OPERATIONS RESEARCH
Articles in Advance, pp. 1–25
ISSN 0030-364X (print) | ISSN 1526-5463 (online)
https://doi.org/10.1287/opre.2016.1548

Intertemporal Pricing Under Minimax Regret

René Caldentey
Booth School of Business, University of Chicago, Chicago, Illinois 60637, rene.caldentey@chicagobooth.edu

Ying Liu, Ilan Lobel
Stern School of Business, New York University, New York, New York 10012 {yliu2@stern.nyu.edu, ilobel@stern.nyu.edu}

We consider the pricing problem faced by a monopolist who sells a product to a population of consumers over a finite time horizon. Customers' types differ along two dimensions: (i) their willingness-to-pay for the product and (ii) their arrival time during the selling season. We assume that the seller knows only the support of the customers' valuations and do not make any other distributional assumptions about customers' willingness-to-pay or arrival times. We consider a robust formulation of the seller's pricing problem that is based on the minimization of her worst-case regret. We consider two distinct cases of customers' purchasing behavior: myopic and strategic customers. For both of these cases, we characterize optimal price paths. For myopic customers, the regret is determined by the price at a critical time. Depending on the problem parameters, this critical time will be either the end of the selling season or it will be a time that equalizes the worst-case regret generated by undercharging customers and the worst-case regret generated by customers waiting for the price to fall. The optimal pricing strategy is not unique except at the critical time. For strategic consumers, we develop a robust mechanism design approach to compute an optimal policy. Depending on the problem parameters, the optimal policy might lead some consumers to wait until the end of the selling season and might price others out of the market. Under strategic customers, the optimal price equalizes the regrets generated by different customer types that arrive at the beginning of the selling season. We show that a seller that does not know if the customers are myopic should price as if they are strategic. We also show there is no benefit under myopic consumers to having a selling season longer than a certain uniform bound, but that the same is not true with strategic consumers.

Keywords: demand uncertainty; strategic consumers; robust optimization; prior-free; worst-case regret.
Subject classifications: marketing: pricing; games/group decisions: noncooperative; programming.
Area of review: Operations and Supply Chains.
History: Received November 2013; revisions received April 2015, April 2016; accepted June 2016. Published online in *Articles in Advance* November 7, 2016.

1. Introduction

Over the last couple of decades, dynamic pricing has been transformed from a curious and somewhat controversial practice used primarily by upstart airlines into a technique that is widely used in a variety of industries. As technology has evolved and reduced menu costs, retailers of all sorts have adopted intertemporal pricing practices. One of the key economic drivers behind the rapid dissemination of dynamic pricing is demand uncertainty: there is enormous value for a firm in being able to change prices over time in situations where the firm does not know how much customers are willing to pay for its products.

In response to the increasing use of dynamic pricing in practice, academics have proposed a variety of techniques for algorithmically determining intertemporal pricing policies. However, the vast majority of these approaches require the firm to know the probability distributions of customer valuations and arrival times. Assuming that the firm knows a full probabilistic model of customer valuations and arrival times is problematic for at least two reasons. The first reason is the obvious one: firms do not have access to such probability distributions; even when they have sales data, they typically do not have access to a data set that is rich enough to estimate valuation and arrival time distributions. The second reason is less obvious but equally important: taking the probability distributions of customer valuations and arrival times as given assumes away a significant part of the lack of knowledge about customer valuations and arrival times.

These issues are especially acute for firms introducing new products into the marketplace. When firms launch new products, they usually have very little information about how much customers are willing to pay for them. This great degree of uncertainty makes new products excellent candidates for dynamic pricing strategies. However, the same uncertainty about customer valuations also hobbles our ability to use established dynamic pricing techniques since they rely on the firm knowing the probability distribution of customer valuations. Pricing of new products is no trivial matter. For instance, a McKinsey study reported that more than a 100,000 new products are introduced yearly into the U.S. retail industry, but 70%–80% of these launches fail. The use of dynamic pricing for new products can reduce the impact of demand uncertainty on the firm and, in this way, help some of these launches succeed.

We approach the problem of intertemporal pricing from the perspective of *robust optimization* (see Bertsimas and Sim 2004, Ben-Tal et al. 2009). Specifically, we assume that the seller only knows the range of customers' valuation (or willingness to pay) for its product and makes no additional assumptions about their distribution or about the customers' arrival process. The formulation we propose is quite parsimonious, but still sufficiently rich to give rise to different types of pricing policies for different sets of problem parameters. Our model formulation can be used both in settings with limited information, where the firm has only a very rough guess of the customer's range of valuations, as well as in an environment that is more data-rich, where the firm can use sales data to estimate an uncertainty set of the customers' valuations. It is also flexible enough to allow us to study optimal pricing policies for myopic as well as strategic customers, the latter being customers who time their purchases to maximize their own discounted utilities.

Under such a minimalist informational structure, the standard expected profit maximization criterion is not appropriate as the seller's objective function (for more details, see Section 2). Instead, we consider the seller's *regret*, which is defined as the difference between her payoff under full demand information and her realized payoff. In this setting, an optimal pricing strategy is one that minimizes the difference between the seller's ex-post payoff and that of a clairvoyant who sets prices knowing customer types (valuation and arrival time) in advance. In particular, we assume the seller chooses a policy that minimizes her worst-case anticipated ex post regret. The seller assumes that nature selects customer types from the uncertainty set to generate as much regret as possible, and that customers behave either myopically or strategically with respect to prices. The first paper to propose a minimax regret criterion for pricing without a prior distribution over customer valuations was Bergemann and Schlag (2008). Our approach can be seen as an intertemporal version of Bergemann and Schlag (2008)'s static model.

We make several modeling assumptions that we wish to highlight before we discuss our contributions. First, we do not include inventory considerations in our model. Having the customers strategically consider availability risk would require us to model customers' beliefs about each other's strategies, and thus make our model less parsimonious. To reduce the number of parameters in our model and maintain tractability, we also assume the seller and the consumers discount the future at the same rate, and that all consumers can be described by the same uncertainty set. We do allow for a mix of strategic and myopic consumers to be present in the market (see Section 6). Furthermore, we assume the firm has full commitment power. Minimax regret without commitment leads to a problem of dynamic inconsistency, which is not amenable to a Bellman-equation approach (see Hayashi 2009), and is thus likely an intractable problem. For a more detailed discussion on our commitment assumption, see Section 2.

1.1. Contributions

The primary contribution of this paper is the development of a robust optimization methodology to compute intertemporal pricing policies that minimize a firm's worst-case regret when selling to myopic or strategic consumers with uncertain valuations and arrival times.

In Section 4, we consider the case in which the firm sells to myopic customers. The regret from a given consumer can be decomposed into two terms: a valuation and a delay regret. Valuation regret captures losses due to undercharging consumers and can be lowered by raising prices. Delay regret captures losses due to consumers waiting for lower prices and can be reduced by lowering prices overall. For any given regret level R, there exists a price path that maintains the valuation regret for customers with high value at a constant R over the selling horizon and there exists another price path that maintains the delay regret for customers that arrive at the beginning of the season at the same constant R. We show that if the time horizon is sufficiently long and the market uncertainty is sufficiently high, the minimax regret is determined by ensuring these two price paths intersect tangentially for a given regret level R. The unique time where these two price paths intersect constitutes what we call the critical time. The optimal price offered at the critical time is uniquely determined, but generically there are multiple optimal price paths. The price paths that maintain constant valuation regret and delay regret determine the boundaries of the set of optimal price paths. Any continuous decreasing price path within these boundaries is optimal, as long as the final price is below a certain value. A typical optimal maximal price path includes an initial *full-markup* period where prices are set equal to the upper limit of customers' valuation range ($\bar{v}$) followed by a *markdown* period. In contrast, the optimal minimal price path has no markup period and has less significant markdowns. In other words, the seller has some flexibility to set prices either aggressively (maximal solution) or conservatively (minimal solution) during the early and late stages of the selling season, but not at an intermediate critical time. When the selling horizon is short, however, the critical time becomes the end of the selling horizon. In this case, the maximal and minimal optimal price paths never intersect. We also show that a selling horizon of length $\ln(3)/r$, where r is the discount factor, is always sufficient to minimize the maximum regret. There is no value for the seller in having a selling season longer than this $\ln(3)/r$.

In Section 5, we consider the case in which the market consists of forward-looking consumers. Under our robust formulation, this problem can be viewed as a three stage game, with the firm acting first and choosing prices, nature responding and selecting customers' valuations and arrival times, and customers acting last and deciding whether and when to buy the firm's product. We develop a methodology based on robust mechanism design to determine an optimal price path for the case of strategic consumers. We show there exist optimal price paths that are decreasing, and that they can take one of three forms. When the market uncertainty is high, prices will be strictly decreasing and customers will be separated into three groups: the ones that buy before the end of the horizon, the mass that will wait until the end of the horizon, and the ones that will be priced out of the market. When market uncertainty is moderate, prices eventually reach the lowest valuation, and the last group ceases to exist. When market uncertainty is low, all consumers will purchase before the end of the horizon. We also show that consumers act "myopically" not with respect to prices, but with respect to modified price path that we call threshold valuations. We further show that, unlike in the myopic consumers case, there does not exist a uniform bound on the maximum useful length of the selling season.

In Section 6, we first compare optimal price paths for myopic and strategic consumers. We show that policies that are tailored for strategic consumers are flatter than policies designed for myopic ones. With strategic consumers, the firm starts from a lower price point than it would with myopic customers, but ends with a higher price than it would end otherwise. This is a consequence of the firm's reduced ability to do price skimming due to the consumers' strategic behavior. We also show that the firm's regret is always worse under strategic customers than under myopic customers. We show that if the firm is unsure of the mix between myopic and strategic consumers, the policy that minimizes the maximum regret is the one that prices as if all consumers were strategic.

2. Related Literature

Starting from the seminal paper by Gallego and van Ryzin (1994), the revenue management community has focused its attention on the problem of how to use dynamic pricing for handling uncertain customer valuations and arrival times over a finite selling season. The early literature is vast and we refer readers to surveys by Bitran and Caldentey (2003), Elmaghraby and Keskinocak (2003), and Talluri and Van Ryzin (2005).

The early models in the dynamic pricing literature all assumed that customers were myopic in how they made their decisions, in that customers would not try to anticipate the firm's future prices when making their decisions. Recently, there has been a major research drive trying to understand the impact of strategic customer behavior on firms using dynamic pricing strategies. Aviv and Pazgal (2008) showed that ignoring forward-looking customers can be costly for the firm and that committing to a fixed price can potentially be more profitable for the firm even in the face of stochastic demand, a counterintuitive result that builds on the insight of the Coase conjecture (see Coase 1972). Su (2007), studying a model where customers are heterogeneous in both their valuations and their degree of patience, shows that markup policies are optimal when high-valuation customers are proportionally more strategic, whereas markdowns are optimal if they are proportionally more myopic. Furthermore, recent papers in the economics and operations management literatures by Hendel and Nevo (2013) and Li et al. (2014) have empirically shown that strategic customer behavior is an important issue that should not be ignored when deciding prices in settings such as retail and airline markets.

Several recent papers have also studied the impact of dynamic pricing when customers are strategic not only about prices, but also about product availability. Liu and van Ryzin (2008) show that understocking can be used by the firm to drive early purchases, at higher prices, when customers are forward-looking. Cachon and Swinney (2009) demonstrate that quick response production is especially valuable in the presence of strategic customers. Yin et al. (2009) recommend sellers to display one item at a time when faced with forward-looking customers to increase the sense of product scarcity in the market. Caldentey and Vulcano (2007) and Osadchiy and Vulcano (2010) propose alternative models for selling to strategic customers, such as running an auction in parallel to regular sales channel and selling with binding reservations, respectively.

Determining a good pricing strategy for selling to strategic customers is challenging, so most of the papers above make one or more simplifying assumption on the pricing problem to keep it manageable. Some papers assume there are only two pricing periods and other papers assume there are only two possible customer valuation levels. Some papers that do allow for general valuation models in multiperiod settings, such as Besbes and Lobel (2015), instead assume there is no uncertainty on customer valuations or arrival times. In contrast to most papers in this literature, we simplify the problem by removing inventory considerations, but offer an intertemporal pricing framework that allows for uncertainty on both customer valuations and arrival times. Our framework can be used to generate optimal dynamic pricing policies for both myopic and strategic customers, enabling us to compare and contrast the two.

Our methodology utilizes a robust optimization approach (see Bertsimas and Sim 2004, Ben-Tal et al. 2009) to model customer valuation and arrival time uncertainty. That is, the firm only assumes that customer valuations and arrival times belong to some uncertainty set, without having Bayesian priors associated with these parameters. We consider the problem of finding the policy that minimizes the maximum regret the firm can incur, where regret is defined as the difference between the seller's payoff under full information and her realized payoff. This minimax regret decision rule was originally proposed by Savage (1951) in his interpretation of Wald (1950). Milnor (1954) proved the existence of decision-theoretic axioms that supported the minimax regret decision rule. We refer the reader to Stoye (2011) for a more recent treatment of the axiomatic underpinnings of minimax regret. The minimax regret criterion captures the fact that the firm would like to earn a profit similar to what it would earn if it knew the customer valuations and arrival times. If the minimax regret is low, the firm will have done almost as well as if it knew the customer valuations and their arrival times. For a more detailed discussion on the merits of minimax regret as a decision rule, we direct the reader to Schlag (2006).

The natural alternative to minimax regret for decision-making without priors is maximin utility, i.e., to maximize the minimum possible utility of the firm. In fact, optimizing assuming the worst possible outcome within an uncertainty set is the standard framework in robust optimization. However, assuming the worst case in our problem would mean assuming customers have the lowest valuation possible. This would lead the firm to set its price at a constant equal to the lowest possible valuation, an exceedingly conservative solution.[1] In the special case where customers valuations are drawn from a set that includes the value zero, this approach would lead to the nonsensical answer of pricing the good at zero.

Our approach builds directly on the robust pricing model proposed by Bergemann and Schlag (2008). That paper studies the static pricing problem faced by a firm that knows nothing about its customers' valuation except that they belong to a given interval. Like us, they study the problem under a minimax regret criterion. They suggest using a randomized pricing scheme as a method of reducing regret, and then determine the optimal randomized pricing rule. Our work can be seen as a dynamic extension of the model considered by Bergemann and Schlag (2008). However, we do not allow for randomizations as they do. In our model, the firm reduces regret by offering different prices at different times instead of using different prices with different probabilities. In Bergemann and Schlag (2011), the same authors consider the case where the seller has multiple priors over the set of consumer valuations, à la Gilboa and Schmeidler (1989). Our work is closer to Bergemann and Schlag (2008) since we consider a scenario where the seller has no prior—or, equivalently, the seller chooses a pricing policy that is robust against all possible priors rather than a particular set of priors.

Our paper utilizes the minimax regret criterion in a dynamic setting, an approach that was axiomatized by Epstein and Schneider (2003) in a multiple priors formulation. In particular, we study a notion of regret proposed by Hayashi (2008) called anticipated ex post regret, which captures the anticipated regret the decision-maker expects to have after the uncertainty is realized. A decision-maker minimizing her maximum regret over time might make choices that are dynamically inconsistent (see Hayashi 2011). There are two natural ways for a decision-maker to address this issue. The first is to impose a dynamic consistency requirement on her own decision-making, à la subgame perfection. Unfortunately, anticipated ex post regret does not satisfy a Bellman-type equation like the one imposed by subgame perfection (see Hayashi 2009). For a generic dynamic decision-making problem, computing the policy that minimizes regret while satisfying dynamic consistency cannot be done using dynamic programming and is likely to be an intractable problem. We thus propose a simpler approach: commitment. We assume the seller has commitment power and that she chooses the policy that minimizes her anticipated ex post regret over the entire horizon. Intertemporal pricing with commitment power is a problem that was first studied by Stokey (1979) and that has recently gained popularity in both the economics (see Board 2008, Pavan et al. 2014, Garrett 2014, Deb 2014) and the revenue management literatures (see Aviv and Pazgal 2008, Borgs et al. 2014, Wang 2016, Besbes and Lobel 2015, Liu and Cooper 2015). We would like to emphasize that we do not tackle the problem of how to do intertemporal pricing without commitment in this paper. Therefore, solutions that we produce do not satisfy subgame perfection and it is plausible that a seller without commitment power would try to deviate from the policies proposed here in the middle of the selling horizon. Recent work by Schlag and Zapechelnyuk (2015) on time-consistent dynamic decision-making under minimax regret could potentially offer some avenues for future researchers to tackle the problem of intertemporal pricing without commitment.

Our paper is also related to Eren and Maglaras (2010), who study how to find a dynamic pricing policy with an optimal competitive ratio for selling to myopic consumers, to Lobel and Perakis (2010), which combines ideas from data-driven and robust optimization to generate robust dynamic pricing policies, to Lim and Shanthikumar (2007), who study robust revenue management in a multiple priors context, and to Perakis and Roels (2010), a paper that proposes robust network revenue management policies using both the maximin utility and the minimax regret criteria.

3. The Model

We consider the pricing problem faced by a monopolist selling durable products to a population of consumers over a continuous-time horizon with length T. Customers differ along two dimensions:[2] (1) their willingness-to-pay for the product and (2) their arrival time during the selling season. We assume that the seller knows only the support $[\underline{v}, \bar{v}]$ of customers' valuations. We do not make any distributional assumptions about customers' willingness-to-pay or arrival times. We consider a robust formulation of the seller's pricing problem, based on the minimization of her worst-case regret, which is defined as the difference between her payoff under full demand information and her realized payoff. In computing these payoffs, we assume that the seller has unlimited capacity and that there are no holding costs or salvage value for unsold units.

In setting up the seller's problem, we first formulate this problem for the special case of a single customer. In Section 3.1, we extend our model to the case with multiple customers and show that the seller's optimal pricing strategy is independent of the number of customers in the marketplace.

In the single-customer case, demand can be modeled by a pair (v, τ), where $v \in [\underline{v}, \bar{v}]$ is the customer's willingness-to-pay and τ is his arrival time. Without loss of generality, we assume that $\tau \in [0, T]$, otherwise there would be no demand during the selling season and the seller's regret would be identically zero. On the supply side, the seller's strategy is given by a price function $p \in \mathcal{P}$, where $\mathcal{P}$ is the set of continuous functions from $[0, T]$ to $[\underline{v}, \bar{v}]$. We assume the seller selects and commits to a price schedule p at time $t = 0$.

To compute the seller's payoffs and corresponding regret, we need to specify how the consumer makes his purchasing decision in response to the seller's pricing strategy. To this end, we introduce a function $d(\cdot)$ that maps the state of the market (v, τ, p) to the time $d(v, \tau, p) \in [0, T] \cup \{\infty\}$ when the customer makes the purchase. We use the convention $d(v, \tau, p) = \infty$ if no purchase is made during the selling season.

We consider two contrasting purchasing behaviors: *myopic* and *strategic*.

—Myopic Consumer: Under a myopic purchasing behavior, the consumer will purchase the product as soon as the price is equal to or falls below his valuation without any consideration of future prices. We denote this myopic purchasing time by $d_M(v, \tau, p)$, which is given by

$$d_M(v, \tau, p) := \min_{\tau \leqslant t \leqslant T} \{t \mid v \geqslant p_t\}. \tag{1}$$

If $v < p_t$ for all $t \in [0, T]$, the consumer leaves the market without making any purchase, i.e., $d_M(v, \tau, p) = \infty$.

Though myopic, the consumers are patient and remain on the market until they make purchases or the end of the sales horizon is reached. It would be possible to consider a different version of a myopic consumer that departs immediately if the price is above his valuation on her arrival date. This different, impatient myopic consumer would be more in line with the consumer behavior assumed by Gallego and van Ryzin (1994). Our model of myopic, but patient consumer is closer in spirit to the consumer behavior assumed in Ahn et al. (2007) and Liu and Cooper (2015).

—Strategic Consumer: As opposed to a myopic consumer, a strategic buyer is forward-looking and optimizes the timing of his purchase to maximize his net discounted utility. We let $d_S(v, \tau, p)$ denote the purchasing time of a strategic consumer, which we define as follows:

$$d_S(v, \tau, p) := \min\left\{\arg\max_{\tau \leqslant t \leqslant T} \{e^{-rt}(v - p_t) \mid v \geqslant p_t\}\right\}, \tag{2}$$

where $r > 0$ is the discount factor. The minimum in the equation above captures the fact that, all else being equal, the consumer would like to get the product as soon as possible. As in the myopic consumer case, if $v < p_t$ for all $t \in [0, T]$ then $d_S(v, \tau, p) = \infty$.

As mentioned above, a distinguishing feature of our model with respect to the revenue management literature on intertemporal pricing is our prior-free approach, where we assume the seller knows only the domain $\mathcal{D} := [\underline{v}, \bar{v}] \times [0, T]$ of the consumer's type (v, τ). There are two standard approaches for dealing with the lack of priors in the robust optimization literature: maximin utility and minimax regret. As pointed out by Bergemann and Schlag (2008), using a maximin utility formulation in a pricing model would lead to a trivial and excessively conservative answer: the firm would price its product at $\underline{v}$. We therefore choose the second option, using a *worst-case regret* criterion for decision-making.

For a given a state of the market (v, τ, p) and a specific consumer's buying behavior $d(v, \tau, p)$, the seller's regret is defined by

$$\mathcal{R}(v, \tau, p) := \Pi_{\mathrm{F}}(v, \tau) - \Pi(v, \tau, p), \tag{3}$$

which is the difference between the supremum of her profit with full information $\Pi_{\mathrm{F}}(v, \tau)$ and her realized profit $\Pi(v, \tau, p)$ with limited information. A perfectly informed seller (or clairvoyant) who knows in advance the buyer's type (v, τ) is capable of extracting all the consumer's surplus by charging a price $p_\tau = v$ at the consumers' arrival time τ and then charging prices $p_t \geqslant v$ for all $t > \tau$. It follows that $\Pi_{\mathrm{F}}(v, \tau) := \sup_p \Pi(v, \tau, p) = e^{-r\tau} v$. On the other hand, the seller's payoff with limited information depends on the consumer's purchasing behavior and is equal to $\Pi(v, \tau, p) = e^{-rd(v, \tau, p)} p_{d(v, \tau, p)}$.[3] We assume in our model that the seller's discount rate is the same as a strategic consumer's discount rate, r.

For a given price path $p \in \mathcal{P}$, we define the seller's worst-case regret $\mathcal{R}(p)$ to be equal to

$$\mathcal{R}(p) := \sup_{(v,\tau)\in\mathcal{D}} \mathcal{R}(v, \tau, p) = \sup_{(v,\tau)\in\mathcal{D}} e^{-r\tau} v - e^{-rd(v,\tau,p)} p_{d(v,\tau,p)}.$$

The seller's worst-case regret problem is then defined as follows:

$$\begin{aligned} R^* &:= \inf_{p\in\mathcal{P}} \mathcal{R}(p) = \inf_{p\in\mathcal{P}} \sup_{(v,\tau)\in\mathcal{D}} \mathcal{R}(v, \tau, p) \\ &= \inf_{p\in\mathcal{P}} \sup_{(v,\tau)\in\mathcal{D}} [e^{-r\tau} v - e^{-rd(v,\tau,p)} p_{d(v,\tau,p)}]. \end{aligned} \tag{4}$$

With a myopic consumer, we can interpret the firm's problem as a zero-sum game between the firm, who selects a price schedule p, and nature, who chooses the customer type (v, τ) to maximize the firm's regret. With a strategic consumer, the game has three separate players acting in sequence: the firm, nature, and the consumer.[4]

In the remainder of this paper, we characterize the solution to the optimization problem in (4) and derive structural properties of the corresponding pricing strategy for various cases in terms of consumers' buying behavior (myopic and strategic) and market size (number of customers). Before we move into this analysis, let us highlight an important feature of the formulation in (4)—one that will prove useful in the derivation of some of our results. The seller's regret can be decomposed into the following two components:

$$\mathcal{R}(v,\tau,p)=\underbrace{e^{-rd(v,\tau,p)}(v-p_{d(v,\tau,p)})}_{\text{valuation regret}}+\underbrace{(e^{-r\tau}-e^{-rd(v,\tau,p)})v}_{\text{delay regret}}. \tag{5}$$

The first term is a *valuation regret*, which is generated by the mismatch between the customer's valuation and the actual price he ends up paying. This is the discounted payoff that the seller "leaves on the table" because she does not know the customer's valuation. The second term is a *delay regret* that captures the time-value of delaying a sale from the customer's arrival time τ to his actual purchasing time $d(v,\tau,p)$. By breaking the seller's regret into these two pieces, one can see that nature has incentives to both postpone and advance the sale (see Figure 1 and the discussion that follows for additional details about the trade-off between valuation and delay regrets). With continuous price paths, an arbitrary myopic customer type (v,τ) can only produce one type of regrets (either valuation or delay) but not both.

3.1. Multiple Customers Case

Up to this point we have characterized the seller's optimal pricing problem under the assumption that there is a single customer (myopic or strategic) who is interested in buying the product. In this section, we extend our previous model to the case in which the marketplace is composed of C customers, for an arbitrary $C\in\mathbb{N}$. We do assume that all these C customers are all either myopic or strategic, an assumption that we relax in Section 6.

Demand in this case is defined by the type (valuation and arrival time) of each customer, that is, by the set $\mathcal{D}^C:=\{(v_i,\tau_i)\in\mathcal{D}\mid i=1,\ldots,C\}$. The seller's worst-case regret problem is given by

$$R^{*C}:=\inf_{p\in\mathcal{P}}\sup_{\{v,\tau\}\in\mathcal{D}^C}\left[\Pi_{\mathrm{F}}^C(v,\tau)-\sum_{i=1}^{C}e^{-rd(v_i,\tau_i,p)}p_{d(v_i,\tau_i,p)}\right], \tag{6}$$

where $\Pi_{\mathrm{F}}^C(v,\tau)$ is the optimal payoff that a clairvoyant seller can obtain knowing in advance (i.e., before selecting the price schedule) the demand realization $\{v,\tau\}\in\mathcal{D}^C$. It follows that

$$\Pi_{\mathrm{F}}^C(v,\tau):=\sup_{p\in\mathcal{P}}\sum_{i=1}^{C}e^{-rd(v_i,\tau_i,p)}p_{d(v_i,\tau_i,p)}.$$

Proposition 1. *The regret R^{*C} is linear in C, i.e., $R^{*C}=C\,R^*$, where R^* is the optimal regret with a single customer. In addition, any optimal pricing strategy for the single customer case is also optimal for any $C\in\mathbb{N}$.*

A similar result is discussed in Section 4 of Bergemann and Schlag (2008) for the case of a static model.

According to Proposition 1, the seller's optimal pricing strategy is independent of the number of customers in the market as long as the firm is aware that customers are all myopic or strategic. As a result, and without loss of generality, in what follows we restrict our analysis to the case of a single customer.

We conclude our model description with two additional remarks.

Remark 1 (Open-Loop vs. Closed-Loop). Although prices in our model are selected at the beginning of the selling season (in an open-loop fashion), the solution is indeed the optimal closed-loop solution under commitment, capturing the evolution of the sales process in the single-customer case. In particular, price p_t can be chosen at the beginning of the horizon assuming that no sales has occurred during $[0,t)$. If the customer purchases the product at some period $s<t$, then prices after time s have no impact on the seller's payoff.

Remark 2 (Notation). Throughout the paper, we often use the subscripts "M" and "S" to differentiate the notation that we used in the models with myopic and strategic consumers, respectively. For instance, $\mathcal{R}_{\mathrm{M}}(v,\tau,p)$ is the seller's regret when nature selects a myopic customer with type (v,τ) and the seller chooses the price path p. Similarly, we will denote by R_{M}^* and R_{S}^* the seller's minimum worst-case regret when facing myopic or strategic customers, respectively.

The symbols $\wedge$ and $\vee$ are used interchangeably with "min" and "max," that is, $a\wedge b=\min\{a,b\}$ and $a\vee b=\max\{a,b\}$. We use the terms "decreasing" and "increasing" to refer to weakly decreasing and increasing functions. We add the modifier "strictly" whenever that is not the case.

4. Selling to Myopic Customers

In this section we characterize pricing strategies that minimize the seller's regret for the case in which the customer uses a myopic purchasing strategy, that is, he buys the product as soon as his valuation is equal to or exceeds the posted price. Though myopic, our customers are assumed to be patient, remaining in the system until they make purchases or the end of the sales horizon is reached.

At the core of the discussion that follows are the notions of valuation and delay regrets. To build some intuition on how these two types of regrets impact the seller's pricing strategy, let us consider the price path p_t depicted in Figure 1. Given this path, nature can select a consumer whose type (v,τ) lies above the price path—e.g., consumer "a" in the figure with type (v_a,τ_a)—or whose type lies below the price path—such as consumer "b" with type (v_b,τ_b) in the figure.

If nature chooses consumer "a," then he buys immediately upon arrival at time τ_a and pays the price p_{τ_a}. The corresponding regret is $e^{-r\tau_a}(v_a-p_{\tau_a})$. However, nature can increase this regret by choosing instead consumer "A" with type $(\bar{v},\tau_a)$. This consumer "A" also buys immediately at time τ_a and the seller's regret increases to $e^{-r\tau_a}(\bar{v}-p_{\tau_a})$.

Figure 1. (Color online) Valuation and delay regrets with myopic customers.

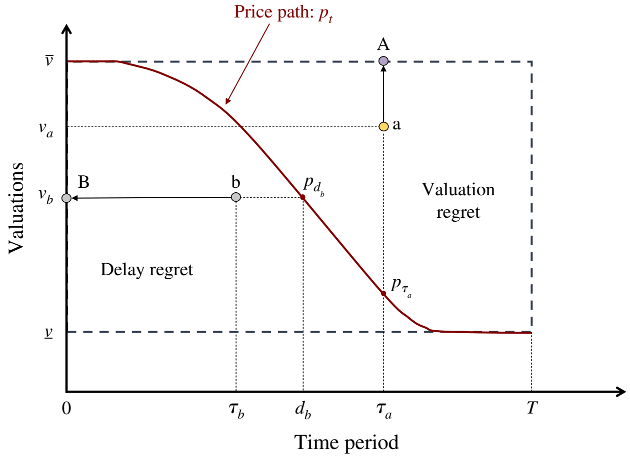

Hence, any type (v, τ) above the price path is dominated (in the sense of increasing the seller's regret) by the type $(\bar{v}, \tau)$ in the upper boundary of the set $\mathscr{D}$, which we denote by $\bar{\mathscr{D}} := \{(v, \tau) \in \mathscr{D}\colon v = \bar{v}\}$. Consumers with type (v, τ) above the price path create *valuation regret*.

Suppose now that nature picks consumer "b" whose type (v_b, τ_b) is below the price path. In this case, the consumer does not buy immediately and must wait until time d_b to purchase the product at price p_{d_b}. In this case, the price p_{d_b} is equal to v_b by the continuity of the price path. Hence, the corresponding regret is equal to $(e^{-r\tau_b} - e^{-rd_b})v_b$. Again, nature can increase this regret to $(1 - e^{-rd_b})v_b$ by choosing consumer "B" instead. This argument requires the price path p to be decreasing. It follows that any type (v, τ) below the price path is dominated by the type $(v, 0)$ in the left boundary of the set $\mathscr{D}$, which we denote by $\mathscr{D}_0 := \{(v, \tau) \in \mathscr{D}\colon \tau = 0\}$. In this case, consumers with type (v, τ) below the price path create *delay regret*.

At this point it should be intuitively clear that if the seller tries to reduce the valuation regret by increasing the price path then the delay regret will increase and vice versa. As a result, an optimal price strategy must balance these two types of regrets as we show below.

Consider an arbitrary price path $p \in \mathscr{P}$ and let us evaluate its performance by looking at the valuation and delay regrets it generates. Based on our previous discussion a myopic customer can only generate one type of regret if the price path is continuous. Indeed, a consumer with type (v, τ) generates the valuation regret $e^{-r\tau}(v - p_\tau)$ if $v \geqslant p_\tau$ or the delay regret $(e^{-r\tau} - e^{-rd_M(v, \tau, p)})v$ if $v < p_\tau$. This delay regret includes the case in which the customer is priced out of the market (i.e., the price path $\{p_t\colon \tau \leqslant t \leqslant T\}$ is strictly greater than v) since in this case $d_M(v, \tau, p) = \infty$.

With the previous decomposition of the regret in mind, let us consider the set of functions $p \in \mathscr{P}$ that have a supremum (or worst-case) valuation regret that is less than or equal to R, for some arbitrary $R > 0$. This is the set of functions p such that $\sup\{e^{-r\tau}(v - p_\tau) \mid (v, \tau) \in \mathscr{D}\} \leqslant R$. Since the argument inside the "sup" is monotonically increasing in v, we have the following result.

Lemma 1. *Let us define* $\underline{p}_t(R) := \{\bar{v} - e^{rt}R\} \vee \underline{v}$ *for all* $t \in [0, T]$. *A function* $p \in \mathscr{P}$ *has a worst-case valuation regret that is bounded above by* R *if and only if it satisfies the condition* $p_t \geqslant \underline{p}_t(R)$ *for all* $t \in [0, T]$.

Next, we would like to produce a similar result that characterizes the set of price paths that have a worst-case delay regret that is less than or equal to R. As we will see, we are able to produce such a result but with one caveat, namely, we have to restrict our characterization to the set of decreasing price paths in $\mathscr{P}$. Fortunately, our next lemma shows that this is not a serious limitation.

Lemma 2. *For any price path* $p \in \mathscr{P}$ *there exists a decreasing path* $\hat{p}$ *such that the seller's worst-case regret under* $\hat{p}$ *is less than or equal to the worst-case regret under* p, *that is,* $\mathscr{R}_{\mathrm{M}}(\hat{p}) \leqslant \mathscr{R}_{\mathrm{M}}(p)$.

We can now establish the following representation of the set of decreasing functions that have a worst-case delay regret bounded above by R.

Lemma 3. *Let us define* $\bar{p}_t(R) := \{\underline{v} \vee R/(1 - e^{-rt})\} \wedge \bar{v}$ *for all* $t \in [0, T]$. *A decreasing function* $p \in \mathscr{P}$ *has a worst-case delay regret that is bounded above by* R *if and only if it satisfies the two conditions*: (i) $p_t \leqslant \bar{p}_t(R)$ *for all* $t \in [0, T]$ *and* (ii) $p_T \leqslant R \vee \underline{v}$.

The first inequality above captures the delay regret originating from a customer delaying his purchase while the second one is associated with the risk of the consumer not purchasing the firm's product at all (which we represent as a delay until time $t = \infty$). Combining these three lemmas, we get the following proposition.

Proposition 2. *A decreasing function* $p \in \mathscr{P}$ *has a worst-case regret that is bounded above by* R *if and only if it satisfies the conditions* $p_T \leqslant R \vee \underline{v}$ *and* $\underline{p}_t(R) \leqslant p_t \leqslant \bar{p}_t(R)$ *for all* $t \in [0, T]$.

Equipped with this result, we can use a geometric argument to derive the set of optimal price paths. Figure 2 depicts three distinctive cases depending on the value of R. In the three panels, the dot at time $t = T = 1$ is located at the level $R \vee \underline{v}$ and is used the check the condition $p_T \leqslant R \vee \underline{v}$.

Consider first the situation in panel (a). In this case the value of R is sufficiently high so that $\underline{p}_t(R) < \bar{p}_t(R)$ for all $t \in [0, T]$. It follows that any decreasing price path $p \in \mathscr{P}$ inside the shaded area with $p_T \leqslant R \vee \underline{v}$ produces a worst-case regret that is bounded above by R.

Starting from the situation in panel (a), the seller has some room to reduce the value of the regret R. As she pushes R down, the function $\underline{p}_t(R)$ moves up and the function $\bar{p}_t(R)$ moves down. However, if she overshoots in this

Figure 2. (Color online) (a) Feasible regret region, (b) infeasible regret region and (c) optimal regret regions.

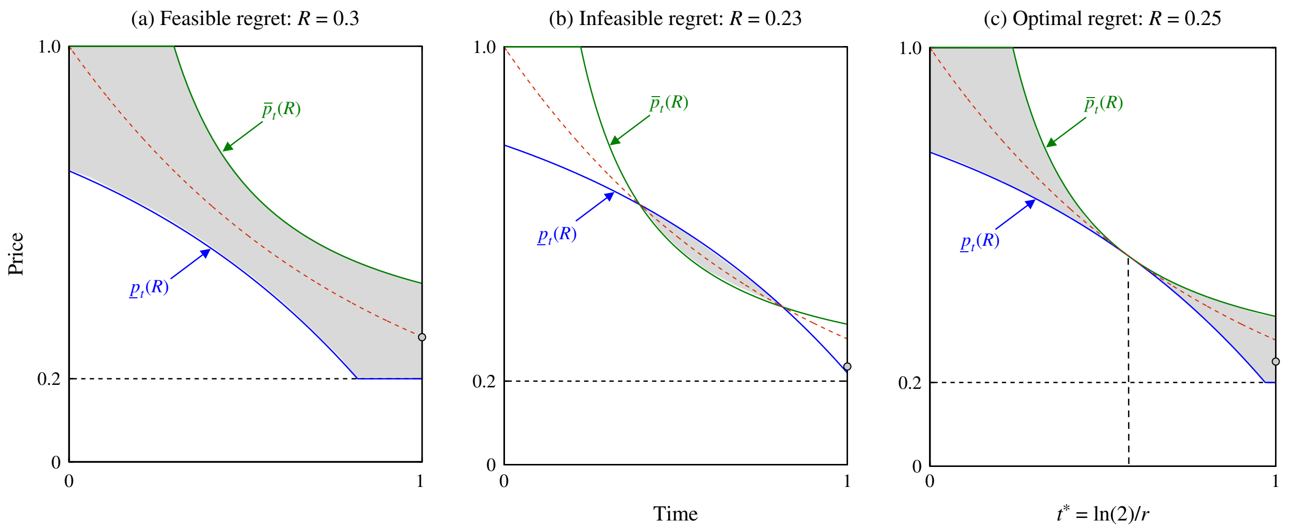

Notes. The decreasing dashed line represents the price path that, at each time point, equalizes the valuation regret for customers with $v = \bar{v}$ and the delay regret for customers with $\tau = 0$. The horizontal dashed line is located at the level $\underline{v}$ and the dot at time $t = 1$ is located at the level $R \vee \underline{v}$. Though within the optimal regret region, the decreasing dashed line in panel (c) is not actually optimal since it ends above the $R \vee \underline{v}$ dot. Data: $\bar{v} = T = 1$, $\underline{v} = 0.2$, $r = 1.2$ and $R = 0.3$, $R = 0.23$ and $R = 0.25$ in panels (a)–(c), respectively.

process and decreases the value of R too much, she can find herself in the situation depicted in panel (b). In this case, there is a region (shaded area) where $\underline{p}_t(R) > \bar{p}_t(R)$ and by Lemma 2 and Proposition 2 we know that there is no price path $p \in \mathcal{P}$ that can achieve a worst-case regret as low as R.

From the results in panels (a) and (b), we conclude that the seller would like to push down the value of R as long as $\underline{p}_t(R) \leqslant \bar{p}_t(R)$ for all $t \in [0, T]$ and $\underline{p}_T(R) \leqslant R \vee \underline{v}$. As R decreases, one of these two constraints will be violated first and will determine the optimal value of R. We can characterize the solution to the seller's problem by studying which constraint is binding for a given set of problem parameters and how precisely this constraint binds the optimization problem. To do so, we need to divide our parameter space into four regions.

The first region, which we denote by A_1, is the one that occurs if the time horizon T is long and the ratio $u := \underline{v}/\bar{v}$ is low. That is, there is a fair amount of uncertainty about the consumer's valuation and the firm has sufficient time to dynamically change prices. For a precise definition of the four regions, including A_1, see Figure 3. Panel (c) in Figure 2 represents the optimal region in the case where the problem parameters T and u belong to the region A_1, though only paths within the optimal region that end at $R \vee \underline{v}$ or below are optimal. In this situation, the two curves

Figure 3. Parameter regions A_1 to A_4.

$$A_1 := \left\{(u, T) \,|\, u \leqslant \frac{1}{2} \text{ and } T \geqslant \min\left\{\frac{\ln 3}{r}, \frac{\ln(4(1-u))}{r}\right\}\right\}$$

$$A_2 := \left\{(u, T) \,|\, u \geqslant \frac{1}{2} \text{ and } T \geqslant \frac{1}{r} \ln \frac{1}{u}\right\}$$

$$A_3 := \left\{(u, T) \,|\, T \leqslant \min\left\{\frac{\ln 3}{r}, \frac{1}{r} \ln\left(\frac{1}{u} - 1\right)\right\}\right\}$$

$$A_4 := [0, \bar{v}] \times \mathbb{R}_+ \backslash \{A_1 \cup A_2 \cup A_3\}$$

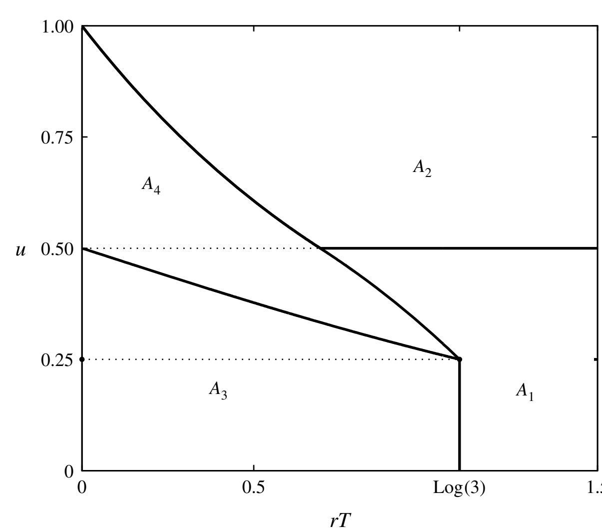

$\underline{p}_t(R)$ and $\bar{p}_t(R)$ intersect tangentially at a time $t = t^*$ without crossing. We call this time t^* the critical time. Any decreasing function p inside the shaded area with $p_T \leqslant R \vee \underline{v}$ is an optimal price path. It is interesting to see that in this situation every optimal price path $p \in \mathcal{P}$ must go through the point at which $\underline{p}_t(R)$ and $\bar{p}_t(R)$ touch. We can derive the coordinates of this point as well as the optimal value of R by imposing the conditions:

$$\underline{p}_{t^*}(R^*_{\mathrm{M}}) = \bar{p}_{t^*}(R^*_{\mathrm{M}}) \quad \text{and} \quad \frac{\mathrm{d}}{\mathrm{d}t}\underline{p}_t(R^*_{\mathrm{M}})\bigg|_{t=t^*} = \frac{\mathrm{d}}{\mathrm{d}t}\bar{p}_t(R^*_{\mathrm{M}})\bigg|_{t=t^*}. \tag{7}$$

After some straightforward calculations, we get that

$$t^* = \frac{1}{r}\ln(2), \quad \underline{p}_{t^*}(R^*_{\mathrm{M}}) = \bar{p}_{t^*}(R^*_{\mathrm{M}}) = \frac{\bar{v}}{2} \quad \text{and} \quad R^*_{\mathrm{M}} = \frac{\bar{v}}{4}. \tag{8}$$

Region A_2 represents the case where the horizon T is still long, but the ratio u is large, i.e., there is not a lot of market uncertainty. This region, with its long horizon T and low market uncertainty, represents a best case scenario for the firm. In this scenario, the binding constraint is still a critical time t^* where the functions $\underline{p}_{t^*}(R^*_{\mathrm{M}})$ and $\bar{p}_{t^*}(R^*_{\mathrm{M}})$ intersect, except that they are no longer tangential at t^*. Instead, what determines t^* is the minimum valuation $\underline{v}$. The minimax regret in region A_2 can be determined by solving the following pair of equations:

$$\underline{p}_{t^*}(R^*_{\mathrm{M}}) = \bar{p}_{t^*}(R^*_{\mathrm{M}}) = \underline{v}. \tag{9}$$

With a bit of algebra, we can show that the regret in this scenario is $R^*_{\mathrm{M}} = u(1-u)\bar{v}$.

Regions A_3 and A_4 are associated with a short selling horizon T. For these two regions, the binding constraint switches to $\underline{p}_T(R) \leqslant R \vee \underline{v}$. That is, there is no time t^* where the curves $\underline{p}_{t^*}(R^*_{\mathrm{M}})$ and $\bar{p}_{t^*}(R^*_{\mathrm{M}})$ intersect. Instead, we call the end of the horizon T the critical time t^* in regions A_3 and A_4. In region A_3, the horizon T is short and the ratio u is low (there is a lot of market uncertainty). Region A_3 is thus the firm's worst case scenario. In this scenario, the precise binding constraint is $\underline{p}_T(R) \leqslant R$, as the main risk facing the seller is the no purchase risk. Region A_4 is characterized by a short T and a high u. In this case, the precise binding constraint is $\underline{p}_T(R) \leqslant \underline{v}$. Perhaps counterintuitively, the seller's regret is either insensitive to small changes in the discount rate r (regions A_1 and A_2), or is decreasing in r (regions A_3 and A_4). In the latter case, the seller's regret decreases in r because the worst-case customer in A_3 and A_4 is one that arrives at time T with valuation $\bar{v}$. When r increases, the regret caused by this customer time is reduced since it occurs at time T.

We are now ready to present the main result of this section, which synthesizes the discussion above.

Theorem 1 (Myopic Consumers). *Let $u = \underline{v}/\bar{v}$. The seller's minimum worst-case regret is equal to*

$$R^*_{\mathrm{M}} = \begin{cases} \dfrac{\bar{v}}{4} & \text{if } (u,T) \in A_1 \\ u(1-u)\bar{v} & \text{if } (u,T) \in A_2 \\ \dfrac{\bar{v}}{1+e^{rT}} & \text{if } (u,T) \in A_3 \\ e^{-rT}(1-u)\bar{v} & \text{if } (u,T) \in A_4. \end{cases}$$

In addition, any decreasing pricing strategy $p \in \mathcal{P}$ that satisfies

$$p_T \leqslant R^*_{\mathrm{M}} \vee \underline{v} \quad \text{and} \quad \underline{p}_t(R^*_{\mathrm{M}}) \leqslant p_t \leqslant \bar{p}_t(R^*_{\mathrm{M}}) \quad \text{for all } t \in [0,T]$$

is optimal and generates a worst-case regret equal to R^*_{M}.

It is not hard to see that the optimal regret R^*_{M} is nonincreasing in both T and u, implying that the seller is (weakly) better off if she increases the length of the selling season or if she has less uncertainty about customer's valuation.

We conclude this section discussing the implications of Theorem 1 on the seller's choice of an optimal price path and duration of the selling season. However, before embarking on this practical discussion, we provide an additional result that shows that our restriction to continuous decreasing price paths is without loss of optimality.

Theorem 2. *Let $\tilde{\mathcal{P}}$ be the set of functions from $[0,T]$ to $[\underline{v},\bar{v}]$ and let R^*_{M} be the optimal regret in Theorem 1, i.e., R^*_{M} is the seller's minimax regret within the set of continuous price paths $\mathcal{P}$. Then, the seller's worst-case regret under any $p \in \tilde{\mathcal{P}}$ is bounded below by R^*_{M}, that is $R^*_{\mathrm{M}} \leqslant \mathcal{R}_{\mathrm{M}}(p)$.*

On the properties of optimal price paths. Possibly, one of the most important insights of Theorem 1 is the fact that a price path that minimizes the maximum regret is not unique. The theorem also provides point-wise upper and lower bounds for the set of optimal price paths. Let us denote by $\bar{p}^*_t := \bar{p}_t(R^*_{\mathrm{M}})$ and $\underline{p}^*_t := \underline{p}_t(R^*_{\mathrm{M}})$ the least upper and greatest lower bounds of the set of the optimal pricing strategies. We now discuss these two extreme optimal price paths in more details.

Figure 4 depicts an example of the price path $\bar{p}^*_t$. In this example, the length of the selling horizon is $T = 30$, customer's valuation belongs to the interval $[0.4, 1]$ and the optimal worst-case regret is $R^*_{\mathrm{M}} = 0.25$. We have decomposed the optimal price path in three distinctive regions: an initial Markup period ($t \leqslant 6.4$ in Figure 4) during which the seller sets the price equal to the upper limit $\bar{v}$, a Markdown period ($t \in (6.4, 21.8)$) during which the price decreases monotonically to the lower limit $\underline{v}$ and a Clearance period ($t \in [21.8, 30]$) during which the price is kept constant at

Figure 4. (Color online) Optimal price path $\bar{p}_t^*$ with $\bar{v} = 1$, $\underline{v} = 0.4$, $T = 30$, and $r = 0.045$.

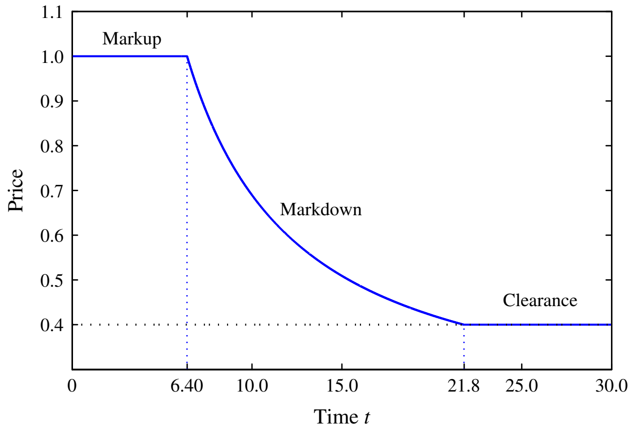

Note. In this example, $R_M^* = 0.25$.

the lower limit $\underline{v}$. The price path (and our choice of the names for these three regions) resembles standard pricing practice for seasonal or perishable products (see Bitran and Caldentey 2003 and references therein).

The presence of a Markup period in the optimal path in Figure 4 is somehow counterintuitive since the seller is effectively delaying sales until after this period. This feature of an optimal price vector suggests that the seller's worst-case regret R_M^* is not affected by high markup prices at the beginning of the selling season but rather by the choices of prices during the Markdown and Clearance periods. This idea is formalized in Theorem 1 that shows that the seller has a significant amount of flexibility in choosing prices before and after the critical time t^*, where t^* is defined as $t^* := \min\{\min\{t \mid \bar{p}_t^* = \underline{p}_t^*\}, T\}$ (see panel (c) in Figure 2). Indeed, any decreasing price function p from $[0, T]$ to $[\underline{v}, \bar{v}]$ is optimal if it is bounded above by $\bar{p}_t^*$ and bounded below by $\underline{p}_t^*$ and $p(T) \leqslant R_M^* \vee \underline{v}$.

On the other hand, the price at time t^* is uniquely determined. Interestingly, this result implies that the seller must be extremely careful in selecting the price at this time t^* since her worst-case regret is defined by her pricing strategy at this point in time. To be more specific, whenever the selling horizon T is sufficiently long (regions A_1 and A_2), the seller's regret is achieved when (i) a myopic customer with valuation $v = \bar{p}_{t^*}^* = \underline{p}_{t^*}^*$ arrives at the very beginning of the selling season and has to wait until the price is sufficiently low (at t^*) to make a purchase, or when (ii) a myopic customer with valuation $v = \bar{v}$ arrives at t^* and buys the product immediately. The seller's optimal pricing strategy in this myopic case is then designed to minimize the regret associated with these two types of events.

The fact that the optimal price path is not unique raises the question of how to select a particular one. From the definition of the lower bound, we have $\underline{p}_0^* < \bar{v}$,[5] that is, there always exists an optimal price path with no Markup period. Hence, we can view the lower bound $\underline{p}_t^*$ as a conservative pricing option that charges low prices throughout the selling season and, therefore, requires less markdowns. On the

Figure 5. (Color online) Upper bound $\bar{p}_t^*$ and lower bound $\underline{p}_t^*$ price paths for $\bar{v} = 1$, $\underline{v} = 0.32$, $T = 35$, and $r = 0.045$.

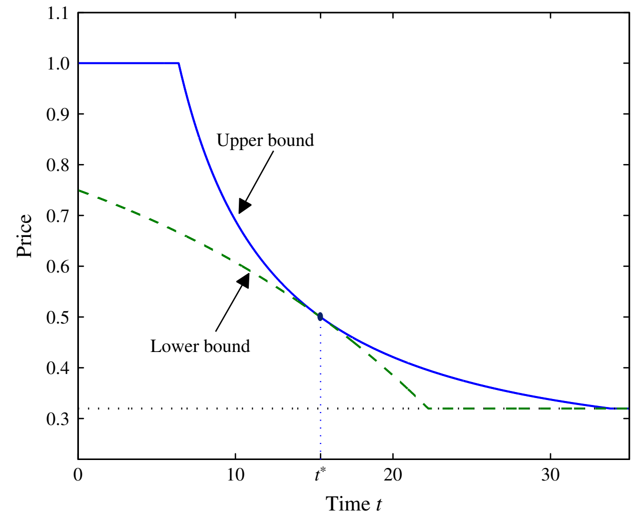

Notes. In this example, $R_M^* = 0.25$ and $t^* = 15.4$. The parameters in this example belong to parameter region A_1.

opposite extreme, the upper bound $\bar{p}_t^*$ is the most aggressive pricing option charging high prices but also imposing significant markdowns. By taking any optimal price path described in Theorem 1, the seller can balance these extreme pricing strategies and impose her particular preferences. Figure 5 depicts the upper and lower bound of the optimal prices, assuming that the problem parameters belong to region A_1.

Figure 6 shows the shapes of the upper and lower bounds of optimal prices in the three other regions of the parameter space. When the selling horizon is relatively long and the market uncertainty is relatively low (region A_2), all optimal price paths decrease to $\underline{v}$ before the sales season ends; in regions A_3 and A_4, the lower and upper bound price paths never intersect, and the critical time t^* is the end of the horizon T. What differentiates these two regions is whether the lower bound price path converges to $\underline{v}$ (region A_4) or not (region A_3).

Finally, there is one additional feature of some of the optimal price paths that is worth highlighting, namely their (quasi) convexity. Take for example the upper bound price path, which is equal to $\bar{p}_t^* = \bar{v} \wedge R_M^*/(1 - \exp(-rt))$, as depicted in Figure 5. This price path is a decreasing and quasi-convex function of t. That is, according to this pricing strategy, after an initial markup period, an optimal markdown policy should be "decelerating" over time in the sense that the magnitude of the markdowns should be decreasing. This recommendation contrasts with some mainstream results in the literature of dynamic pricing that suggest that the speed of markdowns should increase toward the end of the selling season.

To be concrete, consider the seminal work of Gallego and van Ryzin (1994), who consider a similar dynamic

Figure 6. Upper bound (solid curve) and lower bound (dash curve) price paths for $\bar{v}=1$ and $r=1$ with different values of u and T.

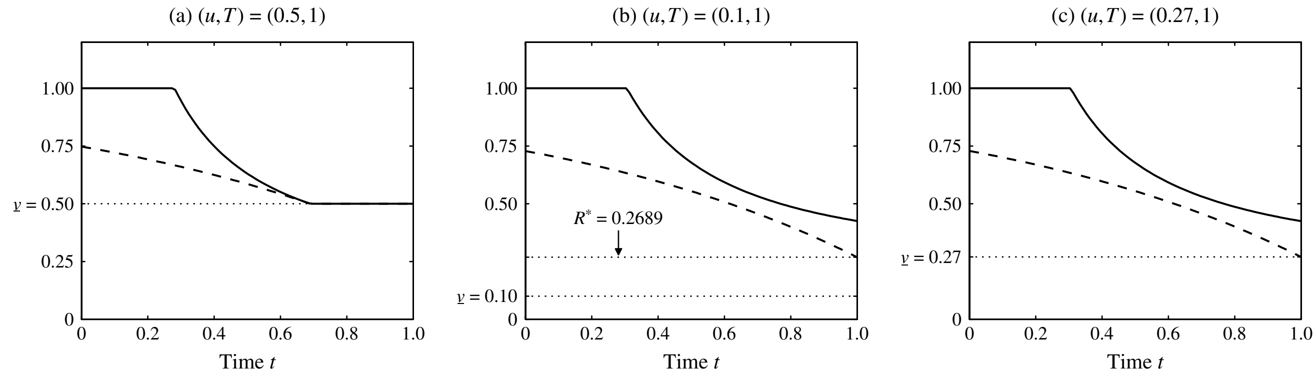

Note. Panels (a)–(c) represent parameter regions A_2, A_3, and A_4 respectively.

pricing problem with fixed selling horizon but assume that demand is driven by a Poisson process and consumers' willingness-to-pay for the product is randomly distributed among the population with a known probability distribution. Under this setting, an optimal pricing strategy can be derived using dynamic programming techniques. One of the key results in Gallego and van Ryzin (1994) is Theorem 1 that shows that the value function is strictly increasing and strictly concave in both remaining selling time and available inventory. Using this result, one can show that for many commonly used demand functions, the optimal price path is decreasing concave in calendar time. As we show in Appendix B, this fact is still true even if we include a seller discount rate to the Gallego and van Ryzin (1994) model.

On the optimal length of the selling horizon. Another important takeaway from Theorem 1 relates to the selection of the length of the selling horizon. If $(u,T)\in A_1\cup A_2$ then marginally increasing or decreasing the length of the selling season T will have no impact on the firm's regret. On the other hand, if $(u,T)\in A_3\cup A_4$ then marginally increasing T will reduce the firm's regret. With this in mind, we define an optimal length T^* for the selling season that is the minimum T that will achieve the minimum possible regret.

Let $R^*_{\mathrm{M}}(T)$ denote the minimax regret as a function of the selling horizon T and denote by $R^{**}_{\mathrm{M}}:=\inf_{T\geqslant 0}\{R^*_{\mathrm{M}}(T)\}$, which is the infimum of the regret as a function of T. We also define $T^*_{\mathrm{M}}:=\inf\{T\geqslant 0\colon R^*_{\mathrm{M}}(T)=R^{**}_{\mathrm{M}}\}$. The following result characterizes the value of T^*_{M} as a function of the discount factor r and market uncertainty $u=\underline{v}/\bar{v}$.

Corollary 1.

$$T^*_{\mathrm{M}}=\begin{cases}\dfrac{1}{r}\ln 3 & \textit{if } u\leqslant\dfrac{1}{4}\\[2mm] \dfrac{1}{r}\ln(4(1-u)) & \textit{if } \dfrac{1}{4}\leqslant u\leqslant\dfrac{1}{2}\\[2mm] \dfrac{1}{r}\ln\dfrac{1}{u} & \textit{if } u\geqslant\dfrac{1}{2}.\end{cases}.$$

One direct consequence of this result is that T^*_{M} is uniformly bounded above by $\ln 3/r$. In other words, independently of the seller's prior believes about the range of consumers' valuations $[\underline{v},\bar{v}]$, a sufficient condition to minimize the seller's regret is to choose a selling season of length $T=\ln 3/r$. The seller does not need to sell to the entire market when demand uncertainty is high. The optimal price path $\bar{p}_t(R)$ will converge to a value strictly above $\underline{v}$ when $\underline{v}$ is low. This occurs because selling to low valuation consumers generates little regret and, therefore, the seller can attain its minimax regret without selling to low valuation customers.

5. Selling to Strategic Customers

We now consider the case in which the firm sells to strategic customers. That is, the firm is selling to customers who are forward-looking in anticipating future prices and who time their purchases with the aim of maximizing their net discounted utility. Recall from Equation (4) that the seller's problem is to select a price path p that minimizes her regret

$$\begin{aligned}R^*_{\mathrm{S}}&=\inf_{p\in\mathscr{P}}\sup_{(v,\tau)\in\mathscr{D}}\mathscr{R}_{\mathrm{S}}(v,\tau,p)\\&=\inf_{p\in\mathscr{P}}\sup_{(v,\tau)\in\mathscr{D}}e^{-r\tau}v-e^{-rd_S(v,\tau,p)}p_{d_S(v,\tau,p)},\end{aligned}\qquad(10)$$

where $d_S(v,\tau,p)$ is the purchasing time of a strategic customer with type (v,τ) when prices are given by p, i.e., $d_S(v,\tau,p)=\min(\arg\max_{\tau\leqslant t\leqslant T}\{e^{-rt}(v-p_t)\mid v\geqslant p_t\})$. The problem of selling to a strategic customer can be interpreted as a three-player game, where the firm moves first selecting a pricing policy p, nature responds by selecting the value v and arrival time τ, and the customer moves last by selecting whether to purchase the firm's product and at which point in time. We solve this game by first looking at the customer's best response for a given triplet (v,τ,p) and computing the associate regret that he generates.

To this end, we would like to use a similar approach as the one we used with myopic consumers to partition the space of consumers' type $\mathcal{D}$ into those that generate valuation regret and those that generate delay regret. However, the strategic nature of consumers' purchasing behavior makes this segmentation less useful for the purpose of analysis. A strategic consumer can generate simultaneously valuation and delay regrets despite the continuity of the price path since such a customer might wait if prices are dropping sufficiently fast. For this reason, it is not just the path p but also its *modulus of continuity* that determine which customer types buy immediately at their arrival time and which ones delay their purchase. To formalize this idea, we introduce the following definition.

Definition 1 (Threshold Valuations). For a given price path $p \in \mathcal{P}$ let us define the real-extended function $\tilde{p}(p)\colon [0,T] \to [\underline{v}, \infty]$ such that

$$\tilde{p}(p)_t := p_t + \sup_{s\in(t,T]} \left\{ \frac{e^{-rs}}{e^{-rt} - e^{-rs}} (p_t - p_s)^+ \right\}. \tag{11}$$

For notational simplicity, we drop the dependence of $\tilde{p}(p)$ on p. The threshold valuation function $\tilde{p}_t$ plays a predominant role in characterizing strategic consumers' purchasing behavior. Its importance lies on the following key property: a strategic customer with type (v,τ) buys immediately at time τ if and only if $v \geqslant \tilde{p}_\tau$. Indeed, it is not hard to see that the condition

$$e^{-r\tau}(v - p_\tau) \geqslant e^{-rs}(v - p_s)^+ \quad \text{for all } s \in (\tau, T],$$

holds if and only if $v \geqslant \tilde{p}_\tau$. As a result, we can view strategic customers as acting "myopically" with respect to $\tilde{p}$ instead of p, that is, they buy the product as soon as their valuations exceeds $\tilde{p}$. Using this property of $\tilde{p}$, it is possible to draw a parallel with the case of myopic consumers to extrapolate the geometric method that we used to characterize optimal pricing strategies in the previous section (see Figure 9 and the discussion that follows). However, this line of attack is less effective in this case because working with threshold valuation paths $\tilde{p}_t$ is more cumbersome than working with price paths p_t. The reason for this is that the set of continuous price paths $\mathcal{P}$ is "too large" in this case in sense that $\tilde{p}$ is ill-defined at those times at which p_t is not differentiable. Hence, further restrictions on the set of price paths would be needed to develop a geometric solution as we did in the myopic case. Instead of using this route, we tackle the seller's minimum regret problem in (10) using a different approach consisting of the following two main steps. In the first step, we consider the special case in which nature is restricted to select consumers from the set $\mathcal{D}_0 = \{(v,\tau) \in \mathcal{D}\colon \tau = 0\}$, i.e., when all consumers are in the market at the beginning of the selling season. We solve this particular instance by adapting the machinery of mechanism design, and in particular the one on optimal screening as in Mirrlees (1971) (see also Mussa and Rosen 1978), to fit our robust minimax formulation. Then, in the second step, we show that if the seller uses the optimal price path identified in the first step then nature will select consumers that arrive at time $t = 0$, i.e., from the set $\mathcal{D}_0$. As a result, we will be able to conclude that the proposed price path is also optimal with respect to the larger set $\mathcal{D}$.

5.1. Selling to Strategic Customers With Types in $\mathcal{D}_0$

We start the analysis outlined in the previous paragraph by looking at the following special instance of problem (10):

$$\begin{aligned} R_0^* &:= \inf_{p\in\mathcal{P}} \sup_{(v,\tau)\in\mathcal{D}_0} \mathcal{R}_S(v,\tau,p) \\ &= \inf_{p\in\mathcal{P}} \sup_{\underline{v}\leqslant v\leqslant\bar{v}} \{v - e^{-rd_S(v,0,p)} p_{d_S(v,0,p)}\}. \end{aligned} \tag{12}$$

We use the subscript "0" to emphasize the fact that the regret R_0^* is calculated under the assumption that nature is restricted to the set $\mathcal{D}_0$.

To solve problem (12) we first look at the consumer's subproblem. Let $U(v,p)$ be the maximum utility that a consumer with valuation $v \in [\underline{v}, \bar{v}]$ arriving at time $t = 0$ can get if the seller chooses the price path $p \in \mathcal{P}$. That is,

$$U(v,p) := \max_{0\leqslant t\leqslant T} \{e^{-rt}(v - p_t)^+\}.$$

The positive part in the definition of $U(v,p)$ captures the consumer's *individual rationality* constraint. That is, if $v < \min\{p_t\colon 0 \leqslant t \leqslant T\}$ then the customer leaves the market without buying the product and gets zero utility. To simplify our notation, let us introduce the following change of variables: $\theta := e^{-rt}$ and $f(\theta) := e^{-rt} p_t$. It follows that the consumer's utility maximization problem can be rewritten as

$$U(v,f) := \max_{\theta_0\leqslant\theta\leqslant 1} (\theta v - f(\theta))^+, \tag{13}$$

where $\theta_0 := e^{-rT}$. In this formulation, the consumer action is represented by θ while the seller's strategy is given by the function $f(\theta)$. Our requirement that price paths belong to $\mathcal{P}$ then translates to requiring that $f \in \mathcal{F}$, which is the set of continuous functions such that $f(\theta)/\theta \in [\underline{v}, \bar{v}]$ for all $\theta \in [\theta_0, 1]$. For a specific choice of $f \in \mathcal{F}$, the seller can divide the range of valuations $[\underline{v}, \bar{v}]$ into those that buy and those that do not buy the product. Indeed, for a given f, let us define

$$\hat{v}_f := \min_{\theta_0\leqslant\theta\leqslant 1} \left\{ \frac{f(\theta)}{\theta} \right\}. \tag{14}$$

It follows that a consumer with valuation $v \in [\underline{v}, \hat{v}_f)$ leaves the market without buying while one with $v \in [\hat{v}_f, \bar{v}]$ buys the product and get a non-negative utility. Also, note that for any $f \in \mathcal{F}$ we have that $\hat{v}_f \in [\underline{v}, \bar{v}]$.

For a given $f \in \mathcal{F}$, and for a given consumer with valuation $v \in [\hat{v}_f, \bar{v}]$, we can define the consumer's optimal strategy

$$\theta_f(v) := \max\left\{ \underset{\theta_0\leqslant\theta\leqslant 1}{\arg\max} \{\theta v - f(\theta)\} \right\} \quad \text{for all } v \in [\hat{v}_f, \bar{v}].$$

The outer "max" in our definition of $\theta_f(v)$ captures our assumption that in case of indifference consumers prefer to buy as early as possible. By the definition of $\hat{v}_f$ in Equation (14), we have that $U(\hat{v}_f, f) = 0$. Also, from convex analysis, we get the following useful properties about $U(v, f)$ and $\theta_f(v)$.

LEMMA 4. *Let $f \in \mathcal{F}$. The function $U(v, f)$ is increasing and convex in $[\underline{v}, \bar{v}]$. In addition, the maximizer $\theta_f(v)$ is right-continuous and nondecreasing in the interval $[\hat{v}_f, \bar{v}]$. Finally, for all $v \in [\hat{v}_f, \bar{v}]$:*

$$U(v, f) = \int_{\hat{v}_f}^{v} \theta_f(x)\,\mathrm{d}x.$$

This lemma is often referred to in mechanism design as the Envelope Theorem. From the definition of $\theta_f(v)$ and the previous lemma, we get that for all $v \in [\hat{v}_f, \bar{v}]$

$$U(v, f) = \theta_f(v)v - f(\theta_f(v)) = \int_{\hat{v}_f}^{v} \theta_f(x)\,\mathrm{d}x. \qquad (15)$$

With this implicit derivation of the consumer's utility and optimal strategy, let us turn to the seller's optimization problem and let us compute the regret generated by a customer with valuation $v \in [\underline{v}, \bar{v}]$. We identify two cases.

—*Case* 1: Suppose that $v \in [\underline{v}, \hat{v}_f)$. Then, the customer leaves the market without buying and generates a regret equal to v.

—*Case* 2: Suppose that $v \in [\hat{v}_f, \bar{v}]$. Then, the customer buys the product and generates a regret equal to $v - f(\theta_f(v))$. From Equation (15) we can rewrite this regret a follows

$$\int_{\hat{v}_f}^{v} \theta_f(x)\,\mathrm{d}x + v(1 - \theta_f(v)).$$

Combining Cases 1 and 2, we conclude that the seller's minimum worst-case regret problem is given by

$$\inf_{f \in \mathcal{F}} \mathcal{R}(f) \quad \text{where}$$

$$\mathcal{R}(f) := \hat{v}_f \mathbb{1}(\hat{v}_f > \underline{v}) \vee \max_{v \in [\hat{v}_f, \bar{v}]} \left\{ \int_{\hat{v}_f}^{v} \theta_f(x)\,\mathrm{d}x + v(1 - \theta_f(v)) \right\}.$$

The indicator $\mathbb{1}(\hat{v}_f > \underline{v})$ is needed because Case 1 above only happens if $\hat{v}_f > \underline{v}$.

Instead of solving the optimization above directly on the set of functions $f \in \mathcal{F}$ we will use the function $\theta_f(x)$ as our decision variable. We will then recover the value of $f(\theta)$ from the value of θ using Equation (15). To proceed, note that according to our definition of $\theta_f(x)$ and to Lemma 4, we must restrict our attention to functions $\theta(x)$ in the following set.

DEFINITION 2. Let $\Theta(\hat{v})$ be the set of right-continuous and nondecreasing functions θ: $[\hat{v}, \bar{v}] \to [\theta_0, 1]$.

The following intermediate result plays a key role in our derivation of an optimal solution.

PROPOSITION 3. *Let us fix $\hat{v} \in [\bar{v}, \underline{v}]$ and consider the optimization problem*

$$R_{\hat{v}} := \hat{v}\mathbb{1}(\hat{v} > \underline{v}) \vee \inf_{\theta \in \Theta(\hat{v})} \max_{v \in [\hat{v}, \bar{v}]} \left\{ \int_{\hat{v}}^{v} \theta(x)\,\mathrm{d}x + v(1 - \theta(v)) \right\}.$$

Then,

$$\theta^*(x) = \theta_0 \vee (1 - \ln(\bar{v}) + \ln(x)), \quad x \in [\hat{v}, \bar{v}],$$

is an optimal solution and

$$R_{\hat{v}} = \hat{v}\mathbb{1}(\hat{v} > \underline{v}) \vee (\tilde{v} + \hat{v}(\ln(\bar{v}) - \ln(\tilde{v}) - 1)),$$
$$\text{where } \tilde{v} := \max\{\hat{v}, \bar{v}\exp(\theta_0 - 1)\}.$$

Equipped with Proposition 3, we can now fully characterize the solution to the case in which nature is restricted to the set $\mathcal{D}_0$. Essentially, this boils down to finding the value of the threshold $\hat{v}$ that minimizes $R_{\hat{v}}$. We can then reverse our change of variables and transform an optimal strategy in the θ-space to an optimal strategy in the original price space. We summarize this derivation in the following result, which uses the following definition.

$$d^*(v) := -\frac{1}{r}\ln(1 - \ln(\bar{v}) + \ln(v)) \wedge T, \quad \text{for all } v \in [\underline{v}, \bar{v}].$$

THEOREM 3. *Suppose nature selects consumers from the set $\mathcal{D}_0 = \{(v, \tau) \in \mathcal{D}\colon \tau = 0\}$. The value of $\hat{v} \in [\underline{v}, \bar{v}]$ that minimizes the regret $R_{\hat{v}}$ in Proposition 3 is given by*

$$\hat{v}^* = \max\left\{\underline{v},\ \frac{\bar{v}\exp(e^{-rT} - 1)}{1 + e^{-rT}}\right\}$$

and the seller's minimax regret is equal to

$$R_0^* = R_{\hat{v}^*} = \tilde{v}^* + \hat{v}^*(\ln(\bar{v}) - \ln(\tilde{v}^*) - 1),$$
$$\text{where } \tilde{v}^* := \max\{\hat{v}^*,\ \bar{v}\exp(e^{-rT} - 1)\}.$$

An optimal pricing strategy for the seller is given by

$$p_S^*(t) = \begin{cases} e^{rt}(\bar{v}\exp(e^{-rt} - 1) - R_0^*) & \text{for all } t \in [0, d^*(\hat{v}^*)] \\ \underline{v} & \text{for all } t \in (d^*(\hat{v}^*), T]. \end{cases}$$

As a result, a customer with valuation $v \in [\underline{v}, \hat{v}^*)$ leaves the market without buying the product while one with valuation $v \in [\hat{v}^*, \bar{v}]$ buys at time $d^*(v)$.

A few remarks about the result in Theorem 3 and properties of an optimal price path are in order. First, the price path equalizes all regrets generated by customers with valuations above $\tilde{v}^*$. Second, the optimal price path we find is a convex and decreasing function of time. Actually, the price path is strictly deceasing until it reaches the lower bound $\underline{v}$ and then remains constant at this value. Whether or not the price path reaches this lower bound during the selling season depends on the particular instance under consideration. To explore this issue in more detail, let us consider

Figure 7. (Color online) Optimal price path $p^*_S(t)$ and corresponding threshold valuation $\tilde{p}^*_S(t)$ for different values of $\underline{v}$.

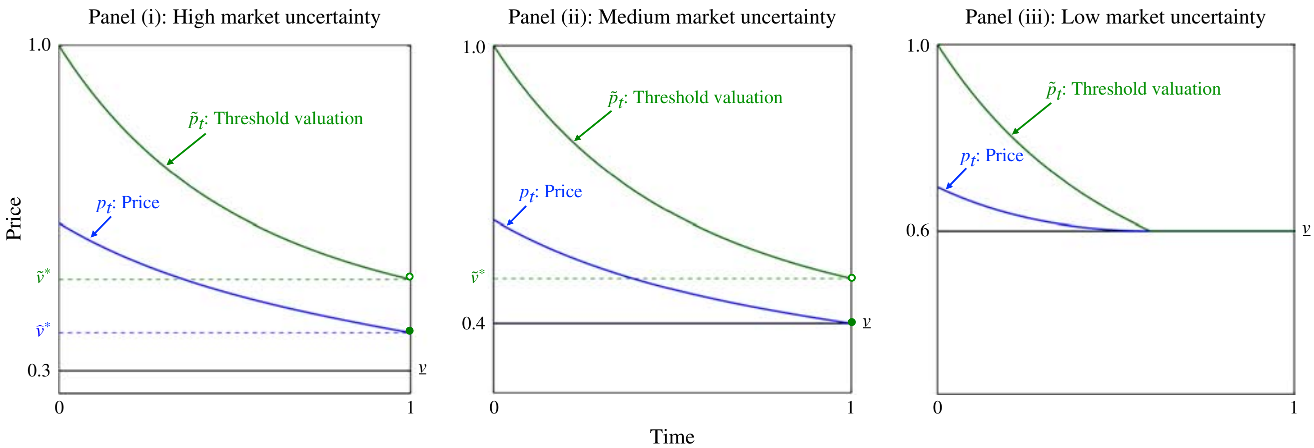

Notes. DATA: $\bar{v} = T = 1$, $r = 1.2$ and $\underline{v} = 0.3$, $\underline{v} = 0.4$ and $\underline{v} = 0.6$ in panels (i)–(iii), respectively. Panels (i)–(iii) correspond to parameter regions B_1, B_2, and B_3.

the numerical examples in Figure 7, which depicts the optimal price path $p^*_S(t)$ and corresponding threshold valuation $\tilde{p}^*_S(t)$ introduced in Definition 1.

We distinguish three different pricing regimes depending on the values of the parameters. These are panels (i)–(iii) in the figure. These cases correspond to three different parameter regions, which we represent by B_1, B_2 and B_3, and formally specify in Figure 8. If we fix the selling horizon T, these regimes depend on the relative value of $\bar{v}$ and $\underline{v}$. Region B_1 (panel (i)) corresponds to the case in which the ratio $u := \underline{v}/\bar{v}$ is low, i.e., when the seller faces *high market uncertainty*. Region B_2 (panel (ii)) represents the case when the value of u is moderate, which we label as *medium market uncertainty*. Finally, region B_3 (panel (iii)) corresponds to high u; we say that the seller faces *low market uncertainty*. In regions B_1 and B_2, the function $\tilde{p}^*_S(t)$ has a discontinuity at $t = T$ since $\lim_{t\uparrow T} \tilde{p}^*_S(t) = \tilde{v}^* > \hat{v}^* = \tilde{p}^*_S(T) = p^*_S(T)$.

A distinctive feature of region B_1, where the seller faces high market uncertainty, is that $p^*_S(T) > \underline{v}$ and so the price path is strictly greater than $\underline{v}$ for all $t \in [0, T]$. As a result, consumers with valuation below $p^*_S(T)$ are priced out of the market. Intuitively, the seller is better off excluding low valuation consumers because otherwise she would have to decrease the price significantly (in a relatively short amount of time) to serve them. It turns out that such a steep price reduction will induce high-valuation customers to delay their purchase time, which increases the seller's overall worst-case regret. Another interesting property of this case is that consumers with valuation in the range $[\hat{v}^*, \tilde{v}]$ will buy only at time T at price $p^*_S(T)$. (Graphically, this feature is reflected on the fact that $\tilde{p}^*_S(t) > \tilde{v}^*$ for all $t \in [0, T)$ while $\tilde{p}^*_S(T) = p^*_S(T) = \hat{v}^*$.) In other words, the seller is inducing a form of last-minute sale despite the fact that the price is decreasing smoothly during the selling season. We can

Figure 8. Parameter regions B_1 to B_3.

$$B_1 = \left\{(u, T) \mid u \leqslant \frac{\exp(e^{-rT} - 1)}{1 + \exp(-rT)}\right\}$$

$$B_2 = \left\{(u, T) \;\middle|\; \frac{\exp(e^{-rT} - 1)}{1 + \exp(-rT)} \leqslant u \leqslant \exp(e^{-rT} - 1)\right\}$$

$$B_3 = \{(u, T) \mid u \geqslant \exp(e^{-rT} - 1)\}$$

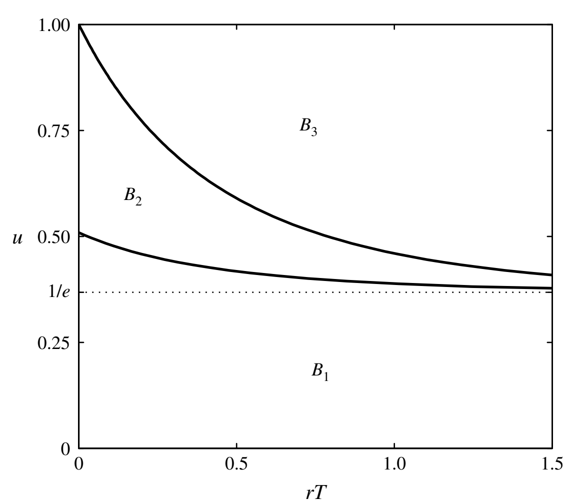

rephrase this property by saying that the optimal price path does not *separate* customers with valuations in $[\hat{v}^*, \tilde{v}]$ and instead *pools* them into a single class that buys at time T.

In region B_2, where market uncertainty is moderate, the price path reaches $\underline{v}$ exactly at time T and so all customer types are served. However, the optimal price path still exhibits the pooling phenomenon of the previous case (since $\tilde{p}^*_S(t) > \tilde{v}^*$ for all $t \in [0, T)$) and all customers with valuation in $[\underline{v}, \tilde{v}^*]$ buy exactly at the end of the selling horizon.

Finally, in region B_3, where market uncertainty is low, the price path also reaches the lower bound $\underline{v}$ as in the previous case and so all customer types are served. The difference here is that this minimum price is reached before the end of the selling season and so the threshold valuation $\tilde{p}^*_S(t)$ also reaches $\underline{v}$ before T (actually, at exactly the same time as $p^*_S(t)$). This means that the price path can now separate all customer types, each one buying at a different time. In other words, in this case the price path can perfectly screen all customers types based on their purchasing time. It is also worth noticing that in this case the seller can stop selling the product at the time the price path reaches $\underline{v}$ without increasing her worst-case regret, that is, the selling season is longer than needed (see Corollary 3 below for further analysis of this issue).

The seller's regret decreases with the discount rate r because the customer is pressured to buy earlier when r is higher. The following corollary deals with the limiting case where the selling horizon T is infinitely long. In this case, region B_2 ceases to exist, but parameter regions B_1 and B_3 do occur depending on the degree of market uncertainty.

Corollary 2. *Suppose nature selects consumers from the set $\mathscr{D}_0 = \{(v, \tau) \in \mathscr{D}\colon\ \tau = 0\}$ and the selling horizon is infinite, $T = \infty$. We distinguish two cases:*

(a) *If $\underline{v} \leqslant \bar{v}/e$, then $R^*_0 = \bar{v}/e$ and $p^*_S(t) = \bar{v}e^{rt-1}(\exp(e^{-rt}) - 1)$ for all $t \geqslant 0$. The price path is strictly decreasing and converges to $p^*_S(\infty) := \lim_{t\to\infty} p^*_S(t) = \bar{v}/e$.*

(b) *If $\underline{v} > \bar{v}/e$, then $R^*_0 = \underline{v}(\ln(\bar{v}) - \ln(\underline{v}))$ and the price path is equal to*

$$p^*_S(t) = \begin{cases} e^{rt}(\bar{v}\exp(e^{-rt} - 1) - R^*_0) & \text{for all } t \in [0, d^*(\underline{v})] \\ \underline{v} & \text{for all } t \in (d^*(\underline{v}), \infty). \end{cases}$$

According to point (a), if $u = \underline{v}/\bar{v}$ is sufficiently low (lower than $1/e$) then the optimal price path will never reach $\underline{v}$, and the price regime identified in region B_1 will persist independently of the length of the selling horizon. On the other hand, when $u > 1/e$, the seller sets an optimal price path that eventually reaches the lower bound $\underline{v}$ as in region B_3.

The value of the minimax regret we obtain in Corollary 2(a), $\bar{v}/e$, is the same as obtained in a static model by Bergemann and Schlag (2008). In our model, if $T = 0$, the minimax regret when valuation uncertainty is high is equal to $\bar{v}/2$ instead (see Theorems 1 and 3). The difference is due to the fact that we do not allow for randomized policies. This coincidence suggests that randomization (which Bergemann and Schlag 2008 allows) and intertemporal pricing (which we allow) might be, in some sense, substitutable strategies.

Following on the previous corollary, our next result investigates the sensitivity of the optimal regret R^*_0 to the length of the selling horizon T. Using a similar notation to the one we used in the myopic case, let $R^*_0(T)$ denote the seller's minimum worst-regret as a function of the selling horizon T and let $R^{**}_0 := \inf_{T>0}\{R^*_0(T)\}$. We also define $T^*_S := \inf\{T \geqslant 0\colon\ R^*_0(T) = R^{**}_0\}$.

Corollary 3. *Suppose nature selects consumers from the set $\mathscr{D}_0 = \{(v, \tau) \in \mathscr{D}\colon\ \tau = 0\}$. The seller's optimal worst-case regret $R^*_0(T)$ is a decreasing function of T and*

$$T^*_S = \begin{cases} -\ln(1 + \ln(\underline{v}) - \ln(\bar{v}))/r & \text{if } \bar{v} < e\underline{v} \\ \infty & \text{if } \bar{v} \geqslant e\underline{v} \end{cases}$$

and

$$R^{**}_0 = \begin{cases} \bar{v}(\ln(\bar{v}) - \ln(\underline{v})) & \text{if } \bar{v} < e\underline{v} \\ \bar{v}/e & \text{if } \bar{v} \geqslant e\underline{v}. \end{cases}$$

It is interesting to contrast this result to Corollary 1. Recall that when the seller faces myopic consumers the optimal selling horizon T^*_M is uniformly bounded by $\ln(3)/r$. In this case with strategy consumers, however, the optimal selling horizon can be infinite if the market uncertain is large, i.e., when $\bar{v} > e\underline{v}$. With $T = \infty$, the minimax regret is strictly decreasing in $\underline{v}$ when the valuation uncertainty is low, but constant in $\underline{v}$ when $\underline{v}$ is below $\bar{v}/e$. In the latter case, the seller does not sell to the entire market despite the infinitely long time horizon.

The discount rate r and the time horizon T appear only as the product rT in Theorems 1 and 3. Therefore, the sensitivity analysis above with respect to T can easily be reinterpreted as sensitivity analysis with respect to r.

The result in Theorem 3 characterizes the solution to the seller's problem when nature is restricted to select customers from the set $\mathscr{D}_0$. Let us now turn to the general case in which nature can also select their arrival times.

5.2. Selling to Strategic Customers with Types in $\mathscr{D}$

The goal of this subsection is to show that the optimal price path $p^*_S(t)$ identified in Theorem 3 (for the special case in which all customers are in the market at $t = 0$) is also optimal for the case in which customers are allowed to arrive throughout the selling season.

Since $p^*_S \in \mathscr{P}$ and $\mathscr{D}_0 \subseteq \mathscr{D}$ the following inequalities hold trivially:

$$R^*_0 = \inf_{p\in\mathscr{P}} \sup_{(v,\tau)\in\mathscr{D}_0} \mathscr{R}_S(v, \tau, p) \leqslant R^*_S$$
$$= \inf_{p\in\mathscr{P}} \sup_{(v,\tau)\in\mathscr{D}} \mathscr{R}_S(v, \tau, p) \leqslant \sup_{(v,\tau)\in\mathscr{D}} \mathscr{R}_S(v, \tau, p^*_S).$$

Hence, to prove the optimality of $p^*_S(t)$ it suffices to show that $R^*_0 = \sup_{(v,\tau)\in\mathscr{D}} \mathscr{R}_S(v, \tau, p^*_S)$, which implies that

Figure 9. (Color online) Valuation and delay regrets with strategic consumers.

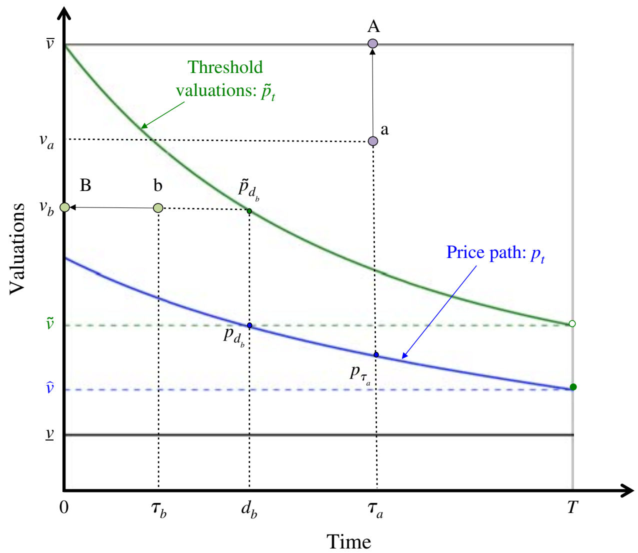

$R_0^* = R_S^*$. To this end, we will use a geometric argument (similar to the one we used in the myopic case) based on the threshold valuations function $\tilde{p}_S^*(t)$ induced by $p_S^*(t)$ (see Definition 1). Let us recall that the distinctive feature of $\tilde{p}_S^*(t)$ is that a consumer with valuation v buys the product as soon as $v \geqslant \tilde{p}_S^*(t)$. That is, the consumer acts myopically with respect to $\tilde{p}_S^*(t)$. Our geometric argument relies on the following monotonicity property.

Lemma 5. *Let $p_S^*(t)$ be the optimal price path in Theorem 3. Then, its corresponding threshold valuations function $\tilde{p}_S^*(t)$ is strictly decreasing in the interval $[0, d^*(\hat{v}^*)]$ and equal to $\underline{v}$ in the interval $(d^*(\hat{v}^*), T]$.*

Using the monotonicity of $\tilde{p}_S^*(t)$, we can now show that if the seller selects the pricing strategy p_S^* then nature would like to choose a customer from the set $\mathcal{D}_0$. For this, let us consider the example in Figure 9, which depicts the price path p_S^* and corresponding threshold valuation $\tilde{p}_S^*$.

Since $\tilde{p}_S^*$ is decreasing we can easily compute the regret generated by any customer type $(v, \tau) \in \mathcal{D}$. Furthermore, in computing the seller's worst-case regret, we can restrict ourselves to types (v, τ) for which either $v = \bar{v}$ or $\tau = 0$. Indeed, consider customer "a" in Figure 9 with type (v_a, τ_a). Since $v_a \geqslant \tilde{p}_{\tau_a}$, the customer buys immediately upon arrival and generates a regret $e^{-r\tau_a}(v_a - p_{\tau_a})$. However, customer "A" with type $(\bar{v}, \tau_a)$ also buys immediately and generates a higher regret. Similarly, consider customer "b" with type (v_b, τ_b). Since, $v_b < \tilde{p}_{\tau_b}$, this customer does not buy immediately and waits until time d_b when $\tilde{p}_{d_b} = v_b$ and generates a regret $e^{-r\tau_b} v_b - e^{-r d_b} p_{d_b}$. Again, it is not hard to see that customer "B" with type $(v_b, 0)$ will also buy at time d_b and generates a larger regret $v_b - e^{-r\, d_b}\, p_{d_b}$.

It follows from the previous discussion that if the seller selects the price path $p_S^*(t)$ in Theorem 3 then nature will choose a customer either from $\mathcal{D}_0 = \{(v, \tau)\colon\ \tau = 0\}$ or $\mathcal{D}_{\bar{v}} := \{(v, \tau)\colon\ v = \bar{v}\}$. Our next result shows that nature is always better off choosing from the set $\mathcal{D}_0$ and as a result $p_S^*(t)$ is indeed an optimal price path for the case with strategic consumers.

Theorem 4. *Let $p_S^*(t)$ and R_0^* be the price path and worst-case regret derived in Theorem 3 when nature is restricted to choose customers from the set $\mathcal{D}_0$. Then, we have that $R_0^* = \sup_{(v,\tau)\in\mathcal{D}} \mathcal{R}_S(v, \tau, p_S^*)$. As a result, we also have that*

$$R_0^* = R_S^* = \inf_{p\in\mathcal{P}} \sup_{(v,\tau)\in\mathcal{D}} \mathcal{R}_S(v, \tau, p),$$

and $p_S^*(t)$ is also an optimal price path for the case with strategic consumers drawn from the set $\mathcal{D}$.

It follows that the worst-case types that nature will select all have arrival time $\tau = 0$ and valuation $v \in [\tilde{v}^*, \bar{v}]$, where $\tilde{v}^*$ is the lowest valuation that purchases during the selling season. The optimal price path p generates the same regret for all customer types in this set.

6. Myopic vs. Strategic Customers

In this section, we compare pricing policies and the minimax regret values for the problems of selling to myopic and selling to strategic consumers. We also consider the question of how to sell to a mix of myopic and strategic consumers.

Figure 10 shows how the optimal price for selling to strategic customers compares to the upper and lower bound optimal prices for selling to myopic customers. In the instance plotted, strategic prices are initially lower

Figure 10. (Color online) Price paths for strategic customers, upper bound and lower bound price paths for myopic customers, with $\bar{v} = 1$, $\underline{v} = 0.4$, $r = 1$, and $T = 1.5$.

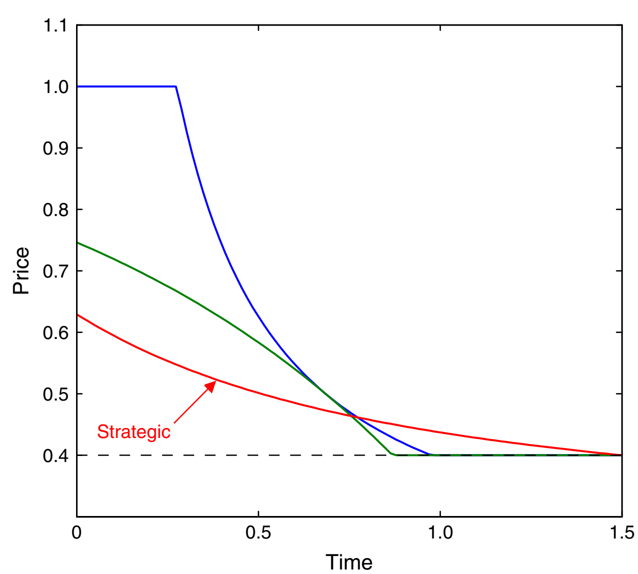

Figure 11. (Color online) Panels (a) and (b) illustrate the values of minimax regrets as functions of T and u, respectively. Panels (c) and (d) illustrate the ratios of minimax regrets from strategic and myopic customers, as functions of T and u, respectively.

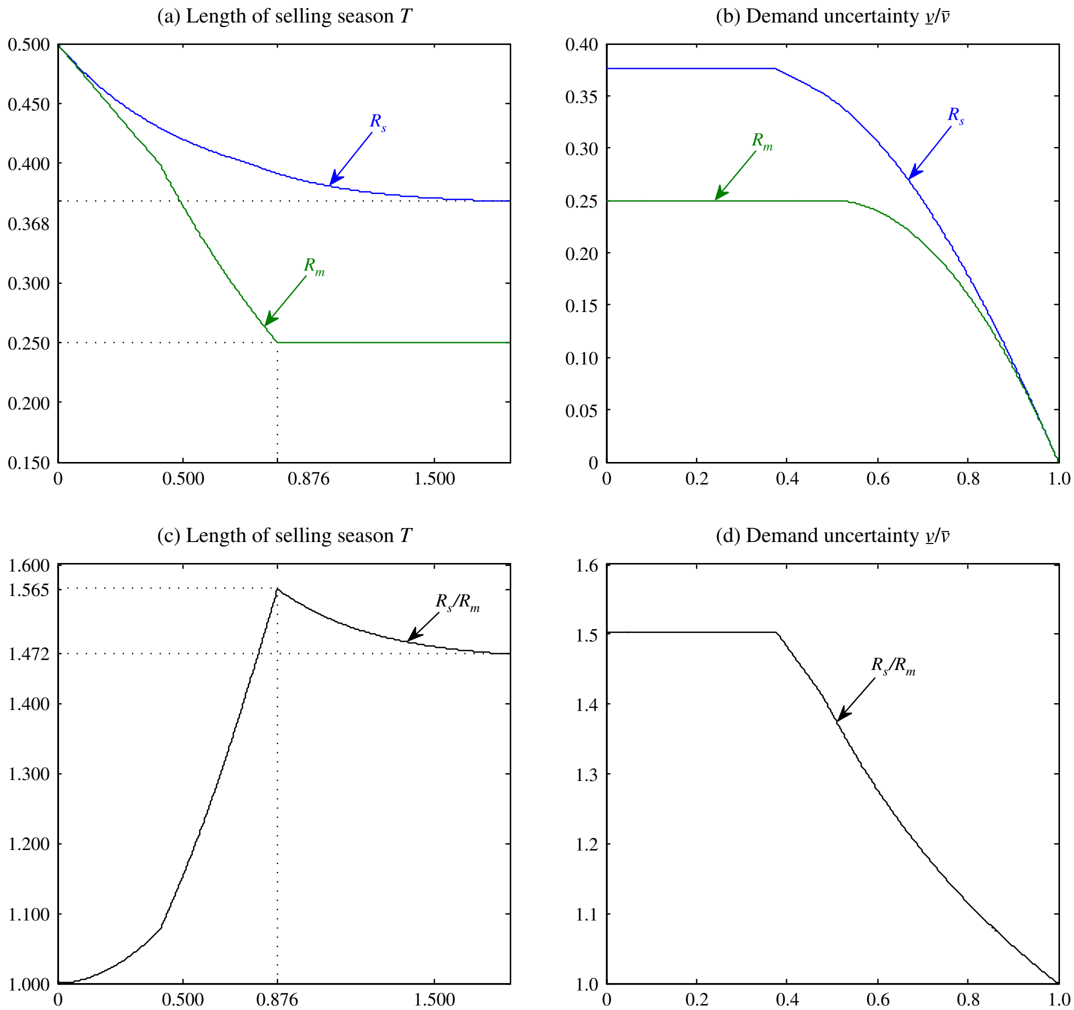

Note. In all the examples, $\bar{v} = 1$ and $r = 1$. In panels (a) and (c), $\underline{v} = 0.4$. In panels (b) and (d), $T = 1.5$.

than myopic prices, but since markdowns occur at a slower rate, strategic prices eventually become higher than myopic ones.

Our next result shows that strategic customers always lead to larger minimax regret than myopic ones. The firm has more freedom to do price skimming when customers are myopic. Our proof is based on showing that strategic customers always buy later and at lower prices than myopic ones.

Proposition 4. *The worst-case regret from selling to strategic customers is higher than the worst-case regret from selling to myopic ones, i.e.,* $R_S^* \geqslant R_M^*$.

Figure 11 contrasts the minimax regret values for the cases of myopic and strategic customers, with panel (a) showing regret values for different lengths of the selling season and panel (b) showing regret levels for different level of demand uncertainty. From Proposition 4, we know that the ratio of the regrets R_S/R_M must be always at least equal to 1. As panel (c) shows, the ratio of regrets is initially increasing in the horizon T as the screening power of dynamic pricing is more powerful with myopic than with strategic consumers. However, as Corollaries 1 and 3 show, arbitrarily long horizons are valuable under strategic consumers but not under myopic ones. Therefore, the ratio R_S/R_M decreases after a certain point in time.

Panel (d) of Figure 11 showcases what happens to the ratio R_S/R_M as market uncertainty decreases. For high market uncertainty parameters (low ratios $\underline{v}/\bar{v}$), price paths for both strategic and myopic consumers are independent of $\underline{v}$. Therefore, the ratio R_S/R_M is constant when market uncertainty is large. As we increase the value of $\underline{v}$, the ratio R_S/R_M is reduced. This occurs because screening is less important when market uncertainty is low, and screening is more effective when customers are myopic.

6.1. Mix of Myopic and Strategic Consumers

So far, we have assumed that there is a single type of consumers (myopic or strategic) in the market and that the seller knows the consumer type. In this subsection, we consider what happens if both types are present in the market simultaneously, and the fraction of consumers that are strategic is unknown to the seller. Results in Su (2007) and Mersereau and Zhang (2012) suggest that the structure of an optimal pricing strategy might be sensitive to the proportions of myopic and strategic consumers in the marketplace. To capture this type of market heterogeneity, in this section we extend our robust formulation by allowing nature to select the mix of myopic and strategic consumers.

Let us consider first the case of a single consumer. In this new setting, the type of this consumer is a triplet (v, τ, θ), where $v \in [\underline{v}, \bar{v}]$ is his valuation, $\tau \in [0, T]$ is his arrival time and $\theta \in \{M, S\}$ is his purchasing behavior, where "M" stands for myopic and "S" for strategic. We define the set of the consumers' types to be $\mathcal{D}_\Theta := [\underline{v}, \bar{v}] \times [0, T] \times \{M, S\}$. We also define the projections $\mathcal{D}_M := \mathcal{D}_\Theta \cap \{\theta = M\}$ and $\mathcal{D}_S := \mathcal{D}_\Theta \cap \{\theta = S\}$, corresponding to the cases when all consumers are myopic or strategic, respectively.

Depending on whether nature can select a consumer's type from the set $\mathcal{D}_\Theta$, $\mathcal{D}_M$ or $\mathcal{D}_S$, the seller's minimax regret is given by

$$\begin{aligned} R_j^* &:= \min_{p \geqslant 0} \max_{(v,\tau,\theta) \in \mathcal{D}_j} \mathcal{R}(v, \tau, \theta, p) \\ &= \min_{p \geqslant 0} \max_{(v,\tau,\theta) \in \mathcal{D}_j} e^{-r\tau} v - e^{-rd(v,\tau,\theta,p)} p_{d(v,\tau,\theta,p)} \\ &\qquad \text{for } j = \Theta, M, S, \end{aligned}$$

where $d(v, \tau, \theta, p)$ is the purchasing time of a (v, τ, θ)-consumer that faces the price path $p \in \mathcal{P}$. We also denote by p_j^* an optimal price path for $j = \Theta, M, S$.

According to Proposition 4, we know that $R_S^* \geqslant R_M^*$. Also, since $\mathcal{D}_S \subseteq \mathcal{D}_\Theta$, we also have that $R_\Theta^* \geqslant R_S^*$. As a result, we can rank the seller's minimax regrets as follows $R_\Theta^* \geqslant R_S^* \geqslant R_M^*$. The next result shows that the first inequality is indeed an equality. That is, even if nature is able to select consumers' types from $\mathcal{D}_\Theta$, the seller can achieve the same minimax regret as if nature was restricted to select types from the set of strategic consumers $\mathcal{D}_S$.

Proposition 5. *Suppose nature can select the buyer's type from the set $\mathcal{D}_\Theta$. Then, an optimal pricing strategy for the seller is to select the vector $p_\Theta^* = p_S^*$ identified in Theorem 3. As a result, $R_\Theta^* = R_S^*$.*

Suppose now that instead of a single buyer, the market is populated by $C \geqslant 1$ consumers and nature is capable of selecting the proportion of myopic and strategic consumers. Let us denote by R_Θ^{*C}, the seller's minimax regret. The following result extends Proposition 1 to this setting.

Proposition 6. *The regret R_Θ^{*C} is linear in C, i.e., $R_\Theta^{*C} = CR_\Theta^*$, where R_Θ^* is the optimal regret with a single customer. In addition, any optimal pricing strategy for the single customer case is also optimal for any $C \in \mathbb{N}$. We conclude from Proposition 5 that p_S^* (an optimal price vector when there is a single strategic consumer in the market) is optimal.*

In summary, the previous result reveals that if the seller is uncertain about the types of consumers in the market then her best strategy is to use the same pricing policy that she would use if she knew that all consumers were strategic.

7. Time-Dependent Valuations

The robust formulation presented in the previous sections assumes that customers' valuations are independent of time in the sense that nature is allowed to select customers' types from the rectangle $\mathcal{D} = [\underline{v}, \bar{v}] \times [0, T]$. In many practical situations, however, we can expect some degree of correlation between a customer's arrival time and his valuation. To capture this correlation between arrival times and valuations, in this section, we consider the case in which nature is restricted to select customer's types from a set $\hat{\mathcal{D}} \subseteq \mathcal{D}$. In particular, we will assume that there exist two continuous functions $\underline{v}_t$ and $\bar{v}_t$ with $\underline{v} \leqslant \underline{v}_t \leqslant \bar{v}_t \leqslant \bar{v}$ and such that

$$\hat{\mathcal{D}} := \{(v, \tau) \in \mathcal{D}\colon 0 \leqslant \tau \leqslant T \text{ and } \underline{v}_\tau \leqslant v \leqslant \bar{v}_\tau\}. \tag{16}$$

In other words, we will consider sets $\hat{\mathcal{D}}$ for which customers' valuations belong to an interval that can change over the selling season. For example, customers arriving at the begin of the season could have higher valuations for the product than late-comers (e.g., as in the fashion or high-tech industries where products exhibit high degree of obsolescence). To model this situation we could consider a set $\hat{\mathcal{D}}$ as in Equation (16) for which the upper level of the range of valuations, $\bar{v}_t$, is a decreasing function of time.

Interestingly, in what follows we show that under some mild conditions on the set $\hat{\mathcal{D}}$ this generalization does not change the structure of our proposed solutions. Let us first consider the case of myopic customers. In this case, the geometric argument that we used in Section 4 to decompose the seller's regret into valuation and delay regrets still holds. As a result, we can import the same solution method to derive the following result.

Proposition 7. *Suppose customers are myopic and nature is restricted to choose types from a set $\hat{\mathcal{D}}$ as in Equation (16) such that $\mathcal{D}_0 \subseteq \hat{\mathcal{D}}$. Then, a decreasing function $p \in \mathcal{P}$ has a worst-case regret that is bounded above by R if and only if it satisfies the conditions $p_T \leqslant R \vee \underline{v}$ and*

$$\underline{\hat{p}}_t(R) := \{\bar{v}_t - e^{rt} R\} \vee \underline{v}_t \leqslant p_t \leqslant \left\{\underline{v}_t \vee \frac{R}{1 - e^{-rt}}\right\} \wedge \bar{v}_t =: \hat{\bar{p}}_t(R) \ \text{for all } t \in [0, T].$$

Furthermore, the minimum worst-case regret R^* is the minimum value of R for which the conditions (i) *$\underline{\hat{p}}_t(R) \leqslant \hat{\bar{p}}_t(R)$ for all $t \in [0, T]$ and* (ii) *$\underline{\hat{p}}_T(R) \leqslant R \vee \underline{v}$ are satisfied. Any decreasing price path p_t that satisfies $\underline{\hat{p}}_t(R^*) \leqslant p_t \leqslant \hat{\bar{p}}_t(R^*)$ and $\underline{\hat{p}}_T(R) \leqslant R \vee \underline{v}$ for all $t \in [0, T]$ is optimal.*

Let us turn now to the case of strategic consumers.

Proposition 8. *Suppose customers are strategic and nature is restricted to choose types from a set $\hat{\mathcal{D}}$ as in Equation (16) such that $\mathcal{D}_0 \subseteq \hat{\mathcal{D}}$. Then, the optimal solution identified in Theorem 3 is still optimal.*

It is worth noticing that this result generalizes to an arbitrary set $\hat{\mathcal{D}}$—not necessary of the form in Equation (16)—as long as $\mathcal{D}_0 \subseteq \hat{\mathcal{D}}$. Intuitively, this result highlights a distinctive feature of the case with strategic customers, namely that an optimal price path is one that protects the seller against the event of a consumer arriving at the very beginning of the selling season, i.e., a customer with type in $\mathcal{D}_0$. Hence, as long as the set $\hat{\mathcal{D}}$ includes $\mathcal{D}_0$, the solution in Theorem 3 is optimal. The condition $\mathcal{D}_0 \subseteq \hat{\mathcal{D}}$ seems reasonable in many practical settings, specially in those in which consumers' valuations tend to decrease over time.

8. Concluding Remarks

In this paper we have investigated the intertemporal pricing problem faced by a monopolist selling to myopic and/or strategic customers over a finite selling horizon. A distinctive feature of our work is that we have endowed the seller with a *minimal* amount of information about market demand. In particular, we have assumed that the seller only knows the range of customers' valuations and has no prior belief on how these valuations are distributed on this range or how customers arrive over the selling season. In this distribution-free environment, we tackle the seller's pricing problem using a robust minimax formulation in which the seller selects a price path that minimizes her worst-case regret, that is, the difference between her payoff under full demand information and her realized payoff.

With myopic customers, we show that the seller's regret admits a decomposition in terms of *delay* and *valuation* regrets. Delay regret occurs when a low-valuation customer arrives early in the selling season and must wait for the price to drop below his valuation level before he can purchase the product. On the other hand, valuation regret occurs when a high-valuation customer buys the product immediately upon arrival. These two types of regrets act as opposing forces that the seller must balance at optimality. Using a geometric argument we show how to solve this trade-off and how to characterize an optimal pricing strategy. Furthermore, we show that an optimal price path is not unique and we provide (point-wise) upper and lower bounds for the set of optimal price paths. We also show that a representative price strategy consists of three distinct phases: an initial markup, a markdown, and a final clearance phase. Our analysis also reveals that despite of the multiplicity of optimal price paths they all coincide at a particular critical time. In other words, while the seller has flexibility to set prices in early and late stages of the selling season, she is restricted in her pricing decision around this particular point in time. Another interesting feature when selling to myopic customers is that the optimal length of the selling season is uniformly bounded, independently of the range of customers' valuations.

With strategic customers, the geometric argument that we use to solve the myopic case is less effective since strategic customers can simultaneously generate delay and valuation regrets. As a result, we use a different approach to solve the seller's problem. In particular, we view the seller's pricing problem as a screening problem and adapt the machinery of mechanism design to fit our robust minimax formulation. Our analysis shows that—depending on the parameters of the model—there exist three types of pricing regimes. First, when the seller faces significant demand uncertainty (i.e., when the range of customers' valuations is wide), it is optimal to set a price path that leaves low-valuation consumers out of the market. On the other hand, for moderate (or intermediate) levels of demand uncertainty, the seller selects a price path that serves every customer type. However, low-valuation consumers are pooled together into a single class that buys the product exactly at the end of the selling season. Finally, when demand uncertainty is low, the seller sets a price path that not only serves every customer type but also separates them based on their purchasing time, i.e., the pricing strategy behaves as a prefect screening device. In general, compared to the myopic case, an optimal price path is "flatter" in the sense that the optimal price path relies less on markdowns. Also, it is no longer true that there is a uniform upper bound on the length of the selling season. That is, with strategic customers, it is possible that the seller wants to extend the selling season indefinitely.

There are many interesting questions that we leave open for future research. Intertemporal pricing under minimax regret without commitment is a topic that we do not address in our paper. Another such question is whether a randomized pricing policy could be used to further reduce the seller's regret. A further interesting follow-up project would be to incorporate inventory into our model. Inventory considerations are an important factor in practice, but, as Correa et al. (2016) show, they generally lead to a multiplicity of equilibria in settings with strategic customers. Allowing consumers to exit the market without purchasing before the end of the sales horizon, for either endogenous or exogenous reasons, would also be an interesting extension to consider. Extending our methodology to more complex uncertainty sets than intervals over the customers' valuations or to consider different demand classes with different uncertainty sets would also be valuable extensions, Section 7 above constitutes a first step in this direction.

Acknowledgments

The authors thank Chung-Piaw Teo, the associate editor and the two referees for their thoughtful recommendations. The paper benefited greatly from their suggestions.

Appendix A. Proofs

Proof of Proposition 1. Consider the seller's optimal regret R^{*C} in (6). A lower bound of this regret is obtained by restricting nature's optimization to the subset $\mathcal{D}_0^C \subseteq \mathcal{D}^C$ such that $\{v, \tau\} \in \mathcal{D}_0^C$ if and only if $v_1 = v_2 = \cdots = v_C$ and $\tau_1 = \tau_2 = \cdots = \tau_C$, that is, all C customers have identical valuations and arrival times. But if all customers are identical then the seller's regret is equal to C times the optimal regret with a single customer, that is, $R^{*C} \geqslant CR^*$.

On the other hand, any fixed pricing strategy p_t leads to an upper bound for R^{*C} in the sense that

$$R^{*C} \leqslant \sup_{\{v,\tau\}\in\mathcal{D}^C} \left[\Pi_F^C(v,\tau) - \sum_{i=1}^{C} e^{-rd(v_i,\tau_i,p)} p_{d(v_i,\tau_i,p)} \right].$$

Furthermore, it is not hard to see that

$$\Pi_F^C(v,\tau) \leqslant \sum_{i=1}^{C} e^{-r\tau_i} v_i$$

since the right-hand side is the payoff that a clairvoyant can get under perfect price discrimination (i.e., charging a different price to every customer). It follows that

$$\begin{aligned} R^{*C} &\leqslant \sup_{\{v,\tau\}\in\mathcal{D}^C} \left[\sum_{i=1}^{C} e^{-r\tau_i} v_i - e^{-rd(v_i,\tau_i,p)} p_{d(v_i,\tau_i,p)} \right] \\ &= \sum_{i=1}^{C} \sup_{\{v_i,\tau_i\}\in\mathcal{D}} \left[e^{-r\tau_i} v_i - e^{-rd(v_i,\tau_i,p)} p_{d(v_i,\tau_i,p)} \right], \end{aligned}$$

hence the maximization above decouples for every customer. But if the seller selects a pricing vector p_t^*, which is an optimal pricing strategy for the case of single customer, then each term in the summation on the right equals R^*, that is, $R^{*C} \leqslant CR^*$. We conclude that $R^{*C} = CR^*$. □

Proof of Lemma 1. Consider a price path p_t such that $p_t \geqslant \underline{p}_t(R)$ for all $t \in [0, T]$. The worst-case valuation regret from p_t satisfies the following inequality

$$\begin{aligned} \max_t e^{-rt}(\bar{v} - p_t) &\leqslant \max_t e^{-rt}(\bar{v} - \underline{p}_t(R)) \\ &\leqslant \max_t e^{-rt}(\bar{v} - (\bar{v} - e^{rt}R)) = R. \quad \square \end{aligned}$$

Proof of Lemma 2. Consider an arbitrary price path $p \in \mathcal{P}$ (i.e., a continuous function from $[0, T]$ to $[\underline{v}, \bar{v}]$) and let

$$\begin{aligned} \mathcal{R}_{\mathrm{M}}(p) &= \sup_{(v,\tau)\in\mathcal{D}} \mathcal{R}_{\mathrm{M}}(v,\tau,p) \\ &= \sup_{(v,\tau)\in\mathcal{D}} \{ e^{-r\tau} v - e^{-rd_M(v,\tau,p)} p_{d_M(v,\tau,p)} \}, \end{aligned}$$

be the corresponding seller's worst-case regret. In what follows, we show that there always exists a (weakly decreasing) price path $\hat{p}$ such that $\mathcal{R}_{\mathrm{M}}(\hat{p}) \leqslant \mathcal{R}_{\mathrm{M}}(p)$.

Indeed, for a given $p \in \mathcal{P}$, we let $\hat{p}$ be the running minimum price path induced by p that is given by $\hat{p}_t := \min\{p_\tau : \tau \in [0, t]\}$ for $t \in [0, T]$ (the "min" is well-defined by the continuity of p and Weierstrass Theorem). By construction, $\hat{p}$ is (weakly) decreasing and continuous (since p is continuous). To show that $\mathcal{R}_{\mathrm{M}}(p) \geqslant \mathcal{R}_{\mathrm{M}}(\hat{p})$ let us partition the customer type's space $\mathcal{D}$ into the following three subsets $\mathcal{D}_1 := \{(v,\tau) \in \mathcal{D}: p_\tau \leqslant v\}$, $\mathcal{D}_2 := \{(v,\tau) \in \mathcal{D}: \hat{p}_\tau \leqslant v < p_\tau\}$ and $\mathcal{D}_3 := \{(v,\tau) \in \mathcal{D}: \hat{p}_\tau > v\}$ and define $R_i(p) := \sup_{(v,\tau)\in\mathcal{D}_i} \mathcal{R}_{\mathrm{M}}(v,\tau,p)$. In what follows, we show that $R_i(p) = R_i(\hat{p})$ for $i = 1, 3$ and that $R_2(\hat{p}) \leqslant R_1(\hat{p})$. As a result,

$$\begin{aligned} \mathcal{R}_{\mathrm{M}}(p) &= \max_{i=1,2,3}\{R_i(p)\} \geqslant \max_{i=1,3}\{R_i(p)\} = \max_{i=1,3}\{R_i(\hat{p})\} \\ &= \max_{i=1,2,3}\{R_i(\hat{p})\} = \mathcal{R}_{\mathrm{M}}(\hat{p}). \end{aligned}$$

—Let $(v,\tau) \in \mathcal{D}_1$. By the definition of $\hat{p}$, there exists a $t \leqslant \tau$ such that $\hat{p}(t) = p(t) \leqslant p(\tau) \leqslant v$. Since consumers are myopic, it follows that $\mathcal{R}_{\mathrm{M}}(v,t,p) = \mathcal{R}_{\mathrm{M}}(v,t,\hat{p}) = e^{-rt}(v - p_t) \geqslant \mathcal{R}_{\mathrm{M}}(v,\tau,p)$. Since $(v,t) \in \mathcal{D}_1$, we conclude that $R_1(p) = R_1(\hat{p})$.

—Let $(v,\tau) \in \mathcal{D}_3$. In this case, it is easy to see that $d_M(v,\tau,p) = d_M(v,\tau,\hat{p})$ (possibly ∞) and $p(d_M(v,\tau,p)) = \hat{p}(d_M(v,\tau,\hat{p}))$. It follows that $\mathcal{R}_{\mathrm{M}}(v,\tau,p) = \mathcal{R}_{\mathrm{M}}(v,\tau,\hat{p})$. We conclude that $R_3(p) = R_3(\hat{p})$.

—Let $(v,\tau) \in \mathcal{D}_2$. Since $p_\tau > \hat{p}_\tau$ there exists a $t \in [0,\tau)$ such that $\hat{p}(t) = p(t) \leqslant v$. It follows that $\mathcal{R}_{\mathrm{M}}(v,\tau,\hat{p}) = e^{-r\tau}(v - \hat{p}_\tau) \leqslant e^{-rt}(v - \hat{p}_t) = \mathcal{R}_{\mathrm{M}}(v,t,\hat{p})$. Furthermore, $(v,t) \in \mathcal{D}_1$. Hence, $R_2(\hat{p}) \leqslant R_1(\hat{p})$.

It follows from the three cases above that $\mathcal{R}_{\mathrm{M}}(p) \geqslant \mathcal{R}_{\mathrm{M}}(\hat{p})$. Hence, without loss of optimality the seller can restrict attention to (weakly) decreasing price functions. □

Proof of Lemma 3. Let $p \in \mathcal{P}$ be a decreasing function. For any customer type (v,τ) with $v < p_\tau$, the delay regret he generates is equal to $(e^{-r\tau} - e^{-rd_M(v,\tau,p)})v$. We distinguish two cases:

—*Case* 1: $d_M(v,\tau,p) \leqslant T$. In this case the customer buys the product at a future time $d_M(v,\tau,p)$ such that $p_{d_M(v,\tau,p)} = v$ and he generates a delay regret $(e^{-r\tau} - e^{-rd_M(v,\tau,p)}) p_{d_M(v,\tau,p)}$ that is bounded above by $(1 - e^{-rd_M(v,\tau,p)}) p_{d_M(v,\tau,p)}$. But, by hypothesis, the price path p satisfies $p_t \leqslant R/(1 - e^{-rt}) \vee \underline{v}$ for all $t \in [0, T]$. Furthermore, from the monotonicity of p and the fact that $v < p_\tau$ it is not hard to see that $R/(1 - e^{-rd_M(v,\tau,p)}) \geqslant \underline{v}$. As a result, $(1 - e^{-rd_M(v,\tau,p)}) p_{d_M(v,\tau,p)} \leqslant R$, that is, the delay regret generated by the customer is less than or equal to R.

—*Case* 2: $d_M(v,\tau,p) = \infty$. In this case the customer is priced out of the market. Since p is a decreasing price path this means that $v < p_T$. But, by hypothesis, $p_T \leqslant R \vee \underline{v}$, hence this case can only occur if $R > \underline{v}$. So, assuming $R > \underline{v}$, the delay regret in this case is equal to $e^{-r\tau}v$, which is bounded above by v, which together with the inequalities $v < p_T < R$ imply that the delay regret generated by the customer is less than R. □

Proof of Proposition 2. This result follows immediately from Lemmas 1–3. □

Proof of Theorem 1. Proposition 2 characterizes the set of minimax regret price paths. Region A_1 is the set of problem parameters where the solution is given by the price paths $\underline{p}_t(R)$ and $\bar{p}_t(R)$ intersecting tangentially as described by Equations (7) and (8). For this solution to be feasible, the ratio $u = \underline{v}/\bar{v}$ must satisfy $u \leqslant 1/2$ to ensure the middle equation of (8) is feasible. The solution must also satisfy $\underline{p}_t(R) = \bar{v} - e^{rT}R \leqslant R \vee \underline{v}$, where $R = \bar{v}/4$ from the right-most equation in (8). Therefore, A_1 is characterized by

$$u \leqslant \frac{1}{2} \quad \text{and} \quad T \geqslant \min\left\{ \frac{\ln 3}{r}, \frac{\ln(4(1-u))}{r} \right\}.$$

In regions A_2, A_3, and A_4, the boundary conditions (u and/or T) become binding and, thus, play a role in determining the minimax regret. If $u \geqslant \frac{1}{2}$ and the selling horizon is sufficiently long, the binding constraint becomes $\underline{p}_{t^*}(R_{\mathrm{M}}^*) = \bar{p}_{t^*}(R_{\mathrm{M}}^*) = \underline{v}$ for some

t^*, as given by Equation (9). We can solve for R^*_{M} and t^* to get $R^*_{\mathrm{M}} = u(1-u)\bar{v}$ and $t^* = (1/r)\ln(1/u)$. Thus, we need $T \geqslant (1/r)\ln(1/u)$ to guarantee to obtain $t^* \leqslant T$. The pair of conditions $u \geqslant \frac{1}{2}$ and $T \geqslant (1/r)\ln(1/u)$ lead to the definition of A_2.

In the regime where $(u, T) \notin A_1 \cup A_2$, the binding constraint becomes $\underline{p}_T(R^*_{\mathrm{M}}) = \max\{R^*_{\mathrm{M}}, \underline{v}\}$. One can solve the equation

$$\bar{v} - e^{rt} R^*_{\mathrm{M}} = \max\{R^*_{\mathrm{M}}, \underline{v}\} \tag{A1}$$

and get the following results: if $T \leqslant \min\{\ln(3)/r, \ln(1/u-1)/r\}$, then the value that maximizes the right-hand side of Equation (A1) is R^*_{M}, generating the solution associated with region A_3; otherwise, the right-hand side of Equation (A1) is maximized by $\underline{v}$, producing the solution for region A_4. □

Proof of Theorem 2. Consider an arbitrary price path p with regret $\mathcal{R}_{\mathrm{M}}(p)$ and let us suppose that $\mathcal{R}_{\mathrm{M}}(p) < R^*_{\mathrm{M}}$. In what follows, we will show that the previous inequality leads to a contradiction. For this, consider the critical time t^* (see the discussion that precedes Theorem 1). It is not hard to see that the lower bound $\underline{p}_t(R^*_{\mathrm{M}})$ is strictly decreasing in $[0, t^*]$. It follows, under the assumption that $\mathcal{R}_{\mathrm{M}}(p) < R^*_{\mathrm{M}}$, that we must have that $p_t > \underline{p}_t(R^*_{\mathrm{M}})$ for all $t \in [0, t^*]$. Otherwise, we would be able to find a customer type that generates a valuation regret strictly greater than R^*_{M} under the price path p, and this would violate our hypothesis $\mathcal{R}_{\mathrm{M}}(p) < R^*_{\mathrm{M}}$.

Now, consider a customer with type $(v, \tau) = (\underline{p}_{t^*}(R^*_{\mathrm{M}}), 0)$. Since $p_t > \underline{p}_t(R^*_{\mathrm{M}})$ for all $t \in [0, t^*]$ and $\underline{p}_t(R^*_{\mathrm{M}})$ is strictly decreasing in this interval, it follows that the purchasing time of this customer under the price path p satisfies $d_M(\underline{p}_{t^*}(R^*_{\mathrm{M}}), 0, p) > t^*$. In addition, we trivially must have that $\mathcal{R}_{\mathrm{M}}(p)$ is greater than or equal to the delay regret generated by this consumer. It follows that

$$\begin{aligned}\mathcal{R}_{\mathrm{M}}(p) &\geqslant (1-\exp(-r d_M(\underline{p}_{t^*}(R^*_{\mathrm{M}}), 0, p)))\,\underline{p}_{t^*}(R^*_{\mathrm{M}})\\ &> (1-\exp(-r\,t^*))\,\underline{p}_{t^*}(R^*_{\mathrm{M}}).\end{aligned} \tag{A2}$$

Let us now evaluate this expression for each of the four parameter regimes (regions A_1 through A_4) identified in Theorem 1:

—In region A_1, we have that $R^*_{\mathrm{M}} = \bar{v}/4$, $t^* = \ln(2)/r$ and $\underline{p}_{t^*}(R^*_{\mathrm{M}}) = \bar{v}/2$. As a result, we have that $(1-\exp(-rt^*)\cdot \underline{p}_{t^*}(R^*_{\mathrm{M}}) = R^*_{\mathrm{M}}$, which together with (A2), contradicts the hypothesis $\mathcal{R}_{\mathrm{M}}(p) < R^*_{\mathrm{M}}$.

—In region A_2, we have that $R^*_{\mathrm{M}} = \underline{v}(1-\underline{v}/\bar{v})$, $t^* = \ln(\bar{v}/\underline{v})/r$ and $\underline{p}_{t^*}(R^*_{\mathrm{M}}) = \underline{v}$. In this case, we get that $(1-\exp(-rt^*)\,\underline{p}_{t^*}(R^*_{\mathrm{M}}) = R^*_{\mathrm{M}}$, which again contradicts $\mathcal{R}_{\mathrm{M}}(p) < R^*_{\mathrm{M}}$.

—In region A_3, we have that $t^* = T$ and so $d_M(\underline{p}_{t^*}(R^*_{\mathrm{M}}), 0, p) = \infty$, i.e., the customer is price out of the market. In addition, $\underline{p}_{t^*}(R^*_{\mathrm{M}}) = R^*_{\mathrm{M}}$. It follows from (A2) that $\mathcal{R}_{\mathrm{M}}(p) \geqslant R^*_{\mathrm{M}}$, which contradicts our hypothesis.

—In region A_4, we have again that $t^* = T$ and $d_M(\underline{p}_{t^*}(R^*_{\mathrm{M}}), 0, p) = \infty$. In this case, we have that $\underline{p}_{t^*}(R^*_{\mathrm{M}}) = \underline{v} \geqslant R^*_{\mathrm{M}}$. Once again, it follows from (A2) that $\mathcal{R}_{\mathrm{M}}(p) \geqslant R^*_{\mathrm{M}}$, which contradicts our hypothesis and the proof is complete. □

Proof of Lemma 4. See Part V in Rockafellar (1997). □

Proof of Proposition 3. Note first that the proposed solution θ^* is feasible in the sense that $\theta^* \in \Theta(\hat{v})$. Let us define $\tilde{v} := \max\{\hat{v}, \bar{v}\exp(\theta_0 - 1)\}$. The proposed solution can be written as:

$$\theta^*(x) = \begin{cases} \theta_0 & \text{if } \hat{v} \leqslant x < \tilde{v} \\ 1 - \ln(\bar{v}) + \ln(x) & \text{if } \tilde{v} \leqslant x \leqslant \bar{v}. \end{cases}$$

It follows that the regret of a customer with valuation $v \in [\hat{v}, \bar{v}]$ under θ^* is equal to

$$\int_{\hat{v}}^{v} \theta^*(x)\,\mathrm{d}x + v(1-\theta^*(v)) = \min\{v, \tilde{v}\} + \hat{v}(\ln(\bar{v}) - \ln(\tilde{v}) - 1).$$

An important property of this regret is that it is increasing in v and it is constant in the interval $[\tilde{v}, \bar{v}]$. Hence, the seller's worst-case regret under θ^* is achieved in the entire interval $[\tilde{v}, \bar{v}]$. We will use this property to prove the optimality of θ^* using a variational arguments. To this end, suppose that θ^* is not optimal. Hence, there is a $\theta \in \Theta(\hat{v})$ that generates a strictly better regret than the one generated by θ^* and let us define $\epsilon(x) := \theta(x) - \theta^*(x)$ for all $x \in [\hat{v}, \bar{v}]$. By feasibility, we need $\epsilon(x)$ to be right-continuous and such that $\theta(x) \in [\theta_0, 1]$ for all $x \in [\hat{v}, \bar{v}]$. For instance, since $\theta^*(\bar{v}) = 1$, we must have $\epsilon(\bar{v}) \leqslant 0$. Similarly, since $\theta^*(x) = \theta_0$ for all $x \in [\hat{v}, \tilde{v})$, we must have $\epsilon(x) \geqslant 0$ in this interval.

Now, the assumption that θ generates a strictly better regret than θ^* and the fact that seller's worst-case regret under θ^* is achieved in the entire interval $[\tilde{v}, \bar{v}]$ imply that $\epsilon(x)$ must satisfy

$$\int_{\hat{v}}^{v} \epsilon(x)\,\mathrm{d}x - v\epsilon(v) < 0 \quad \text{for all } v \in [\tilde{v}, \bar{v}]. \tag{A3}$$

Since $\epsilon(x) \geqslant 0$ in the interval $x \in [\hat{v}, \tilde{v})$, we must have $\epsilon(\tilde{v}) > 0$, since $\epsilon(\tilde{v}) > 1/\tilde{v}\int_{\hat{v}}^{\tilde{v}} \epsilon(x)\,\mathrm{d}x \geqslant 0$.

Let us show that $\epsilon(x) \geqslant 0$ for all $x \in [\tilde{v}, \bar{v})$. Suppose, by contradiction, that this is not the case. Then, there exist a $\delta < 0$ and $y = \inf\{v \in (\tilde{v}, \bar{v})\colon \epsilon(v) \leqslant \delta\}$. Also, since $\epsilon(x)$ is right-continuous and $\epsilon(\tilde{v}) > 0$, we must have $y > \tilde{v}$ and

$$\int_{\hat{v}}^{y} \epsilon(x)\,\mathrm{d}x > \delta(y - \hat{v}) \geqslant \epsilon(y)(y - \hat{v}) \geqslant y\epsilon(y).$$

This contradicts (A3) and we must have $\epsilon(x) \geqslant 0$ for all $x \in [\hat{v}, \bar{v})$. Combining the non-negativity and right-continuity of $\epsilon(x)$ and the facts that $\epsilon(\tilde{v}) > 0$ and $\epsilon(\bar{v}) \leqslant 0$, we get

$$\int_{\hat{v}}^{\bar{v}} \epsilon(x)\,\mathrm{d}x - \bar{v}\epsilon(\bar{v}) > 0.$$

This again contradicts (A3), which proves the optimality of θ^*. □

Proof of Theorem 3. From Proposition 3 we have that for an arbitrary $\hat{v} \in [\underline{v}, \bar{v}]$

$$R_{\hat{v}} = \hat{v}\mathbb{1}(\hat{v} > \underline{v}) \vee \tilde{v} + \hat{v}\,(\ln(\bar{v}) - \ln(\tilde{v}) - 1), \qquad \text{where } \tilde{v} := \max\{\hat{v}, \bar{v}\exp(e^{-rT} - 1)\}.$$

Hence, it is a matter of relatively straightforward calculations to show that

$$\hat{v}^* := \underset{\underline{v} \leqslant \hat{v} \leqslant \bar{v}}{\arg\max}\{R_{\hat{v}}\} = \max\left\{\underline{v}, \frac{\bar{v}\exp(e^{-rT} - 1)}{1 + e^{-rT}}\right\} \quad \text{and}$$

$$R^*_0 = R_{\hat{v}^*} = \tilde{v}^* + \hat{v}^*\,(\ln(\bar{v}) - \ln(\tilde{v}^*) - 1), \qquad \text{where } \tilde{v}^* := \max\{\hat{v}^*, \bar{v}\exp(e^{-rT} - 1)\}.$$

Recall that $\hat{v}^*$ represents is the threshold that separates consumers' valuations into those that buy and those that do not buy the product. From Proposition 3, we have that the corresponding optimal purchasing strategy of those customers that buy the product is given by $\theta^*(v) = \theta_0 \vee 1 - \ln(\bar{v}) + \ln(v)$ for $v \in [\hat{v}^*, \bar{v}]$. We can reverse the change of variable, $\theta = e^{-rt}$, to rewrite this consumers' purchasing strategy in term of purchasing time. Indeed, if we let

$d^*(v)$ be the time at which a consumer with valuation $v \in [\hat{v}^*, \bar{v}]$ buys the product, we have that

$$\begin{aligned} d^*(v) &= -\frac{1}{r}\ln(\theta^*(v)) \\ &= -\frac{1}{r}\ln(1-\ln(\bar{v})+\ln(v)) \wedge T \quad \text{for all } v \in [\hat{v}^*, \bar{v}]. \end{aligned}$$

The last step in the proof is a *verification* step. Essentially, we need to verify that the proposed pricing strategy implements the consumers' optimal strategy $\theta^*(v)$ (or equivalently $d^*(v)$). The optimality of the pricing strategy will then follow from the optimality of $\theta^*(v)$ in Proposition 3.

Let us recall from the statement of the theorem that our proposed pricing strategy is equal to

$$p_S^*(t) = \begin{cases} e^{rt}(\bar{v}\exp(e^{-rt}-1) - R_0^*) & \text{for all } t \in [0, d^*(\hat{v}^*)] \\ \underline{v} & \text{for all } t \in (d^*(\hat{v}^*), T]. \end{cases}$$

To verify that $p_S^*(t)$ implements $d^*(v)$ we need to show that the following conditions hold

(a) For all $v \in [\underline{v}, \hat{v}^*)$: $\max_{0 \leqslant t \leqslant T}\{e^{-rt}[v - p_S^*(t)]\} < 0$,

(b) For all $v \in [\hat{v}^*, \bar{v}]$: $d^*(v) = \max\{\arg\max_{0\leqslant t\leqslant T}\{e^{-rt}[v - p_S^*(t)]\}\}$ and

(c) $\hat{v}^* = p_S^*(d^*(\hat{v}^*))$.

Condition (a) guarantees that customers with valuation below $\hat{v}^*$ leave the market without buying. Condition (b) guarantees that the price path p_t implements $d^*(v)$ for $v \geqslant \hat{v}^*$. Finally, condition (c) ensures that the lowest type that buys the product $v = \hat{v}^*$ gets zero utility (recall the definition of $\hat{v}$ in Equation (14)).

To verify the three conditions above we identify three cases depending on the values of $\underline{v}$, $\bar{v}$, and $\theta_0 = e^{-rT}$: (i) $\underline{v} \leqslant \bar{v}\exp(\theta_0 - 1)/(1+\theta_0)$, (ii) $\bar{v}\exp(\theta_0-1)/(1+\theta_0) \leqslant \underline{v} \leqslant \bar{v}\exp(\theta_0 - 1)$, and (iii) $\underline{v} \geqslant \bar{v}\exp(\theta_0 - 1)$. We will verify case (i) and leave the other two to the reader as they follow similar steps. These cases correspond to the regions B_1, B_2 and B_3 in Figure 3 respectively.

For case (i), we have that $\hat{v}^* = \bar{v}\exp(\theta_0 - 1)/(1+\theta_0)$, $\tilde{v}^* = \bar{v}\exp(\theta_0 - 1)$, $R_0^* = \hat{v}^*$ and

$$d^*(v) = \begin{cases} T & \text{for all } v \in [\hat{v}^*, \tilde{v}^*] \\ -\ln(1-\ln(\bar{v})+\ln(v))/r & \text{for all } v \in [\tilde{v}^*, \bar{v}]. \end{cases}$$

As a result, the proposed price path takes the form

$$p_S^*(t) = e^{rt}\left(\bar{v}\exp(e^{-rt}-1) - \bar{v}\frac{\exp(\theta_0 - 1)}{1+\theta_0}\right), \quad \text{for all } t \in [0, T].$$

Hence the consumer's utility maximization problem is equal to

$$\begin{aligned} &\max_{0\leqslant t\leqslant T}\{e^{-rt}[v - p_S^*(t)]\} \\ &\quad = \max_{0\leqslant t\leqslant T}\left\{e^{-rt}v - \bar{v}\exp(e^{-rt}-1) - \bar{v}\frac{\exp(\theta_0-1)}{1+\theta_0}\right\}. \end{aligned}$$

Since the function $e^{-rt}v - \bar{v}\exp(e^{-rt}-1)$ is unimodal (inverted U-shaped) for $t \in [0, T]$, we can use first-order condition to solve the consumer's problem. This FOC is given by

$$re^{-rt}(\bar{v}\exp(e^{-rt}) - v) = 0 \implies t(v) = -\frac{1}{r}\ln(1-\ln(\bar{v})+\ln(v)).$$

It is a matter of simple calculations to show that $t(v) \leqslant T$ if and only if $v \geqslant \tilde{v}^*$ and so for values of v in $[\hat{v}^*, \tilde{v}^*]$ the consumer's utility is maximized at $t = T$. Hence, we conclude that the optimal purchasing time of consumers with valuation in $[\hat{v}^*, \bar{v}]$ is precisely

$$d^*(v) = \begin{cases} T & \text{for all } v \in [\hat{v}^*, \tilde{v}^*] \\ -\ln(1-\ln(\bar{v})+\ln(v))/r & \text{for all } v \in [\tilde{v}^*, \bar{v}] \end{cases}.$$

as required. Finally, it is easy to verify that $p_S^*(d^*(\hat{v}^*)) = \hat{v}^*$ and so customer with valuation equal to $\hat{v}^*$ makes zero utility and all customers with valuation strictly below $\hat{v}^*$ make a negative utility under $p_S^*(t)$ and hence leave the market without purchasing the product. □

Proof of Lemma 5. To alleviate the notation, let us write $p_t = p_S^*(t)$ and $\tilde{p}_t = \tilde{p}_S^*(t)$.

The first step in the proof is to show that the function p_t is continuously differentiable in $[0, T]$. This holds trivially for the case in which $d^*(\hat{v}^*) = T$. On the other hand, if $d^*(\hat{v}^*) < T$, what we need to show is that the left-derivative of p_t at $t = d^*(\hat{v}^*)$ is equal to 0. To see this, note that for all $t \in [0, d^*(\hat{v}^*)]$,

$$\dot{p}_t = r(p_t - R_0^*) - re^{-rt}p_t.$$

Now, the condition $d^*(\hat{v}^*) < T$ implies that $d^*(\hat{v}^*) = -\ln(1-\ln(\bar{v})+\ln(\underline{v}))/r$, $p_{d^*(\hat{v}^*)} = \underline{v}$ and $R_0^* = \underline{v}(\ln(\bar{v}) - \ln(\underline{v}))$. Plugging these values, we get that $\dot{p}_t = 0$ at $t = d^*(\hat{v}^*)$ and so p_t is continuously differentiable in $[0, T]$.

Now, from the fact that p_t is decreasing and continuously differentiable in $[0, T]$, the definition of $\tilde{p}_t$ in (11) implies that

$$\begin{aligned} \tilde{p}_t &= \sup_{s\in(t,T]}\left\{\frac{e^{-rt}p_t - e^{-rs}p_s}{e^{-rt} - e^{-rs}}\right\} \\ &= \sup_{s\in(t,T]}\left\{\frac{\int_t^s re^{-ru}(p_u - \dot{p}_u/r)\,du}{\int_t^s re^{-ru}\,du}\right\} \quad \text{for all } t \in [0, T) \\ \tilde{p}_T &= p_T. \end{aligned}$$

But $p_t - \dot{p}_t/r = R_0^* + e^{-rt}p_t$, which is a decreasing function of t. It follows that the supremum above is attained as $s \downarrow t$. Finally, from L'Hôpital's rule, we get then that $\tilde{p}_t = p_t - \dot{p}_t/r$, which (as we just showed) is decreasing in $[0, T]$. □

Proof of Theorem 4. Since the optimal price path $p_S^*(t)$ was derived assuming that all customer arrive at time $t = 0$ (i.e., customers such as customer "B" in Figure 9), we only need to show that the worst-case regret of a customer with type (v, t) in the upper boundary $\mathcal{D}_{\bar{v}} = \{(\bar{v}, t)\colon t \in [0, T]\}$ is less than or equal to the regret R_0^* for all $t \in [0, T]$. However, this condition follows trivially since the regret generated by customer $(\bar{v}, t)$ if the seller uses price $p_S^*(t)$ is equal to

$$e^{-rt}[\bar{v} - p_S^*(t)] = \begin{cases} R_0^* - \bar{v}(\exp(e^{-rt}-1) - e^{-rt}) \\ \qquad \text{for all } t \in [0, d^*(\hat{v}^*)] \\ e^{-rt}(\bar{v} - \underline{v}) \\ \qquad \text{for all } t \in (d^*(\hat{v}^*), T]. \end{cases}$$

It is not hard to see that $\exp(e^{-rt}-1) - e^{-rt} \geqslant 0$ for all $t \geqslant 0$ and so $e^{-rt}[\bar{v} - p_S^*(t)] \leqslant R_0^*$ for all $t \in [0, d^*(\hat{v}^*)]$. On the other hand, the price path $p_S^*(t)$ is continuous at $t = d^*(\hat{v}^*)$ and $e^{-rt}[\bar{v} - p_S^*(t)] = e^{-rt}(\bar{v} - \underline{v}) \leqslant e^{-rd^*(\hat{v}^*)}(\bar{v} - \underline{v}) \leqslant R_0^*$ for all $t \in (d^*(\hat{v}^*), T]$. □

Proof of Proposition 4. For a given price path $p \in \mathcal{P}$, let us partition the space of customer types $\mathcal{D}$ into $\mathcal{D}_b$ and $\mathcal{D}_l$, where

$\mathcal{D}_b$ represents the set of types that buy the product and $\mathcal{D}_l$ the set of types that are priced out of the market under p. The sets $\mathcal{D}_b$ and $\mathcal{D}_l$ are independent of whether the customers are myopic or strategic. Hence, the regret generated by a customer with type $(v,\tau)\in\mathcal{D}_l$ under p is equal to $\mathcal{R}_\theta(v,\tau,p)=e^{-r\tau}v$ for any $\theta\in\{M,S\}$.

For a given consumer type $(v,\tau)\in\mathcal{D}_b$ and price path p, it is straightforward to see that a strategic consumer will purchase later than a myopic consumer, i.e., $d_S(v,\tau,p)\geqslant d_M(v,\tau,p)$. We now show that a strategic consumer also buys at a lower price, i.e., $p_{d_S(v,\tau,p)}\leqslant p_{d_M(v,\tau,p)}$. Suppose $p_{d_S(v,\tau,p)}>p_{d_M(v,\tau,p)}$. Then the utility for the strategic customer at time $d_S(v,\tau,p)$ is $e^{-rd_S(v,\tau,p)}(v-p_{d_S(v,\tau,p)})$, which is strictly less than $e^{-rd_M(v,\tau,p)}(v-p_{d_M(v,\tau,p)})$. This implies that the strategic customer could have gained a higher utility if he purchased the product at time $d_M(v,\tau,p)$, which contradicts to the fact that $d_S(v,\tau,p)$ is the time when the utility of strategic customer is maximized.

Therefore, the seller can always extract a higher surplus from a myopic consumer than from a strategic one, i.e., $e^{-rd_M(v,\tau,p)}p_{d_M(v,\tau,p)}\geqslant e^{-rd_S(v,\tau,p)}p_{d_S(v,\tau,p)}$. Thus, the regret from selling to strategic customer is higher, or $\mathcal{R}_S(v,\tau,p)\geqslant\mathcal{R}_M(v,\tau,p)$ for all v,τ and p. Since, $R^*_\theta=\inf_{p\in\mathcal{P}}\sup_{(v,\tau)\in\mathcal{D}}\mathcal{R}_\theta(v,\tau,p)$, we obtain that $\mathcal{R}^*_S\geqslant\mathcal{R}^*_M$. □

Proof of Proposition 5. From the proof of Proposition 4, we know that for under any given price path, a strategic customer yields higher regret than a myopic customer with the same type. We conclude that for any $p\in\mathcal{P}$ and $(v,\tau)\in\mathcal{D}$, nature will select a strategic buyer $\theta=S$, and the result follows. □

Proof of Proposition 6. The proof follows the same steps as the proof of Proposition 1 and is therefore omitted. □

Proof of Proposition 7. The proof follows the same steps as the proof of Proposition 2 and is therefore omitted. □

Proof of Proposition 8. From the identity $\mathcal{D}_0\subseteq\hat{\mathcal{D}}\subseteq\mathcal{D}$ we have that

$$R^*_0=\inf_{p\in\mathcal{P}}\sup_{(v,\tau)\in\mathcal{D}_0}\mathcal{R}_S(v,\tau,p)\leqslant\inf_{p\in\mathcal{P}}\sup_{(v,\tau)\in\hat{\mathcal{D}}}\mathcal{R}_S(v,\tau,p)\leqslant\inf_{p\in\mathcal{P}}\sup_{(v,\tau)\in\mathcal{D}}\mathcal{R}_S(v,\tau,p)=R^*_S.$$

However, from Theorem 4 we have that $R^*_0=R^*_S$. □

Appendix B. Modified Gallego and van Ryzin (1994) Model

In this appendix, we compare the optimal price path resulting from our robust minimax approach to one of the most well-cited models of dynamic pricing within the OR/MS literature. It is probably fair to say that one of the most influential papers on this topic is the *Management Science* paper by Gallego and van Ryzin (1994), who consider a monopolist selling a finite inventory of a product over a finite selling horizon. At the core of Gallego and van Ryzin's pricing model is a dynamic programming formulation that produces the following optimality condition for an optimal price $p^*(t,n)$ as function of time-to-go t and available inventory n:

$$p^*(t,n)=\arg\max\{\lambda(p)[p-(J(t,n)-J(t,n-1)]\},\qquad\text{(B1)}$$

where $\lambda(p)$ is the demand intensity (a non-negative and decreasing function of p) and $J(t,n)$ is the value function, that is, the seller's expected optimal revenues as a function of the state (t,n). (In Gallego and van Ryzin (1994), t denotes time-to-go as opposed to calendar time).

One of the key results in Gallego and van Ryzin (1994) is Theorem 1 that shows that the value function is strictly increasing and strictly concave in both t and n. Using this result and Equation (B1) above, one can show that for many commonly used demand functions, the optimal price path is decreasing concave in calendar time. For example, in Section 2.3 in Gallego and van Ryzin (1994), the authors consider the special case of an exponential demand function and show that $p^*(t,n)=1+J(t,n)-J(t,n-1)$, where

$$J(t,n)=\log\left(\sum_{i=0}^{n}\frac{(\lambda^*(T-t))^i}{i!}\right),$$

$$\text{where }\lambda^*:=\lambda(p^*)\quad\text{and}\quad p^*:=\arg\max\{p\lambda(p)\}.$$

Figure B.1. (Color online) Optimal price paths for the Gallego and van Ryzin model when consumers have exponential willingness to pay.

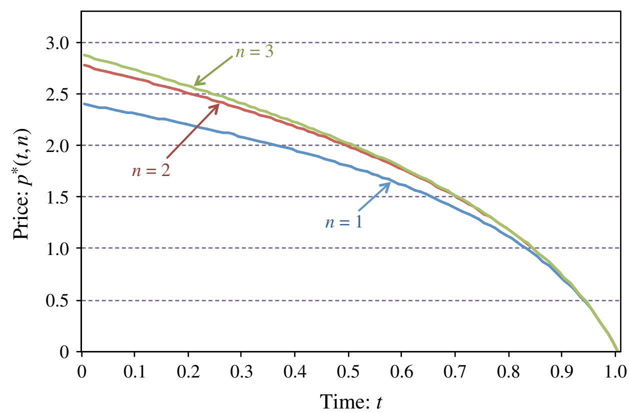

Figure B.1 depicts $p^*(t,n)$ for $n=1,2,3$ under an exponential demand model with $T=1$ and $\lambda^*=10$. A distinctive feature of these prices paths in the Gallego and van Ryzin model is that they exhibit "*accelerating markdowns*." This is in contrast to some of the optimal price paths in our robust minimax setting that have "*decelerating markdowns*" (e.g., upper bound price path depicted in Figure 5).

The attentive reader may have have noticed that there is one important aspect of our model that is not captured by the Gallego and van Ryzin (1994) framework, namely, the fact that the seller uses a discount factor, r, to penalize future cash-flows. Since the use of a discount factor creates pressure on the seller to try to expedite sales and possibly engage in early markdowns, it is plausible that an optimal price path in a modified Gallego and van Ryzin setting that includes discounting could exhibit decelerating markdowns. In what follows we show that this is not the case.

As we argued in the previous paragraph, for the purpose of having a more direct and fair (apples-to-apples) comparison between the results in Gallego and van Ryzin and our model, we next modify their setting to account for this difference. We would also like to select the values of n and $\lambda(p)$ in a way that it is consistent with our distribution-free analysis. To this end, we first note that

Figure B.2. (Color online) Comparison of robust price paths to modified Gallego and van Ryzin solution.

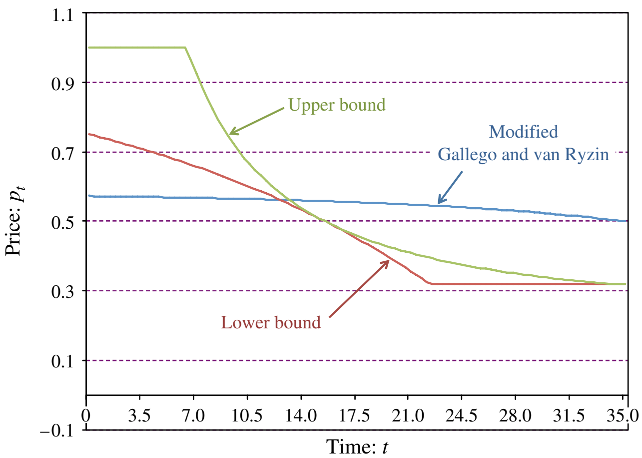

regarding our robust prior-free formulation, the "natural" choice for modeling consumers' willingness-to-pay is a Uniform distribution in $[\underline{v}, \bar{v}]$, which is equivalent to assuming that the demand function $\lambda(p)$ is linear, or $\lambda(p) = \Lambda(\bar{v}-p)/(\bar{v}-\underline{v})$ for $p \in [\underline{v}, \bar{v}]$ and for some positive scalar Λ. In addition, by the linearity of the regret on the number of customers (Proposition 1), we can restrict our analysis (w.l.o.g.) to the case of a single consumer. Therefore, we set $n = 1$ and select the values of Λ and T (the selling horizon) so that $\Lambda T = 1$. Under these conditions, one can show that the *modified Gallego and van Ryzin model* (that incorporates discounting) leads to the following optimality (HJB) condition for the value function $J(t)$

$$-\frac{\mathrm{d}J(t)}{\mathrm{d}t} = \max_p \left\{ \Lambda \frac{\bar{v}-p}{\bar{v}-\underline{v}} [p - J(t)] \right\} - rJ(t),$$

with boundary condition $J(T) = 0$.

After the solving the maximization on p, this HJB equation reduces to a standard Riccati equation that leads to following solution.

$$J(t) = \bar{v} + 4\frac{\bar{v}-\underline{v}}{\Lambda} \frac{A\alpha_1 e^{\alpha_1 t} + \alpha_2 e^{\alpha_2 t}}{A e^{\alpha_1 t} + e^{\alpha_2 t}},$$

where

$$A = -\frac{4\alpha_2(\bar{v}-\underline{v}) + \Lambda\bar{v}}{4\alpha_1(\bar{v}-\underline{v}) + \Lambda\bar{v}} e^{(\alpha_2-\alpha_1)T} \quad \text{and}$$

$$\alpha_{1,2} = \frac{r \pm \sqrt{r^2 + \Lambda r/(\bar{v}-\underline{v})}}{2}.$$

Finally, the optimal price policy is given by

$$p^*(t) = \frac{\bar{v} + J(t)}{2}.$$

Figure B.2 plots this *modified Gallego and van Ryzin* price path using the data in Figure 5 in the paper, namely $\bar{v} = 1$, $\underline{v} = 0.32$, $T = 35$ and $r = 0.045$. The figure also depicts our distribution-free upper bound $\bar{p}_t^*$ and lower bound $\underline{p}_t^*$ price paths.

Again, we can see that despite the additional pressure to start marking down early created by the discount factor, optimal price paths in the Gallego and van Ryzin framework remain concave, i.e., exhibit *accelerating markdowns*. We also note that the optimal price path in the modified Gallego and van Ryzin model is more stable (flat) than the optimal robust counterpart.

Endnotes

1. We are not the first ones to argue that maximin utility leads to excessively conservative solutions. In fact, avoiding excessive conservativeness is a common thread in the robust optimization literature (see, e.g., the survey by Bertsimas et al. 2011). Within the context of robust pricing, the criticism of the maximin utility criterion offered here is a paraphrasing of the one in Bergemann and Schlag (2008).

2. We use the words "consumer" and "customer" interchangeably. We do the same with the words "firm" and "seller." We use feminine pronouns to refer to the firm and masculine ones for the consumers.

3. Such a notion of regret for dynamic decision-making environments has been labeled anticipated ex post regret by Hayashi (2008).

4. If we were to consider a static problem ($T = 0$) and allow the firm to randomize prices, the model we study would coincide with Bergemann and Schlag (2008).

5. Unless $t^* = 0$ and $\bar{p}_t^* = \bar{v}$ that can only happen if $R_{\mathrm{M}}^* = 0$ or equivalently $\bar{v} = \underline{v}$.

References

Ahn H-S, Gümüs M, Kaminsky P (2007) Pricing and manufacturing decisions when demand is a function of prices in multiple periods. *Oper. Res.* 55(6):1039–1057.

Aviv Y, Pazgal A (2008) Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management* 10(3):339–359.

Ben-Tal A, Ghaoui LE, Nemirovski A (2009) *Robust Optimization* (Princeton University Press, Princeton, NJ).

Bergemann D, Schlag K (2008) Pricing without priors. *J. Eur. Econom. Assoc.* 6(2–3):560–569.

Bergemann D, Schlag K (2011) Robust monopoly pricing. *J. Econom. Theory* 146(6):2527–2543.

Bertsimas D, Brown DB, Caramanis C (2011) Theory and applications of robust optimization. *SIAM Rev.* 53(3):464–501.

Bertsimas D, Sim M (2004) The price of robustness. *Oper. Res.* 52(1): 35–53.

Besbes O, Lobel I (2015) Intertemporal price discrimination: Structure and computation of optimal policies. *Management Sci.* 61(1): 92–110.

Bitran G, Caldentey R (2003) An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* 5(3): 203–230.

Board S (2008) Durable-goods monopoly with varying demand. *Rev. Econom. Stud.* 75(2):391–413.

Borgs C, Candogan O, Chayes J, Lobel I, Nazerzadeh H (2014) Optimal multiperiod pricing with service guarantees. *Management Sci.* 60(7): 1792–1811.

Cachon GP, Swinney R (2009) Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Sci.* 55(3): 497–511.

Caldentey R, Vulcano G (2007) Online auction and list price revenue management. *Management Sci.* 53(5):795–813.

Coase RH (1972) Durability and monopoly. *J. Law Econom.* 15(1): 143–149.

Correa J, Montoya R, Thraves C (2016) Contingent preannounced pricing policies with strategic consumers. *Oper. Res.* 64(1):251–272.

Deb R (2014) Intertemporal price discrimination with stochastic values. Working paper, University of Toronto, Toronto.

Elmaghraby W, Keskinocak P (2003) Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Sci.* 49(10):1287–1309.

Epstein LG, Schneider M (2003) Recursive multiple-priors. *J. Econom. Theory* 113(1):1–31.

Eren SS, Maglaras C (2010) Monopoly pricing with limited demand information. *J. Revenue and Pricing Management* 9(1):23–48.

Gallego G, van Ryzin G (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* 40(8):999–1020.

Garrett D (2014) Dynamic mechanism design: Dynamic arrivals and changing values. Working paper, Toulouse School of Economics, Toulouse, France.

Gilboa I, Schmeidler D (1989) Maxmin expected utility with non-unique prior. *J. Math. Econom.* 18(2):141–153.

Hayashi T (2008) Regret aversion and opportunity dependence. *J. Econom. Theory* 139(1):242–268.

Hayashi T (2009) Stopping with anticipated regret. *J. Math. Econom.* 45(7):479–490.

Hayashi T (2011) Context dependence and consistency in dynamic choice under uncertainty: The case of anticipated regret. *Theory and Decision* 70(4):399–430.

Hendel I, Nevo A (2013) Intertemporal price discrimination in storable goods markets. *Amer. Econom. Rev.* 103(7):2722–2751.

Li J, Granados N, Netessine S (2014) Are consumers strategic? Structural estimation from the air-travel industry. *Management Sci.* 60(9): 2114–2137.

Lim A, Shanthikumar G (2007) Relative entropy, exponential utility, and robust dynamic pricing. *Oper. Res.* 55(2):198–214.

Liu Q, van Ryzin GJ (2008) Strategic capacity rationing to induce early purchases. *Management Sci.* 54(6):1115–1131.

Liu Y, Cooper W (2015) Optimal dynamic pricing with patient customers. *Oper. Res.* Forthcoming.

Lobel R, Perakis G (2010) Dynamic pricing through sampling based optimization. Working paper, MIT, Cambridge, MA.

Mersereau A, Zhang D (2012) Markdown pricing with unknown fraction of strategic customers. *Manufacturing Service Oper. Management* 14(3):355–370.

Milnor J (1954) Games against nature. Thrall RM, Coombs CH, Davis RL, eds. *Decision Processes* (John Wiley & Sons, New York), 49–59.

Mirrlees J (1971) An exploration in the theory of optimum income taxes. *Rev. Econom. Stud.* 38(2):175–208.

Mussa M, Rosen S (1978) Monopoly and product quality. *J. Econom. Theory* 18(2):301–317.

Osadchiy N, Vulcano G (2010) Selling with binding reservations in the presence of strategic consumers. *Management Sci.* 56(12): 2173–2190.

Pavan A, Segal I, Toikka J (2014) Dynamic mechanism design: A myersonian approach. *Econometrica* 82(2):601–653.

Perakis G, Roels G (2010) Robust controls for network revenue management. *Manufacturing Service Oper. Management* 12(1):56–76.

Rockafellar RT (1997) *Convex Analysis* (Princeton Landmarks in Mathematics, Princeton, NJ).

Savage L (1951) The theory of statistical decision. *J. Amer. Statist. Assoc.* 46(253):55–67.

Schlag K (2006) Why minimax regret? Unpublished manuscript, University of Vienna, http://homepage.univie.ac.at/karl.schlag/research/learning/whyMR.pdf.

Schlag K, Zapechelnyuk A (2015) Dynamic benchmark targeting. Working paper, University of Vienna, Vienna.

Stokey NL (1979) Intertemporal price discrimination. *Quart. J. Econom.* 93(3):355–371.

Stoye J (2011) Axioms for minimax regret choice correspondences. *J. Econom. Theory* 146(6):2226–2251.

Su X (2007) Intertemporal pricing with strategic customer behavior. *Management Sci.* 53(5):726–741.

Talluri KT, Van Ryzin GJ (2005) *The Theory and Practice of Revenue Management* (Springer, Berlin).

Wald A (1950) Statistical decision functions.

Wang Z (2016) Technical Note—Intertemporal price discrimination via reference price effects. *Oper. Res.* 64(2):290–296.

Yin R, Aviv Y, Pazgal A, Tang CS (2009) Optimal markdown pricing: Implications of inventory display formats in the presence of strategic customers. *Management Sci.* 55(8):1391–1408.

René Caldentey is a professor of operations management at the University of Chicago's Booth School of Business. His primary research interests include stochastic modeling with applications to revenue and retail management, queueing theory, and finance.

Ying Liu received her Ph.D. in operations management from the Stern School of Business at New York University. Her research focuses on revenue management, dynamic pricing and demand estimation.

Ilan Lobel is an associate professor of information, operations and management sciences at the Stern School of Business at New York University. He is interested in questions of pricing, learning, and contract design in dynamic and networked markets.