, both empirical and from theory, to believe that it can be of benefit to a seller to negotiate prices individually with each buyer, and that the equilibirum strategy can result in a negotiation delay in transactions. The main goal of this article is to take a step in operationalizing the problem from the seller's side. We discuss how the seller can estimate the aggregate buyers' strategy from sales history, and determine the optimal response under the practical constraint of discrete price revisions.

Our model is motivated by business-to-business transactions of medium or large value. For transaction of this type, the seller typically does not publicly post a price. A buyer who is interested in purchasing a given good is required to contact the seller for a price quote. We assume that the seller and buyer are not engaged in repeated interaction, and that the seller can maintain the price and cost information private. In fact, in large transactions, sellers often require buyers to contractually commit to maintain the details of the transaction private. For the most part, the prices to be quoted are taken as fixed and inventory considerations disregarded. Solving for the optimal timing of price-quote revisions, we find that the seller can obtain substantial gains in expected revenue from correctly managing individual discounts.

Consider, say, a small company that calls a computer equipment manufacturer asking for a price quote for a high-end server with a particular configuration. After receiving the quote, evaluating the offer, and possibly considering alternative suppliers, the decision-maker at the buying firm may, after some time has elapsed, authorize the terms of the purchase. However if, say, a week has elapsed and the buyer has not made the purchase yet, the seller may call back to inquire about the buyer's interest and offer a lower price. It is now common for call-center software to record all customer contact information: when the buyer first asked for a price quote; what price was quoted; when subsequent contacts occurred; if and when the quote was revised and to what price; whether the buyer purchased the item and, if so, when and for what price.

A key feature of this example, and a motivation for our model, is that the buyer's evaluation and decision process is not instantaneous, nor is the search for alternatives. In markets with limited information and imperfect competition, the seller will have significant monopoly power. This is often the case for specialized inputs to production or specialized services, say an OEM negotiating the pricing for the supply of a semiconductor component of which there are few suppliers. However, some large consumer purchases also fit this framework, such as in car dealerships, where the search for alternatives is time-consuming and costly. Car salesmen will typically offer individual discounts, but only after significant time and effort has been expended by the buyer.

Models used in the revenue management literature generally presume the availability of at least approximate estimates (or, more commonly, exact knowledge) of the demand function. However, if prices are set on an individual basis, the price-elasticity of demand cannot be estimated from historical data. The number of sales at each price do not necessarily reflect the distribution of valuations but, rather, the outcomes of the negotiation processes, which is to say, how often salespeople settled on different price discounts. However, given a sufficient number of sales, much information is available that makes the problem amenable to a quantitative revenue management approach. If information about all customer contacts is extant, the seller can estimate the probability distribution of the delay in buyer purchases, and estimate the effect of an earlier or later timing of the offer of a discount. This can be interpreted as estimating a *time-elasticity of demand*.

In the examples above, the seller has most of the control over the negotiation process. We make the modeling assumption that only the seller makes offers. We also assume that the buyer's valuation for the good and the buyer's expected time until purchase are private. A fraction of buyers have high valuation for the good and are willing to pay the seller's initial price quote. Another fraction of buyers has a lower valuation, and will only buy at a reduced price. Both types of buyer conduct a search for alternatives. From the point of view of the seller, there is a random distribution for when the buyer terminates the search. Variability in the purchasing delay is likely to arise for a number of reasons. Different buyers will be operating with different internal decision-making procedures and with different time-constraints. There may be heterogeneity in search costs and in beliefs about the probability of finding a preferred alternative. Also, different buyers may make different strategic calculi about the likelihood that the seller will offer a lower price.

High-valuation buyers may either find an alternative or accept the seller's offer. Low-valuation buyers can only terminate their search if a lower-cost alternative is found. It can then be expected that high-valuation buyers will terminate the search at a higher rate than low-valuation buyers. If this is the case, conditional on a buyer not having yet made a purchase, the probability that the buyer is low-valuation increases over time, and the seller has an increasing incentive to offer a lower price. A partial form of third-degree price discrimination is achieved by making low-valuation buyers wait for a discount. Our analysis shows, with plausible model parameters, significant gains in expected revenue for the seller from such a policy and, which is crucial to justify the widespread use of negotiated prices and individual discounts, that these gains can be robust to model uncertainty.

We make the pragmatic assumption that buyers are not strategic, which we argue is adequate for a revenue management model here. More precisely, we assume that the distribution of high-valuation buyers' time-to-purchase already incorporates, along with other considerations, the buyers' assessment of the probability that the seller will offer a lowered price within an acceptable time-frame. For the transactions under consideration, with limited public information and where interactions are infrequent, there is in any case little scope for learning on the part of the buyers. Also, since transacted prices are assumed to be private, high-valuation buyers who have shorter purchase delays never see the individual discount and, unless they conduct exploration, cannot learn the seller's policy. While it would be straigthforward to add to our model, we assume no discounting on the part of the buyer, a simplification which allows us to keep results and notation simpler. In most transactions, the negotiation delay is likely to be too short for any sizeable discounting to apply.

Coase (1972) conjectures that a durable-goods monopolist who is unable to commit to not reducing prices in the future loses all market power. Buyers would otherwise wait for the price reduction, and in equilibrium there is no bargaining delay in transactions. However, in addition to the empirical fact that protracted negotiations and individual discounts are ubiquitous, a substantial literature in sequential bargaining has established that delay can occur for diverse reasons. In simpler discrete-time models, price differentiation over time arises only with discounting. For instance, Sobel and Takahashi (1983) solve the problem of sequential bargaining with two time periods with private information and discounting, where only the seller makes offers. Gul, Sonnenschein, and Wilson (1986) consider a sequence of finite-time periods, including the case where the period-length goes to zero. They find that, in continuous time, there is no delay in the transaction, verifying the Coase conjecture. Admati and Perry (1987) consider the case where players can credibly commit to delay and, even in continuous time, find equilibria where the players use strategic delays. Delay in the transaction is found to be the equilibrium outcome of several other models. Abreu and Gul (2000) study the reputation effects of obstinate "irrational types". Yildiz (2004) considers players with Bayesian learning who are optimistic about their bargaining power. Cho (1990) finds equilibria with delay in a model with two-sided uncertainty where the seller makes offers. Other work, for instance Deneckere and Liang (2006), has looked at the case where the seller's cost and buyer's valuation are interdependent. Perhaps most closely related to our model, Fuchs and Skrzypacz (2006) allow for the random arrival of alternative buyers and sellers, which relates to the search for alternative suppliers in our model, and find that the equilibrium strategy in continuous time involves delay. Our model can be seen as an operational version from the point of view of the seller, with the goal of finding the seller's optimal strategy given the available data to estimate the strategy currently being played by the buyer. We formulate the problem in continuous time, and the initial assumption of two prices is later extended to an arbitrary number or prices.

Models of search have received attention in the consumer behavior literature, and to some extent in the industrial organization literature. This work, while related to our problem, does not directly apply to it, as existing models assume that the seller does not have repeated access to buyer. Each customer either buys the good or goes look for another seller. For instance, Stahl (1989) looks at the optimal price as a function of search cost, with consumers who conduct a sequential search. A model of price promotions with consumer search is proposed in Banks and Moorthy (1999), who find that higher search costs result in deeper promotions. In this model, consumers search for promotions and there is no decision on the part of the seller on whether to offer a discount to a particular buyer. In the price-discrimination literature, Salop (1977) showed that price dispersion can be profitably used to obtain price discrimination, with more efficient searchers obtaining a lower price (for price-dicrimination references see also, *e.g.*, Phlips

(1988) and Armstrong (2006)).

Inter-temporal price discrimination is also a feature of price-skimming models (*e.g.*, Stokey (1979), Besanko and Winston (1990)). In this literature it is generally assumed that the buyer has exact knowledge of the seller's price schedule (or can solve or learn the equilibrium schedule). There is no uncertainty about the buyer's valuation and no delay in purchase once the price drops below the valuation. Sales occur continuously over time only if the price is continuously decreasing. High-valuation buyers will delay their purchase if the rate of decrease in the price is too slow, so that how quickly the price is lowered is determined by the buyer's discount rate, rather than being due to the search for and availability of alternatives.

Another related field is that of bargaining, both game-theoretic and behavioral (*i.e.*, what negotiators 'actually do'). See Hausken (1997) for a review. In the context of revenue management, Bhandari and Secomandi (2007), with a simplified model of negotiation, endogeneize the calculation of the best alternative to a negotiated agreement based on depletion of inventory, in discrete time with finite horizon, and derive some stuctural properties of the resulting Markov decision process.

Other recent revenue management work has looked at dynamic pricing and timing issues. In most of this work, the price at any give time is common to all buyers. It can be argued that this is not the case for auctions, such as in Vulcano, van Ryzin, and Maglaras (2002), who consider sequential auctions with an inventory constraint. Netessine (2004) considers the dynamic pricing of inventory where a limited number of price changes is allowed. Ovchinnikov and Milner (2005) consider last-minute deals with a strategic response from customers who learn to wait for the discount. Other recent work has sought to introduce behavioral models. Popescu and Wu (2007) study a dynamic model (where strategic interaction is not considered) with consumers that use the good's price history to form their valuation.

Contributions

We introduce a model for managing price-quote revisions in transactions where the price is individually negotiated with each buyer, as is the case for most business-to-business transactions. A number of structural and technical results are of interest. We provide a detailed characterization of the two-price model, including a closed-form solution under the assumption of exponential distributions on the random times. The model provides support for a result which may be counter-intuitive, but is consistent with anecdotal evidence; that if altenative suppliers are harder to find and the seller has more monopoly power, then the seller is more likely to offer substantial individual discounts to some buyers.

For the multiple-price problem, we provide a dynamic-programming formulation, and establish detailed structural properties. We are able to propagate quasi-concavity of the maximand in the decision variable by propagating a number of properties of the value function including, somewhat surprisingly, convexity in one of the parameters.

We also discuss implementation and statistical issues of identifiability and estimation, which are often neglected. Models are only usable (and, arguably, can only be realistic representations of reality) to the extent that players can estimate the model parameters, and to the extent that the results are robust to uncertain knowledge of those parameters. While we find that there are identifiability concerns in the problem under study, with the appropriate estimation procedures, and with decisions that take the uncertainty into account, the seller can nevertheless derive substantial gains in expected revenue.

Organization of the article

The article is organized as follows. Section 2 introduces the model and characterizes the optimal quote-revision times, first for two prices with abitrary probability distributions for the random times, then with exponential distributions for two prices and for an abritrary number of prices. Section 3 considers the case where the demand function is known and the seller optimally selects the quote prices jointly with the quote-revision time. In Section 4 we discuss practical estimation procedures for the model parameters, and the optimal selection of the quote-revision time based on the probability distribution of the model parameters conditional on a set of historical observations. The discussion includes numerical examples. Section 5 provides some final remarks.

2 Model

We first introduce the model for the problem of timing the change between two prices. For the case of general distributions we argue that, with reasonable conditions on the probability distributions, the solution of the timing problem is unique. More detailed results are then provided for the two-price problem with exponential distributions on the random times, and an example is described. We conclude the section with the problem of timing the price-quote revisions with multiple prices, and with exponential distributions.

2.1 Two prices with general distributions

A buyer requests a price quote for a good at $t = 0$. The seller quotes price $\pi_1 \in]0, +\infty[$. After a period of time $\tau_1 \in [0, +\infty]$ has elapsed, if the customer has not yet made a purchase, the seller revises the price quote to $\pi_2 \in]0, \pi_1[$. The quote-revision time τ_1 is 0 if the initial quote is π_2, and $+\infty$ if π_2 is never quoted. The buyer's valuation for the good is private. The probabilities that the buyer will purchase the good at each price are as follows.

- If the price is held constant at π_1, a fraction $p_1 \in [0, 1]$ of buyers will eventually purchase the good.
- If the price is held constant at π_2, a fraction $p_1 + p_2$ of buyers will eventually purchase the good, where $p_2 \in [0, 1 - p_1]$.

The p_i define the values of the effective demand curve at two specific points. The π_i, which specify those points, can be alternatively defined as the profit contribution of a sale at each price. The problem will then be one of maximizing the expected variable profit rather than maximizing revenue. Note that the purchase probabilities p_1 and p_2 are a function not only of the distribution of buyer valuations for the good, but also of the buyers' search behavior and of the availability of alternatives. A buyer may request, say, two or three price quotes from different sellers and then select the most favorable. This search for alternatives may be costly, time-consuming, of uncertain outcome, and extended by strategic delay on the part of the buyer. We define the following two random time distributions.

- For buyers who would purchase the good at price π_1, $F_1(t)$ specifies the cumulative probability of purchase by time t.
- For buyers who would purchase the good at price π_2 but not at price π_1, $F_2(t)$ specifies the cumulative probability that customer is lost to an alternative if price π_2 is not offered by time t.

The time variable may be measured in, say, business days. It may also be measured as a function of the interaction history, taking into account the number of customer contacts, which may be indicative of the intensity of the search being performed by the buyer. Note that no individual buyer need behave according to the specified distributions. Rather, they describe the aggregated distribution over the behavior of the population of potential buyers.

If the quoted price is revised at time τ_1, the expected revenue to the seller is $\mathbf{E}R = \pi_1 p_1 F_1(\tau_1) + \pi_2 \left[p_1(1 - F_1(\tau_1)) + p_2(1 - F_2(\tau_1))\right]$. The optimal-timing problem faced by the seller is to

$$\text{maximize} \quad \pi_2 p_1 + (\pi_1 - \pi_2) p_1 F_1(\tau_1) + \pi_2 p_2 (1 - F_2(\tau_1)) \tag{1}$$

over $\tau_1 \in [0, +\infty]$. This problem can, in general, have multiple maxima. Even if the global maximum is unique, it can have multiple local maxima. For the following result, we assume that the densities f_1 and f_2 are finite and continuous a. e., and denote by $f^-(t)$ and $f^+(t)$ the limits from the left and from the right of f at t.

Result 1 *If the ratio of the density f_1 of purchases by high-valuation buyers to the density f_2 of losses of low-valuation buyers is decreasing, then problem (1) is quasi-concave and has a unique solution τ_1^* characterized by*

$$(\pi_1 - \pi_2) p_1 f_1^-(\tau_1^*) - \pi_2 p_2 f_2^-(\tau_1^*) \geq 0 \geq (\pi_1 - \pi_2) p_1 f_1^+(\tau_1^*) - \pi_2 p_2 f_2^+(\tau_1^*),$$

by $\tau_1^* = 0$ if $(\pi_1 - \pi_2) p_1 f_1(\tau_1) - \pi_2 p_2 f_2(\tau_1) \leq 0$ for all $\tau_1 > 0$, and by $\tau_1^* = +\infty$ if $(\pi_1 - \pi_2) p_1 f_1(\tau_1) - \pi_2 p_2 f_2(\tau_1) \geq 0$ for all $\tau_1 \geq 0$.

Proof. For continuous densities, a first-order optimality condition for a non-zero finite solution is

$$(\pi_1 - \pi_2) p_1 f_1(\tau^*) - \pi_2 p_2 f_2(\tau^*) = 0. \tag{2}$$

If f_1/f_2 is decreasing, then (2) has at most one solution, is positive to the left of that solution and negative to the right, and is necessary and sufficient for non-zero finite solutions. The maximand is increasing to the left of this solution and decreasing to the right. The extension to a. e. continuous densities is straightforward. ■

A decreasing f_1/f_2 ratio is required for the first-order condition in Result 1 to be necessary and sufficient for all choices of problem parameters π_1, π_2, p_1 and p_2. If, further, the densities are differentiable, the assumption that f_1/f_2 is decreasing can be stated as

$$\frac{-f_1'(t)}{f_1(t)} > \frac{-f_2'(t)}{f_2(t)}, \quad \text{for all } t \in [0, +\infty[. \tag{3}$$

That is, the density f_1 of purchases by high-valuation buyers decays faster than does the density f_2 of losses of low-valuation buyers to alternatives. This is reasonable with a model where buyers conduct, over time, a search for alternatives. Buyers with a valuation above the quoted price can exit by purchasing from the seller, and they can also exit if a preferable alternative is found. Buyers with a valuation below the price being offered, on the other hand, can only exit if an acceptable alternative is found. High-valuation buyers should therefore exit at a higher rate than low-valuation buyers. This relates also to the following result.

Result 2 *If the hazard rate for the purchase by high-valuation buyers is greater than the hazard rate for the loss to alternatives of low-valuation buyers, it becomes increasingly probable over time that a customer who has not yet purchased is a low-valuation buyer. This condition is implied by f_1/f_2 decreasing.*

Proof. Consider that the quoted price at time t is still π_1, and that the buyer has not yet purchased the good. Conditional on the buyer being such that she would purchase the good if price π_2 was offered, the probability that the buyer at time t is a low-valuation buyer is

$$\mathbf{P}_2(t) = \frac{p_2(1 - F_2(t))}{p_1(1 - F_1(t)) + p_2(1 - F_2(t))}.$$

This is increasing if

$$-f_2(t)(1 - F_1(t)) + (1 - F_2(t))f_1(t) > 0,$$

or

$$\frac{f_1(t)}{1 - F_1(t)} > \frac{f_2(t)}{1 - F_2(t)}. \tag{4}$$

Note that this condition on the hazard rates is not sufficient for the solution of (1) to be unique (counter-examples can be constructed), but is implied by f_1/f_2 decreasing. This is shown by seeing that if (4) is violated at some t_1, there must exist a $t_2 > t_1$ such that $f_1(t_2)/f_2(t_2) > f_1(t_1)/f_2(t_1)$, so that f_1/f_2 must be increasing at some $t \in [t_1, t_2]$. ■

Note that, should conditions with the reverse inequality hold, the probability that a buyer is high-valuation would increase over time. It would then be optimal to increase the price over time, rather than decrease. Increasing the price over time to a particular buyer is only feasible in practice if a deadline for the price-quote is announced, and if the buyer cannot chose to be anonymous. Attaching a deadline to a price quote ('exploding offers') and collecting detailed information about the buyer before providing a price quote are not uncommon practices. However, increasing a price quote if the purchase is delayed seems to be a far less common practice than negotiated discounts, which are the focus of our model.

2.2 Two prices with exponential distributions

We now consider the particular case of exponentially-distributed random times, for which we provide closed-form results and numerical examples. In later sections, we also use exponential distributions for the generalization to the timing of n prices, and for the generalization where the seller optimally chooses prices based on a known demand curve.

The distribution of buyer valuations for the good is such that:

- A fraction $q_1 \in [0, 1]$ of buyers have a valuation $w \in [\pi_1, +\infty[$.
- A fraction $q_2 \in [0, 1 - q_1]$ of buyers have a valuation $w \in [\pi_2, \pi_1[$.

Consider that a price quote of π_i is in effect for a customer with valuation $w \in \mathbf{R}$.

- If $w < \pi_i$, the buyer searches for an alternative. Buyers find alternatives at rate $\beta > 0$. Denoting by x the probability that the buyer is still available to the seller after time t has elapsed since the price quote, we have $dx/dt = -\beta x$, and $x = e^{-\beta t}$.
- If $w \geq \pi_i$, the buyer may find a preferred alternative, at rate β, or may give up searching for an alternative and accept the seller's quote, at rate $\alpha > 0$. Denoting by x the probability that the buyer is still searching after time t, we have $dx/dt = -(\alpha+\beta)x$, and $x = e^{-(\alpha+\beta)t}$. Denoting by z the cumulative probability that the buyer has accepted the quoted price after time t, we have $dz/dt = \alpha x$, so that $z = \frac{\alpha}{\alpha+\beta}\left(1 - e^{-(\alpha+\beta)t}\right)$.

It may be more intuitive to think of these parameters in terms of average times. The average time that a low-valuation buyer takes to either find an alternative or exit the market is $1/\beta$. For buyers who eventually accept a given price quote, the average time until the offer is accepted is $1/(\alpha+\beta)$. For example, suppose $\alpha = 0.4$ and $\beta = 0.1$, and the unit of time is one business day. If the price is acceptable, a buyer will make a decision to purchase, on average, in about two days (the good may be needed for production, the buyer may wish to complete the purchase to continue some project, *etc.*) If the price is not acceptable, a buyer will take, on average, ten days to find an alternative supplier (it is a low-competition or low-information market, there are few alternatives or high search costs).

With quote revision at time τ_1, the cumulative probability that the customer has purchased the good at price π_1 by time t is

$$F_{\pi_1}(t) = q_1 \frac{\alpha}{\alpha+\beta}\left(1 - e^{-(\alpha+\beta)\min(t,\tau_1)}\right), \tag{5}$$

and the cumulative probability that the customer has purchased the good at price π_2 by time t is

$$F_{\pi_2}(t) = \left(q_1 e^{-(\alpha+\beta)\tau_1} + q_2 e^{-\beta\tau_1}\right)\frac{\alpha}{\alpha+\beta}\left(1 - e^{-(\alpha+\beta)\max(0,t-\tau_1)}\right). \tag{6}$$

In the notation of Section 2.1, the probabilities of purchase are $p_1 = q_1\frac{\alpha}{\alpha+\beta}$ and $p_2 = q_2\frac{\alpha}{\alpha+\beta}$, and the random time distributions are $F_1(t) = 1 - e^{-(\alpha+\beta)t}$ and $F_2(t) = 1 - e^{-\beta t}$. The expected revenue to the seller with price-quote revision at time τ_1 is

$$\begin{aligned} \mathbf{E}R &= \pi_1 F_{\pi_1}(\tau_1) + \pi_2 F_{\pi_2}(+\infty) \\ &= \pi_1 q_1 \frac{\alpha}{\alpha+\beta}\left(1 - e^{-(\alpha+\beta)\tau_1}\right) + \pi_2\left(q_1 e^{-(\alpha+\beta)\tau_1} + q_2 e^{-\beta\tau_1}\right)\frac{\alpha}{\alpha+\beta} \end{aligned}$$

$$= \frac{\alpha}{\alpha+\beta}\left[\pi_2(q_1+q_2)+(\pi_1-\pi_2)q_1\left(1-e^{-(\alpha+\beta)\tau_1}\right)-\pi_2 q_2\left(1-e^{-\beta\tau_1}\right)\right]. \quad (7)$$

The first term in the square brackets is the revenue if the price is held constant at π_2 from $t = 0$. The second term is the gain from delaying the introduction of the lower price, derived from sales to high-valuation customers at price π_1. The third term is the loss from delaying the introduction of the lower price, which arises from low-valuation customers finding alternative sellers in the interim. In numerical examples we will provide as a reference the following simple upper bound on the expected revenue, which is based on the expected revenue if full third-degree price discrimination were possible.

Result 3 *An upper bound on the expected revenue given optimal timing of the quote revision is given by the expected revenue under full price discrimination,*

$$\mathbf{E}R^* \leq \frac{\alpha}{\alpha+\beta}\left(q_1\pi_1+q_2\pi_2\right).$$

Proof. From the independent maximization in τ_1 of the terms of (7). ■

We next establish that the optimal price-revision time is unique and finite, and provide closed-form expressions for the optimal price-revision time and for the optimal expected revenue.

Result 4 *The expected revenue is quasi-concave in the quote-revision time τ_1. The optimal quote-revision time is unique and finite and equal to*

$$\begin{aligned}\tau_1^* &= \operatorname*{argmax}_{\tau\geq 0}\mathbf{E}R \\ &= \max\left(0,\ \frac{1}{\alpha}\Big(\log q_1-\log q_2+\log(\pi_1-\pi_2)-\log\pi_2+\log(\alpha+\beta)-\log\beta\Big)\right). \quad (8)\end{aligned}$$

The optimal expected revenue, given optimal timing of the quote revision, is

$$\begin{aligned}\mathbf{E}R^* &= \max_{\tau_1\geq 0}\mathbf{E}R \\ &= \begin{cases}\dfrac{\alpha}{\alpha+\beta}\pi_2(q_1+q_2), & \text{if } \tau_1^*=0 \\ \dfrac{\alpha}{\alpha+\beta}\left(\pi_1 q_1+\dfrac{\alpha}{\alpha+\beta}\pi_2 q_2\left(\dfrac{\pi_2 q_2\beta}{(\pi_1-\pi_2)q_1(\alpha+\beta)}\right)^{\frac{\beta}{\alpha}}\right), & \text{if } \tau_1^*>0.\end{cases} \quad (9)\end{aligned}$$

If $\tau_1^* = 0$ the optimal policy is to never quote price π_1 and open with price π_2. The optimal quote-revision time τ_1^* is increasing with q_1, and π_1, and decreasing with q_2, π_2, and β. The optimal expected revenue $\mathbf{E}R^*$ is increasing with q_1, q_2, π_1, π_2 and α, and decreasing with β.

Proof. While quasi-concavity and uniqueness are immediate from either Result 1 or Result 5, we provide here a short alternative proof. With the monotonic transformation $z = e^{-(\alpha+\beta)\tau_1}$,

$$\mathbf{E}R = \frac{\alpha}{\alpha+\beta}\left(\pi_1 q_1-(\pi_1-\pi_2)q_1 z+\pi_2 q_2 z^{\frac{\beta}{\alpha+\beta}}\right)$$

is strictly concave in $z \in [0,1]$ and increasing near zero, so that there is a well-defined maximizer $z^* \in]0,1]$. The expected revenue is therefore strictly quasi-concave in τ_1, with a well-defined maximizer $\tau_1^* \in [0,+\infty[$. The optimal time and optimal revenue follow from the first-order optimality condition and algebraic manipulation. The proofs of the monotonicities are by differentiation and algebraic manipulation, using the condition $\tau_1^* \geq 0$ where needed. ■

When the seller is capacity-constrained, it may be optimal to set $\tau_1 > \tau_1^*$. If buyers arrive at rate μ and the seller has an inventory c to be sold by deadline T (measured from the moment of arrival of the buyer), the desired target probability of sale to the buyer is

$$\widetilde{\mathbf{P}} = \frac{c}{\mu T}.$$

The probability of a sale at either price to the buyer is

$$\mathbf{P} = \frac{\alpha}{\alpha+\beta}\left(q_1 + q_2 e^{-\beta\tau_1}\right),$$

and the quote-revision time that achieves the target probability of sale is

$$\widetilde{\tau}_1 = -\frac{1}{\beta}\log\left[\frac{1}{q_2}\left(\frac{\alpha+\beta}{\alpha}\frac{c}{\mu T} - q_1\right)\right].$$

The seller's quote-revision time is then $\max(\tau_1^*, \widetilde{\tau}_1)$. Note that this is an approximate solution, since the optimal policy should take into account the possible arrival of other buyers before τ_1 (and if an arrival does occur, the time should be revised). This policy will, however, be a good approximation if $\frac{1}{\mu} \gg \tau_1$.

Example

Consider an example. A potential buyer asks the seller for a price quote for a given good, and the seller initially quotes the price $\pi_1 = 600$. After a delay, if the buyer has not yet purchased the good, the seller may decrease the price to $\pi_2 = 100$. Consider the following the three cases:

(i) 5% of the potential buyers have a valuation above π_1, and 55% of the potential buyers have a valuation between π_2 and π_1 ($q_1 = 0.05$, $q_2 = 0.55$). There is a relatively large proportion of low-valuation buyers and, if price discrimination is not possible, π_2 is preferred.

(ii) 10% of the potential buyers have a valuation above π_1, and 50% of the potential buyers have a valuation between π_2 and π_1 ($q_1 = 0.1$, $q_2 = 0.5$). If price discrimination is not possible, π_1 and π_2 are equally profitable.

(iii) 20% of the potential buyers have a valuation above π_1, and 40% of the potential buyers have a valuation between π_2 and π_1 ($q_1 = 0.2$, $q_2 = 0.4$). There is a relatively large proportion of high-valuation buyers and, if price discrimination is not possible, π_1 is preferred.

In Figure 1, we plot the expected revenue as a function of the time at which the quote is revised from π_1 to π_2. Without loss of generality, we set $\alpha = 1$. For each case we plot the expected revenue for β equal to 0.2, 0.5, 1, 2, and 5. When $\beta = 0.2$, a buyer whose valuation is above the quoted price accepts the seller's offer at a rate five times higher than the rate at which she finds an alternative solution (the seller has significant monopoly power, buyers' search costs are high, *etc.*) When $\beta = 5$, a buyer whose valuation is above the quoted price accepts the seller's offer at a rate five times lower than the rate at which she finds an alternative solution (the market is competitive, alternatives are readily available search costs are low, *etc.*)

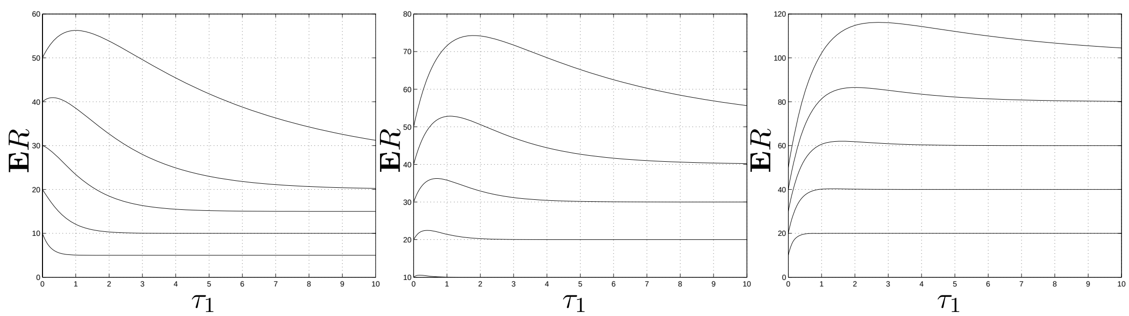

Figure 1: Expected revenue $\mathbf{E}R$ as a function of τ_1, the timing of the price-quote revision, for cases (i) large proportion of low-valuation buyers, (ii) balanced, and (iii) large proportion of high-valuation buyers. Each is plotted for different levels of β, the availability of alternatives to buyers.

When there are many low-valuation buyers and alternatives are readily available, it is optimal to offer the low price π_2 as the initial quote. The expected gain of the optimal policy relative to the best static posted price is then zero. Table 1 summarizes the gains that the seller obtains from optimally timing the price-quote revision. The percentual gain in expected revenue is compared to the expected revenue with constant pricing (at π_2 in case (i), and at π_1 in case (iii); note that, in this example, if the suboptimal constant price is chosen the percentual gains are identical in case (ii) but significantly larger in cases (i) and (iii), on the order of 100%).

2.3 Multiple prices

Consider the extension of the problem with exponential distributions to a sequence of n price quotations $\pi_1 > \pi_2 > \ldots \pi_n > 0$. Price π_i is valid for a period of time τ_i, and then replaced by price π_{i+1}. Defining $t_0 = 0$ and $t_i = \tau_1 + \tau_2 + \ldots + \tau_i$ for $i = 1, \ldots, n$, the price quote π_i is in effect during $t \in [t_{i-1}, t_i[$. If $\tau_1, \ldots, \tau_{i-1}$ are finite and $\tau_i = +\infty$, then π_i is the ending price and $\pi_{i+1}, \ldots, \pi_n$ are never quoted. No sales are allowed after t_n, although will see that, with

	(i) $q_1 = 0.05$ $q_2 = 0.55$	(ii) $q_1 = 0.10$ $q_2 = 0.50$	(iii) $q_1 = 0.20$ $q_2 = 0.40$
Upper bound with full price discrimination	41.7%	83.3%	33.3%
$\beta = 0.2$	12.5%	48.5%	16.2%
$\beta = 0.5$	2.3%	32.1%	8.1%
$\beta = 1$	0.0%	20.8%	3.3%
$\beta = 2$	0.0%	12.3%	0.8%
$\beta = 5$	0.0%	5.6%	0.0%

Table 1: Gains from full and partial price discrimination, for different distributions of customer valuations, and different availability of alternatives (these are the cases plotted in Figure 1).

optimal timing, $t_n^* = +\infty$ and all the other t_i^* are finite.

The distribution of random times is specified as before. Buyers with feasible valuation purchase the good at rate α, and are lost to alternatives at rate β. Buyers with infeasible valuation are lost to alternatives at rate β. The distribution of valuations for the good is such that a fraction q_i of the buyers has a valuation $w \in [\pi_i, \pi_{i-1}[$ (where we define $\pi_0 = +\infty$). We simplify notation by defining the probabilities of purchase $p_i = \frac{\alpha}{\alpha+\beta} q_i$. The expected revenue function $V_n : \mathcal{P}_n \times [0, +\infty[^n \times [0, +\infty]^n \to \mathbf{R}$, where $\mathcal{P}_n = \{\pi \in \mathbf{R}^n \mid \pi_1 > \pi_2 > \dots \pi_n > 0\}$, is then

$$V_n(\pi, p, \tau) = \sum_{i=1}^{n} p_i e^{-\beta(\tau_1+\ldots+\tau_{i-1})} \sum_{j=i}^{n} \pi_j e^{-(\alpha+\beta)(\tau_i+\cdots+\tau_{j-1})} \big(1 - e^{-(\alpha+\beta)\tau_j}\big). \tag{10}$$

The optimal expected revenue function $R : \mathcal{P}_n \times [0, +\infty[^n \to \mathbf{R}$ is defined as

$$R_n(\pi, p) = \max_{\tau \in [0, +\infty]^n} V_n(\pi, p, \tau) = V_n(\pi, p, \tau^*(\pi, p)).$$

We will show that the maximum exists and is unique, so that the maximizing $\tau^*(\pi, q) : \mathcal{P}_n \times [0, +\infty[^n \to [0, +\infty]^n$ is well defined. The remainder of this section is the proof of the following result, with ancillary results stated as lemmas where the structural properties are of independent interest.

Result 5 *The optimal times τ^* are unique, with $\tau_1^*, \ldots, \tau_{n-1}^*$ finite, and $\tau_n^* = +\infty$.*

Proof. We rewrite the problem in terms of the optimality equations

$$R_n(\pi, p) = \max_{\tau_1 > 0} \ \pi_1 p_1 \big(1 - e^{-(\alpha+\beta)\tau_1}\big) + R_{n-1}\big(\widetilde{\pi}(\pi), \widetilde{p}(p, \tau_1)\big), \tag{11}$$

where

$$\widetilde{\pi}(\pi) = (\pi_2, \pi_3, \ldots, \pi_n)^T, \tag{12}$$

$$\widetilde{p}(p, \tau_1) = \Big(p_1 e^{-(\alpha+\beta)\tau_1} + p_2 e^{-\beta\tau_1}, p_3 e^{-\beta\tau_1}, \ldots, p_n e^{-\beta\tau_1}\Big)^T. \tag{13}$$

Also,

$$R_1(\pi, p) = \pi p, \tag{14}$$

since, for $n = 1$, we have trivially that $\tau_1^* = +\infty$. We then have, for the general problem, $\tau_n^* = +\infty$. Before completing the proof, we establish the following four lemmas, all of which assume $n \geq 2$.

Lemma 1 *For any non-negative $a, b \in \mathbf{R}$,*

$$R_n(a\,\pi, b\,p) = a\,b\,R_n(\pi, p). \tag{15}$$

Proof. Immediate from (10), or from (14) and the recursive application of (11), (12) and (13).
■

Lemma 2 *The optimal times τ^* depend on p_1 as follows.*

a) The optimal time τ_1^* is of the form

$$\tau_1^* = \max\Big(0, \frac{1}{\alpha}\Big(\log p_1 - \log \underline{p}_1\Big)\Big),$$

for some $\underline{p}_1 > 0$, which is a function of $p_2, \ldots, p_n$ but not of p_1. Note that if $\tau_1^* > 0$, τ_1^* is increasing concave in p_1.

b) If $\tau_1^* > 0$, $\tau_2^*, \ldots \tau_n^*$ are constant in p_1.

Proof. We rewrite the maximand in R_n as

$$S_n = \pi_1 p_1 \big(1 - e^{-(\alpha+\beta)\tau_1}\big) + e^{-\beta\tau_1} R_{n-1}\big(q_1 e^{-\alpha\tau_1} + q_2, q_3, \ldots, q_n\big)$$

(we will from here on omit $\widetilde{\pi}$ in the parameters of R_{n-1}). The first-order condition is

$$\begin{aligned} \pi_1 p_1 (\alpha+\beta) e^{-(\alpha+\beta)\tau_1} - \beta e^{-\beta\tau_1} R_{n-1}\big(p_1 e^{-\alpha\tau_1} + p_2, p_3, \ldots, p_n\big) \\ - p_1 \alpha e^{-(\alpha+\beta)\tau_1} \frac{dR_{n-1}}{dp_1}\big(p_1 e^{-\alpha\tau_1} + p_2, p_3, \ldots, p_n\big) = 0, \end{aligned}$$

which, assuming τ_1^* is unique and finite, is necessary and sufficient (we later apply this lemma to R_{n-1} to show that the optimal time τ_1^* for R_n is in fact unique and finite). By Lemma 1, the first-order condition is equivalent to

$$\begin{aligned} \pi_1 (\alpha+\beta) f(p_1, \tau_1) - \beta R_{n-1}\big(f(p_1, \tau_1) + p_2, p_3, \ldots, p_n\big) \\ - f(p_1, \tau_1) \frac{dR_{n-1}}{dp_1}\big(f(p_1, \tau_1) + p_2, p_3, \ldots, p_n\big) = 0, \end{aligned}$$

with $f(p_1, \tau_1) = p_1 e^{-\alpha\tau_1}$. If the first-order condition is satisfied by p_1^0 and τ_1^0, it is also satisfied by p_1 and τ_1^* such that $f(p_1, \tau_1^*) = f(p_1^0, \tau_1^0)$, which leads to

$$\tau_1^* = \tau_1^0 + \frac{1}{\alpha}\left(\log p_1 - \log p_1^0\right).$$

Finally, since except for a uniform scaling by $e^{-\alpha\tau_1}$ the $\widetilde{p}$ vector for R_{n-1} does not depend on p_1, we conclude from Lemma 1 that the optimal times $\tau_2^*, \tau_3^*, \ldots$ are constant in p_1. ■

Without loss of generality, we will from here on assume that all times are strictly positive. If the optimal τ_i^* is zero, we can redefine p as $[p_1, \ldots, p_{i-1}, p_i + p_{i+1}, p_{i+1}, \ldots, p_n]$ (that is, we assign valuation π_{i+1} to buyers with valuation π_i).

Lemma 3 *R_n is increasing convex in p_1.*

Proof. The full derivative of $R_n(\pi, p) = V_n(\pi, p, \tau^*(\pi, p))$ with respect to p_1 is

$$\begin{aligned}
\frac{dR_n}{dp_1} &= \frac{dV_n}{dp_1} + \sum_{j=1}^{n} \frac{d\tau_j^*}{dp_1}\frac{dV_n}{d\tau_j} \\
&= \frac{dV_n}{dp_1} \\
&= \pi_1\left(1 - e^{-(\alpha+\beta)\tau_1^*}\right) + \pi_2\left(e^{-(\alpha+\beta)\tau_1^*} - e^{-(\alpha+\beta)(\tau_1^*+\tau_2^*)}\right) + \ldots \qquad (16) \\
&= \pi_1 - (\pi_1 - \pi_2)e^{-(\alpha+\beta)\tau_1^*} - (\pi_2 - \pi_3)e^{-(\alpha+\beta)(\tau_1^*+\tau_2^*)} - \ldots \qquad (17)
\end{aligned}$$

The monotonicity of R_n is obtained from the positivity of (16). Since, from Lemma 2, τ_1^* is increasing in p_1 and $\tau_2^*, \tau_3^*, \ldots$ are constant, we see from (17) that $\frac{dR_n}{dp_1}$ is increasing in p_1, so that R_n is convex in p_1. ■

Lemma 4 *If $\tau_1^* > 0$,*

$$R_n = \pi_1 p_1 + c_n p_1^{-\frac{\beta}{\alpha}}, \qquad (18)$$

where c_n is non-negative and does not depend on p_1.

Proof. Using Lemma 2,

$$\begin{aligned}
R_n &= S_n(\tau_1 *) \\
&= \pi_1 p_1\left(1 - e^{-(\alpha+\beta)\frac{1}{\alpha}(\log p_1 - \log \underline{p}_1)}\right) + e^{-\beta\frac{1}{\alpha}(\log p_1 - \log \underline{p}_1)} R_{n-1}\left(p_1 e^{-\alpha\frac{1}{\alpha}(\log p_1 - \log \underline{p}_1)} + p_2, p_3, \ldots, p_n\right) \\
&= \pi_1 p_1\left(1 - \underline{p}_1^{\frac{\alpha+\beta}{\alpha}} p_1^{-\frac{\alpha+\beta}{\alpha}}\right) + \underline{p}_1^{\frac{\beta}{\alpha}} p_1^{-\frac{\beta}{\alpha}} R_{n-1}\left(\underline{p}_1 + p_2, p_3, \ldots, p_n\right) \\
&= \pi_1 p_1 + c_n p_1^{-\frac{\beta}{\alpha}},
\end{aligned}$$

where

$$c_n = -\pi_1 \underline{p}_1^{\frac{\alpha+\beta}{\alpha}} + \underline{p}_1^{\frac{\beta}{\alpha}} R_{n-1}\left(\underline{p}_1 + p_2, p_3, \ldots, p_n\right)$$

is non-negative from Lemma 3. ■

We now establish that τ_1^* is unique and finite. Applying Lemma 4 to R_{n-1}, we can write S_n as a function of τ_1 as

$$\begin{aligned} S_n &= \pi_1 p_1 \left(1 - e^{-(\alpha+\beta)\tau_1}\right) + e^{-\beta\tau_1}\left[\pi_2\left(q_1 e^{-\alpha\tau_1} + q_2\right) + c_{n-1}\left(q_1 e^{-\alpha\tau_1} + q_2\right)^{-\frac{\beta}{\alpha}}\right] \\ &= \pi_1 p_1 - (\pi_1 - \pi_2) p_1 e^{-(\alpha+\beta)\tau_1} + \pi_2 p_2 e^{-\beta\tau_1} + c_{n-1}\left(p_1 + p_2 e^{\alpha\tau_1}\right)^{-\frac{\beta}{\alpha}}. \end{aligned}$$

The derivative is

$$\frac{dS_n}{d\tau_1} = (\alpha+\beta)(\pi_1 - \pi_2) p_1 e^{-(\alpha+\beta)\tau_1} - \beta\pi_2 p_2 e^{-\beta\tau_1} - \beta p_2 c_{n-1} e^{\alpha\tau_1}\left(p_1 + p_2 e^{\alpha\tau_1}\right)^{-\frac{\beta}{\alpha}}.$$

Taking the limits $\tau_1 \to +\infty$ and $\tau_1 \to -\infty$, we see that $\frac{dS_n}{d\tau_1}$ is negative for τ_1 sufficiently large, and positive for τ_1 sufficiently small, so that the solution must be finite. The first-order condition is then equivalent to

$$k_1 e^{\alpha\tau_1} + k_2 e^{(2\alpha+\beta)\tau_1}\left(p_1 + p_2 e^{\alpha\tau_1}\right)^{-\frac{\alpha+\beta}{\alpha}} = 1,$$

where $k_1 = \frac{\beta\pi_2 p_2}{(\alpha+\beta)(\pi_1-\pi_2)p_1}$ and $k_2 = \frac{\beta p_2 c_{n-1}}{(\alpha+\beta)(\pi_1-\pi_2)p_1}$ are positive constants. Rewriting this as

$$k_1 e^{\alpha\tau_1} + k_2\left(p_1 e^{-\frac{\alpha(2\alpha+\beta)}{\alpha+\beta}\tau_1} + p_1 e^{-\frac{\alpha^2}{\alpha+\beta}\tau_1}\right)^{-\frac{\alpha+\beta}{\alpha}} = 1,$$

we see that the left-hand side is increasing in τ_1, and the solution must be unique. If the solution to the first-order condition is negative, then $\tau_1^* = 0$, and we can redefine the problem as discussed above. We can conclude by induction that, for the n-price problem, τ^* is unique, $\tau_1^*, \ldots, \tau_{n-1}^*$ are finite, and $\tau_n^* = +\infty$. This completes the proof of Result 5. ■

3 Jointly selecting prices and timing with a known demand curve

Consider again the two-price problem with exponential distributions for the random times of (7), and where the entire distribution of buyer valuations is known to the seller. The q_i are then a known function of the prices π_i, and the seller accordingly selects $(\pi_1^*, \pi_2^*, \tau_1^*)$, two sequential quote prices that are jointly optimal with a quote-revision time. We first discuss the general case, then provide a numerical example with a linear demand function. Some insights are derived from considering the optimal policy and expected revenue for limit values of the parameter β, the availability of alternatives to buyers.

We make the following assumptions about the distribution of valuations $F_w(w)$. The single-price revenue function $r_1(w) = wF_w(w)$ has a unique maximizer w^* (the *monopoly price*), in the sense that for every $\delta > 0$ there is an $\epsilon > 0$ such that $r_1(w^*) - r_1(w) < \epsilon \Rightarrow |w - w^*| < \delta$. The two-price, full-discrimination revenue function $r_2(w_1, w_2) = w_1(F_w(w_1) - F_w(w_2)) + w_2 F_w(w_2)$ has a unique maximizer (w_1^*, w_2^*), in the sense that for every $\delta > 0$ there is an $\epsilon > 0$ such that $r_2(w_1^*, w_2^*) - r_2(w_1, w_2) < \epsilon \Rightarrow \|(w_1, w_2) - (w_1^*, w_2^*)\| < \delta$.

Substituting $1 - F_w(\pi_1)$ for q_1 and $F_w(\pi_1) - F_w(\pi_2)$ for q_2 in (9), the resulting expression is maximized to obtain π_1^* and π_2^*. The maximization should take into account the domain of validity for each expression, which is computed with the same substitutions for q_1 and q_2 in the expression for τ_1^* given by (8). Note that, assuming F_w continuous and with finite expectation, $(\pi_1^*, \pi_2^*, \tau_1^*)$ is well-defined and finite. The maximum of (9) is seen to be interior for any α and β by verifying that $\mathbf{E}R$ is either equal to zero or positive and increasing on the boundaries of the set $\mathcal{P}_2$ ($q_2/(\pi_1 - \pi_2)$ goes to a finite limit at $\pi_1 = \pi_2$ from f bounded). If multiple maxima exist, define $(\pi_1^*, \pi_2^*, \tau_1^*)$ as the first by lexicographic order.

Result 6 *For any distribution of buyer valuations satifying the above assumptions:*

a) As $\beta \to 0$ (buyers' alternatives are limited), the optimal price-quote-revision time τ_1^* goes to $+\infty$, and π_1^* and π_2^* go to the optimal prices under full two-price discrimination.

b) As $\beta \to \infty$ (alternatives are readily available to buyers), the optimal price-quote-revision time τ_1^* goes to zero, and π_1^* and π_2^* both go to the monopoly price.

Proof. For $\beta \to 0$, and for fixed π_1, π_2, q_1, and q_2, τ_1^* as given by (8) goes to $+\infty$. The second expression for $\mathbf{E}R^*$ in (9) goes to $\pi_1 q_1 + \pi_2 q_2$. This has a unique maximizer by assumption. The monotonicity of $\mathbf{E}R^*$ in β implies convergence of π_1^* and π_2^*.

For $\beta \to \infty$, the first branch of (9), $\frac{\alpha}{\alpha+\beta}\pi_2(q_1 + q_2)$, is maximized with the unique maximizer of r_1. The second branch of (9) becomes arbitrarily close to $\frac{\alpha}{\alpha+\beta}\pi_1 q_1$ and, from the monotonicity in β, its maximizer converges to the unique maximizer of r_1. If there is an ϵ such that $\tau_1^* > \epsilon$ for all β, we must have $\pi_1^* - \pi_2^* \to 0$, otherwise the first branch of (9) becomes larger than the second branch for β small enough. But, from (8), $\pi_1^* - \pi_2^* \to 0$ implies that $\tau_1^* \to 0$. ■

In a non-competitive market, where it's difficult for buyers to find alternatives, the seller can achieve greater price discrimination. The optimal policy is to offer significant individual discounts, but only after a substantial delay. In a competitive market, where buyers have easy-to-find alternatives, the seller is unable to price-discriminate. The optimal policy is to offer at most a small discount after a short delay.

This result may appear somewhat counter-intuitive, as it predicts that in less competitive markets price negotiation is more prevalent and sellers are more willing to offer sizeable individual discounts. It seems, however, consistent with anedoctal evidence that firms with substantial monopoly power engage in protracted negotiations, and regularly offer large discounts. Also consistent with our model, such discounts are commonly offered under confidentiality requirements regarding the negotiation process and terms of deal, which limits the ability of other buyers to engage in strategic behavior. Finally, note that the converse is also true. In markets that are competitive and where information is widely available, transactions are mostly posted-price.

Example

Consider the case of linear demand. Without loss of generality, we assume that the distribution of valuations is uniform in the $[0,1]$ interval, $F_w(w) = w,\ w \in [0,1]$, so that $q_1 = 1 - \pi_1$ and $q_2 = \pi_1 - \pi_2$. Substituting in (9), we obtain an expression for the expected revenue with optimal timing as a function of the opening and revised prices, $\mathbf{E}R^*(\pi_1, \pi_2)$. We solve numerically for the maximizing π_1^* and π_2^*, and compute the corresponding τ_1^*. As before, we normalize α to one. The results, summarized in Table 2, are consistent with Result 6.

β	π_1^*	π_2^*	τ_1^*	$\Delta\mathbf{E}R^*$
$\to 0$	$\to 0.67$	$\to 0.33$	$\to +\infty$	$\to 33.33\%$
0.2	0.60	0.33	1.98	15.98%
0.5	0.58	0.35	1.30	8.76%
1.0	0.55	0.37	0.88	4.48%
2.0	0.53	0.40	0.55	1.88%
5.0	0.52	0.44	0.27	0.45%
$\to +\infty$	$\to 0.50$	$\to 0.50$	$\to 0.00$	$\to 0.00\%$

Table 2: Optimal opening price quote π_1^*, revised price quote π_2^*, and time of revision τ_1^*, as a function of the availability of alternatives to the buyer. Larger β corresponds to alternatives being more readily available. The gain in expected revenue $\Delta\mathbf{E}R^*$ is given relative to the expected revenue with the optimal constant price $\pi^* = 0.5$.

4 Estimating the model parameters from observations

In order to be able to estimate the model parameters from historical observations, the seller has to systematically record information about all customer contacts. The minimum information needed for model estimation to be feasible is: the time of the first contact by each buyer requesting a price-quote; the time, if any, when the price quote was revised; and the selling price if a transaction took place. In addition, the ability to obtain accurate estimates is greatly improved if the seller also keeps a record of the time when each buyer agreed to purchase the good. For the most part, we will assume that the time-of-sale is known. Likewise, the ability to estimate the model parameters is improved if the seller has access to the time when buyers who did not purchase the good either found an alternative or made a decision not to purchase. However, this information is less likely to be available to the seller and, to the extent that it is available, might not be reliable. We will assume that it is not available.

While we do not discuss the issue further, note that there are a number of potential sources of bias in the customer-contact and transaction data. In particular, it is important to have the customer contact history for all buyers, including those who did not make a purchase. Other sources of bias may be less obvious. If a salesperson can reliably detect impatient buyers in a way which is not measured by the data and responds accordingly by offering a discount earliear, the quote-revision times may be correlated with different values of α and β.

4.1 Dispersion of quote-revision times and identifiability

As a preface to the discussion of estimation procedures for the two-price model with exponential random times, we note that the model parameters α, β, q_1 and q_2 cannot be fully identified without variance in the price-quote revision times. This is the case even if information about the timing of each sale is available. The parameter β that characterizes the availability of alternatives is, in a sense, a time-elasticity of demand. It cannot be adequately estimated from observations without historical dispersion in the timing decisions. This is analogous to the situation faced when estimating a demand curve, where the price-elasticity of demand cannot be adequately estimated from observations without historical dispersion in the price decisions.

Suppose that the quote-revision time τ_1 is held constant. Given enough observations, including how much time elapsed since each buyer's first contact until a sale was made, we can estimate F_{π_1} and F_{π_2} in (5) and (6). For F_{π_1}, this results in estimates of the quantities $\alpha + \beta$ (the initial rate of F_{π_1}) and $\frac{\alpha}{\alpha+\beta} q_1$ (the limit of F_{π_1}). For F_{π_2} this results in estimates of the quantities $\alpha + \beta$ (the initial rate of F_{π_2}) and $\frac{\alpha}{\alpha+\beta} \left(q_1 e^{-(\alpha+\beta)\tau_1} + q_2 e^{-\beta\tau_1} \right)$ (the limit of F_{π_2}). Since the two initial rates are the identical, we only have three estimates for four parameters and, without variation in the quote-revision times τ_1, the model is not fully identifiable.

As a side-note, if information about timing of the sales in not available (or not reliable), the unidentifiability problem is more severe when there is no variation in the timing of the quote revision. However, the model is still fully identifiable if there is dispersion in the quote-revision times (and if information about the timing of the quote revision for each sale and for each lost sale is recorded).

Consider an example, with $\pi_1 = 600$, $\pi_2 = 100$, $\alpha = \beta = 1$, $q_1 = 0.05$, and $q_2 = 0.25$, for which $\tau_1^* = 0.693$. We compute the manifold of models that result in the same distribution of observations. Data generated from any of the models is equally consistent with all models in this set. With $\tau_1 = 0.693$, all models in the set have $\mathbf{E}R = 18.1$. The optimal quote-revision time for each model and the corresponding optimal expected revenue are plotted in Figure 2.

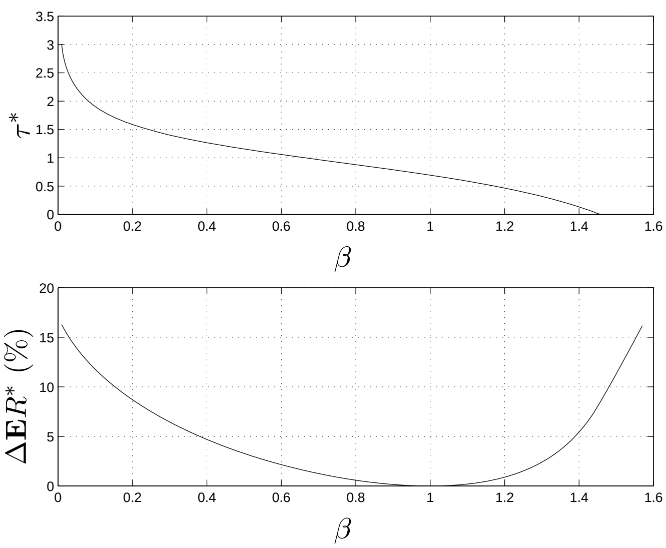

Figure 2: Optimal quote-revision time and corresponding optimal expected revenue for a set of models for which the distribution of observations is identical (parametrized over the value of β for each model).

The difference in expected revenue in this example can be in the order of 10% (when discussing estimation procedures, we will return to this example and see a case where dispersion in the quote-revision times allows for an expected gain of 6%). Since, for any of these models, the observations are consistent with beliefs about models parameters for which the policy is optimal, the seller might then persist in a sub-optimal policy. Another way of describing the model ambiguity given the observations is that the seller cannot determine what fraction of buyers who don't make a purchase find the good too expensive, and what fraction found an alternative supplier for the good. It may be possible to resolve the ambiguity by surveying customers who did not buy the product, and inquire whether they found an alternative supplier for the good or chose not to make a purchase. If this can be reliably done, the seller can fully identify all parameters without timing dispersion. Also, if the distribution of valuations $F_w(\cdot)$ is known (so that q_1 and q_2 are known), we can identify α and β. In fact, it is sufficient to know the elasticity $\frac{(q_1+q_2)\pi_2}{q_2\pi_1}$, which adds a fourth constraint and allows the parameters to be uniquely identified.

4.2 Maximum likelihood estimates

We now assume that elasticity information is not available, and that there is historical dispersion in quote-revision times. We describe the maximum-*a-posteriori* (MAP) estimates given a set of observations, with and without knowledge of the time of sale.

Let n be the number of buyers who requested a price quote. The observations are specified by the n-vectors τ and s, where the k-th entries are denoted by $\tau^k \in [0, +\infty]$, the time at which the price quote was revised from π_1 down to π_2 to the k-th buyer, and $s^k \in [0, +\infty]$ is the time of sale to the k-th buyer. We define $s^k = +\infty$ when no sale took place. Both times, τ^k and s^k, are measured from the k-th buyer's initial price-quote request. The problem parameters are collected in the vector $\theta = (\alpha, \beta, q_1, q_2)^T$.

If time-of-sale information is available, the posterior likelihood is

$$\mathbf{P}(\theta|s,\tau) = \frac{\mathbf{P}(\tau)}{\mathbf{P}(s,\tau)}\mathbf{P}(s|\theta,\tau)\mathbf{P}(\theta). \tag{19}$$

Since $\mathbf{P}(\tau)$ and $\mathbf{P}(s,\tau)$ are constant for a given set of observations, $\mathbf{P}(\theta|s,\tau) \propto \mathbf{P}(s|\theta,\tau)\mathbf{P}(\theta)$, and we can then obtain maximum-likelihood estimates by maximizing $\mathbf{P}(s|\theta,\tau)\mathbf{P}(\theta)$. In Section 4.3 we use a Markov-chain Monte Carlo procedure to draw samples from $\mathbf{P}(s|\theta,\tau)\mathbf{P}(\theta)$ which, scaled to integrate to one, provide an estimate for the distribution of $\mathbf{P}(\theta|s,\tau)$.

With the buyers assumed independent, the density of $\mathbf{P}(s|\theta,\tau)$ is constructed from the multiplication of

$$\begin{aligned}
\mathbf{P}\big(s^k, s^k \in [0,\tau^k[\;\big|\; \theta, \tau^k\big) &= q_1\alpha e^{-(\alpha+\beta)s^k} && (20)\\
\mathbf{P}\big(s^k, s^k \in [\tau^k,+\infty[\;\big|\; \theta, \tau^k\big) &= \Big(q_1 e^{-(\alpha+\beta)\tau^k} + q_2 e^{-\beta\tau^k}\Big)\,\alpha e^{-(\alpha+\beta)(s^k-\tau^k)} && (21)\\
\mathbf{P}\big(s^k = +\infty \;\big|\; \theta, \tau^k\big) &= 1 - q_1\frac{\alpha}{\alpha+\beta} - q_2\frac{\alpha}{\alpha+\beta}e^{-\beta\tau^k}. && (22)
\end{aligned}$$

For the case where time-of-sale information is not available, the data vector s is replaced by the vector π, where π^k is price of sale to the k-th customer, with, say, $\pi^k = 0$ if no sale took place. The posterior distribution is now

$$\mathbf{P}(\theta|\pi,\tau) = \frac{\mathbf{P}(\tau)}{\mathbf{P}(\pi,\tau)}\mathbf{P}(\pi|\theta,\tau)\mathbf{P}(\theta).$$

The likelihood for each buyer is obtained by integrating (20) and (21) over the corresponding range for s^k,

$$\mathbf{P}\big(\pi^k = \pi_1 \;\big|\; \theta, \tau^k\big) = q_1\frac{\alpha}{\alpha+\beta}\Big(1 - e^{-(\alpha+\beta)\tau^k}\Big)$$

$$\mathbf{P}\big(\pi^k = \pi_2 \mid \theta, \tau^k\big) = \frac{\alpha}{\alpha+\beta}\Big(q_1 e^{-(\alpha+\beta)\tau^k} + q_2 e^{-\beta\tau^k}\Big)$$
$$\mathbf{P}\big(\pi^k = 0 \mid \theta, \tau^k\big) = 1 - q_1\frac{\alpha}{\alpha+\beta} - q_2\frac{\alpha}{\alpha+\beta}e^{-\beta\tau^k}.$$

Without time-of-sale information, more data will be required to achieve the same estimation accuracy. Also, the issues with unidentifiability if there is insuficient dispersion in quote-revision times will be more severe.

The estimates can benefit from the use of a prior distribution on the model parameters, $\mathbf{P}(\theta)$, to incorporate the seller's knowledge of the market and buyers. This prior may be based on previous experience with the same or similar goods, or on market and consumer research. The simulation in our numerical example below is performed with a uniform 'uninformative' prior (which is improper for α and β). Note that, without time-of-sale information and with an improper uniform prior, the posterior distribution for α and β can be shown to be improper, which makes the numerical computation of estimates problematic. With time-of-sale information, however, the posterior is always proper. Possible priors for α and β are, for instance, $\mathbf{P}(\alpha) = 1/\widehat{\alpha}\, e^{-\alpha/\widehat{\alpha}}$ and $\mathbf{P}(\beta) = 1/\widehat{\beta}\, e^{-\beta/\widehat{\beta}}$, with $1/\widehat{\alpha}$ set to the seller's best estimate of the average time until a purchase is made by a buyer who has no alternative suppliers and whose valuation is above the quoted price. Likewise, $1/\widehat{\beta}$ is the seller's best estimate of the average time until a buyer whose valuation is below the quoted price finds an alternative. A possible prior for the valuation probabilities is to have q_1 and q_2 distributed according to $(q_1, q_2, q_3) \sim \text{Dirichlet}(c\,\widehat{q}_1, c\,\widehat{q}_2, c\,(1-\widehat{q}_1-\widehat{q}_2))$, where $\widehat{q}_1$ and $\widehat{q}_1+\widehat{q}_2$ are the seller's best estimates for the values of the demand curve at π_1 and π_2. The value c is the strength of the prior, which can be informally interpreted as the number of observations that would be needed to significantly revise the existing best estimate.

An example of a likelihood function with time-of-sale information is given in Figure 3 (details of the parameters in this example and of how the data was generated are given in Section 4.4). Depicted are equally-spaced level sets of the log-likelihood. The function is plotted over a plane of two parameters, with the remaining two parameters set at the 'true' value used to generate the simulated historical data. While the estimation problem is well-conditioned in the q_1 and q_2 plane, the same is not the case in the q_2 and β plane. That is, there are many values for the model parameters that are nearly-equally likely given the observations, which is related to the identifiability issues discussed in the previous section.

A consequent question is how to use the model-parameter estimates to select the quote-revision time. A simple approach is to substitute the estimated parameter values for the true parameters in equation (8). We denote the resulting quote-revision time by τ_1^{MAP}. As will

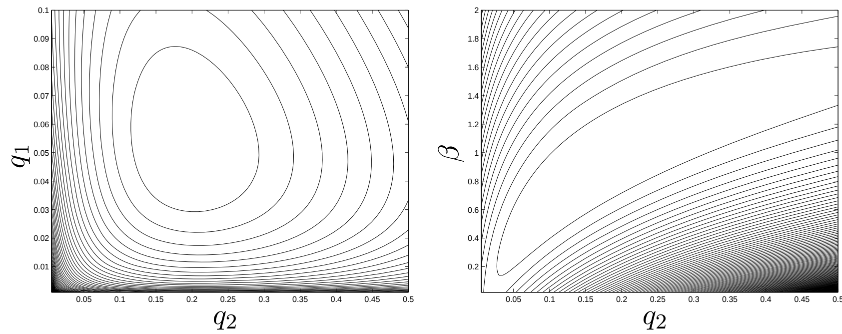

Figure 3: Equally-spaced level sets for an example of the log-likelihood function, plotted over a plane of two parameters (with the remaining parameters at the 'true' value used to generate the simulated observations). The estimation problem is well-conditioned in the q_1 and q_2 plane, but not in the q_2 and β plane.

be illustrated in the numerical example of Section 4.4, this 'certainty equivalent' policy can be considerably suboptimal. An approximation of the optimal policy accounting for the uncertainty in the model-parameter estimates can be obtained from the MAP estimates as follows. Denote by H the Hessian of $\mathbf{E}R$, as defined in (7), with respect to θ, which can be computed in closed form. Denote by Σ_θ the second moment of the log-likelihood of θ, computed at the MAP estimate $\widehat{\theta}$. The quote-revision time is then selected to maximize

$$\mathbf{E}_\theta \mathbf{E}_{s|\theta} R \;\approx\; \mathbf{E}R\big(\,\widehat{\theta}\,\big) + \frac{1}{2}\,\mathrm{Tr}\big(\,\Sigma_\theta H\,\big).$$

However, especially given that the problem is small in the sense of having few parameters, a better approach is to estimate the posterior distribution by simulation in order to directly compute this expectation by numerical integration.

4.3 Full estimate of the posterior density

Markov-chain Monte Carlo simulation allows for the approximate computation of the posterior distribution of the model parameters conditional on the available observations. For a problem with only four parameters, effective results are obtained with a simple Metropolis-Hastings procedure with symmetrical jump distribution, which we now describe.

The MAP estimates can be used as a starting point for the chain. However, the transient from any starting point was observed to be short. Since the log-likelihood function has an ill-conditioned Hessian, the efficiency and reliability of the procedure is enhanced by using an adaptive jump distribution. The adaptation of the jump distribution is as follows. Over the first

5000 iterations, the sample covariance matrix of θ over the previous 1000 iterations is computed, and used as the covariance of the jump distribution. The covariance of the jump distribution is kept constant over the final 5000 iterations, and the first 5000 iterations are discarded. We subsample the second half of the chain by a factor of 10. (An alternative to adaptation based on the sample covariance is to use a scaled version of the inverse of Hessian of the log-likelihood function, which can be computed in closed form, to define the covariance matrix of the jump distribution, which will then no longer be symmetrical.)

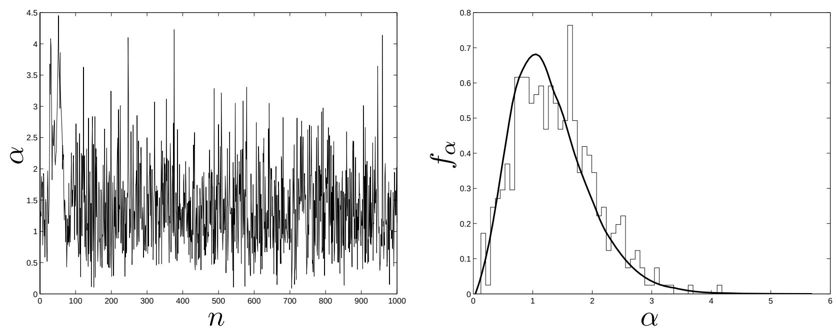

Figure 4: A sample chain for α and the estimated posterior distribution of α, based on 10,000 and on one million trials.

Sample output of the procedure is given in Figure 4 (details of the example used are given in Section 4.4). The transitory regime is over after about 1000 iterations (subsampled to 100 in the figure), and the chain is well-mixed. Also plotted is the estimated posterior distribution of α, based on ten thousand and on one million trials (both also downsampled by a factor of ten).

For each quote-revision time τ_1, the simulated posterior distribution is used to compute a numerical approximation of the expectation over the model uncertainty. The optimal quote-revision time accounting for model uncertainty is selected as

$$\tau_1^{\text{MCMC}} = \operatorname*{argmax}_{\tau>0} \mathbf{E}_\theta \mathbf{E}_{s|\theta} R(\tau).$$

The one-dimensional maximization can be conducted by simple gridding.

4.4 Example

Consider the following two sets of model parameters.

(i) $\alpha = 1, \beta = 1, \pi_1 = 600, \pi_2 = 100, q_1 = 0.05, q_2 = 0.25, N = 1000.$

These parameters are chosen as a case where the expected revenue is sensitive to the relative values of q_2 and $1/\beta$. To reliably estimate these parameters, a significant amount of data with dispersion in the quote-revision times is needed. However, we will see that even without the benefit the parameter estimation accuracy provided by timing dispersion, the two-price policy significantly outperforms a static price.

(ii) $\alpha = 1, \beta = 0.1, \pi_1 = 1000, \pi_2 = 100, q_1 = 0.05, q_2 = 0.45, N = 20$.

These parameters are chosen as a case where the expected revenue is not sensitive to the relative values of q_2 and $1/\beta$. The two-price policy does well even with few historical observations (as the other three dimensions of the parameter space can be estimated with few data). Also, given the small amount of data, the robustness gained by taking into account the uncertainty about the problem parameters results in a significant improvement in the expected revenue.

For each of these two sets of parameters, we consider two cases, with and without dispersion in the historical quote-revision times.

(a) $\tau_1 = \frac{1}{\beta}$ (indicated in the results of Table 3 by $D = 0$).

(b) τ_1 uniformly distributed in $[0, \frac{2}{\beta}]$ (indicated by $D = 1$).

For each of the four resulting distributions of observations, we perform a simulation as follows.

1. Generate 1000 different sample histories of sales according to the 'true' model parameters, each with N buyers.
2. For each sample sales history, compute maximum-posterior-likelihood point estimates for the model parameters, as described in Section 4.2. Select the quote-revision time which is optimal for the estimated parameters, τ_1^{MAP}. For this time, compute the expected revenue according to the true model parameters, $\mathbf{E}R^{\text{MAP}}$.
3. For each sample sales history, run the MCMC procedure as described in Section 4.3. Select the quote-revision time that maximizes the expected revenue over the posterior distribution of the model parameters, τ_1^{MCMC}. Compute the expected revenue according to the true model parameters, $\mathbf{E}R^{\text{MCMC}}$.

An example of the log-likelihood for a particular sample sales history generated according to the distribution of observations of case ($i.a$) is given in Figure 3. The results of the Markov-chain Monte Carlo simulation are illustrated in Figure 4, also for case ($i.a$) and for same sample sales history.

Table 3 summarizes the results. The revenue with fixed price and the expectations of $\mathbf{E}R^{\text{MAP}}$ and $\mathbf{E}R^{\text{MCMC}}$ are compared with the expected revenue with optimal timing based on perfect information about the model parameters. In case (i), the seller can benefit from historical

	$(i.a)$ $\alpha = 1,\ \beta = 1,$ $\pi_1 = 600,\ \pi_2 = 100,$ $q_1 = 0.05,\ q_2 = 0.25,$ $N = 1000,\ D = 0$	$(i.b)$ $\alpha = 1,\ \beta = 1,$ $\pi_1 = 600,\ \pi_2 = 100,$ $q_1 = 0.05,\ q_2 = 0.25,$ $N = 1000,\ D = 1$	$(ii.a)$ $\alpha = 1,\ \beta = 0.1,$ $\pi_1 = 1000,\ \pi_2 = 100,$ $q_1 = 0.05,\ q_2 = 0.45,$ $N = 20,\ D = 0$	$(ii.b)$ $\alpha = 1,\ \beta = 0.1,$ $\pi_1 = 1000,\ \pi_2 = 100,$ $q_1 = 0.05,\ q_2 = 0.45,$ $N = 20,\ D = 1$
Fixed price	83%	83%	61%	61%
Maximum likelihood	90%	96%	82%	82%
Posterior distribution	93%	96%	96%	96%

Table 3: Expected revenue as % of expected revenue with optimal timing based on perfect information about the model parameters.

dispersion in quote-revision times, as well as from the robustness provided by using MCMC to estimate the posterior parameter distribution.

In case (*ii*), given the short sales history and how little data is available, one might expect that not much improvement is possible over the expected revenue with fixed price. This is not the case, and both τ_1^{MAP} and τ_1^{MCMC} lead to substantial gains over either $\tau_1 = 0$ or $\tau_1 = +\infty$. Also, due to the robustness provided by accounting for the uncertainty in the parameter estimates, τ_1^{MCMC} substantially outperforms τ_1^{MAP}. There is no benefit from dispersion in quote-revision times since, even when dispersion is present, there isn't enough data to learn about the relative values of β and q_2.

5 Discussion

Revenue management models have, for the most part, relied in one form or another on price-elasticities of demand. When transacted prices depend on the outcome of the interaction with each individual buyer, the price-elasticity cannot be directly estimated. It can, however, be estimated with a model that controls for time-sensitivity, which is made feasible with current sales-tracking software that records the complete history of customer contacts.

A number of extensions and related issues may be of interest to investigate, the more immediate being to validate the model against empirical data, as well as determine the most adequate measure of (chronological, number of contacts, *etc.*) Another issue we did not address that may be of practical importance is the effect of imperfect information about lost sales. This will arise if, say, records are not kept for an unknown percentage of buyers who do not make a purchase. Estimation procedures can also be developed for the case when there is simply no recorded information about lost sales.

In practice, the seller may use any other available information to segment the buyers. Sellers often request information from each buyer about their expected time of purchase. This information may be used to manage the negotiation process and avoid offering discounts earlier than appropriate. While buyers may be strategic in their answers, the seller can use historical data to estimate model parameters for each segment, based on these or other data, without any presumption of truthfulness. The segmentation can also be based on, for instance, how often the buyer contacts the seller. Note that the sign of the appropriate response cannot be predetermined: a higher frequency of contacts by buyers (*e.g.*, with counter-offers) may signal a more active search behavior (higher β), or may signal that fewer alternatives are available to the buyer (lower β).

Sellers have a strong incentive to keep their strategy private, which is reflected in non-disclosure agreements that companies sometimes require as a pre-condition to negotiate a discount. If information about the seller's strategy becomes public, buyers will hold out longer in hopes of obtaining a discount (lower α). It is also common for sellers to make the request for a price quote time-consuming, or in some other way costly. In a less competitive market, where a policy of individualized pricing is more attractive to the seller, this is not likely to lead to a significant loss of buyers. It allows the seller to remove the incentive for buyers to request quotes when they have no need for the good, or earlier than when the need for the good arises.

Several alternative models can be explored, such as parameters α and β that are dependent on the buyer's valuation for the good. The problem with continuous price revisions can be formulated as an optimal control or calculus of variations problem, or as the limit of the multiple-price case when the number of quote revisions becomes arbitrarily large. While impractical in practice, this problem may be useful as a benchmark. In the limit where $\beta \to 0$, with an infinite number of price-revision times, we obtain perfect (or first-degree) price discrimination, so that the expected revenue should go to the expected valuation, $\mathbf{E}_w w$.

Acknowledgments

The initial motivation for this work was a talk by Andy Boyd of PROS Pricing at the INFORMS Revenue Management conference, making the case for the need for the Revenue Management field to address individually-negotiated prices ("not just modelling pricing, but the actual sales process" (Boyd 2007b)). Giuseppe Lopomo provided helpful early direction for the review of the economic literature.

References

Abreu, D. and F. Gul (2000, January). Bargaining and reputation. *Econometrica 68*, 85–117.

Admati, A. R. and M. Perry (1987, July). Strategic delay in bargaining. *The Review of Economic Studies 54*, 345–364.

Armstrong, M. (2006). *Recent Developments in the Economics of Price Discrimination*, Volume II, Chapter 4. Cambridge University Press.

Banks, J. and S. Moorthy (1999, April). A model of price promotions with consumer search. *International Journal of Industrial Organisation* (17), 371–398.

Besanko, D. and W. Winston (1990, May). Optimal price skimming by a monopolist facing rational consumers. *ms 36*(5), 555–567.

Bhandari, A. and N. Secomandi (2007, June). Dynamic revenue management with negotiated prices.

Boyd, A. (2007a). *The Future of Pricing: How Airline Ticket Pricing Has Inspired a Revolution*. Palgrave Macmillan.

Boyd, A. (2007b). The science of selling. *Journal of Revenue and Pricing Management*. Forthcoming.

Cho, I. K. (1990). Uncertainty and delay in bargaining. *Review of Economic Studies 57*, 575–596.

Coase, R. (1972, April). Durability and monopoly. *Journal of Laws and Economics 15*(1), 143–149.

Deneckere, R. and M. Y. Liang (2006, September). Bargaining with interdependent values. *Econometrica 74*(5), 1309–1364.

Fuchs, W. and A. Skrzypacz (2006, September). Bargaining with arrival of new traders or new information.

Gul, F., H. Sonnenschein, and R. Wilson (1986, June). Foundation of dynamic monopoly and the coase conjecture. *Journal of Economic Theory 39*, 155–190.

Hancock, D. and D. Humphrey (1997, December). Payment transactions, instruments, and systems: A survey. *Journal of Banking and Finance 21*(11-12), 1573–1624.

Hausken, K. (1997, December). Game-theoretic and behavioural negotiation theory. *Group Decision and Negotiation 6*(6), 511–528.

Netessine, S. (2004, October). Dynamic pricing of inventory/capacity with infrequent price changes. *European Journal of Operational Research 174*(1), 553–580.

Ovchinnikov, A. and J. M. Milner (2005, September). Strategic response to wait-or-buy : Revenue management through last minute deals in the presence of customer learning.

Phlips, L. (1988, June). Price discrimination : A survey of the theory. *Journal of Economic Surveys 2*, 135–167.

Popescu, I. and Y. Wu (2007). Dynamic pricing with reference effects. *Operations Reaserch*. forthcoming.

Salop, S. (1977, October). The noisy monopolist : Imperfect information, price dispersion and price discrimination. *The Review of Economic Studies 44*(3), 393–406.

Sobel, J. and I. Takahashi (1983, July). A multistage model of bargaining. *The Review of Economic Studies 50*(3), 411–426.

Stahl, D. (1989, September). Oligopolistic pricing with sequential consumer search. *The American Economic Review 79*(4), 700–712.

Stokey, N. (1979, August). Intertemporal price discrimination. *The Quarterly Journal of Economics 93*(3), 355–371.

Vulcano, G., G. van Ryzin, and C. Maglaras (2002, November). Optimal dynamic auctions for revenue management. *Management Science 48*(11), 1388–1407.

Yildiz, M. (2004, February). Waiting to persuade. *Quarterly Journal of Economics 119*, 223–248.