Monopoly Pricing with Limited Demand Information

Serkan S. Eren Costis Maglaras[1]

February 28, 2008

Traditional monopoly pricing models assume that firms have full information about the market demand and consumer preferences. In this paper we study a prototypical monopoly pricing problem for a seller with limited market information and different levels of demand learning capability under relative performance criterion of the competitive ratio. We provide closed-form solutions for the optimal pricing policies for each case and highlight several important structural insights. We note the following: a) From the firm's viewpoint the worst-case operating conditions are when it faces a homogeneous market where all customers value the product equally, but where the specific valuation is unknown. In cases with partial demand information, the worse case cumulative willingness-to-pay distribution becomes piecewise-uniform as opposed to a point mass. b) Dynamic (skimming) pricing arises naturally as a hedging mechanism for the firm against the two principal risks that it faces: first, the risk of foregoing revenue from pricing too low, and second, the risk of foregoing sales from pricing too high. And, c) even limited learning, e.g., market information at a few price points, leads to significant performance gains.

1 Introduction

Classical models from the economics and revenue management literature study monopoly pricing problems under the assumption that firms have accurate characterizations, potentially probabilistic, of the market demand and consumer preferences. In practice, however, there are many settings, such as introduction of new and innovative products, where one rarely has such full and accurate demand information. This source of model uncertainty may lead to significant revenue loss and may be insufficiently hedged against through the use of pricing policies that do not explicitly incorporate it in their derivation. This paper studies these two issues for a monopolist operating in settings with limited market information and different degrees of learning capability, and where model uncertainty is captured through the relative performance criterion of the competitive ratio.

As a motivating example consider a monopolist firm that offers a new product to a set of risk-neutral, heterogenous consumers, each endowed with a private willingness-to-pay (**WtP** or valuation), which is an independent draw from a common distribution. The market information is summarized by the number of potential consumers, i.e., the market size, and the WtP distribution.

[1]Graduate School of Business, Columbia University, Uris Hall, 3022 Broadway, New York, NY 10027, USA
se2027@columbia.edu, c.maglaras@gsb.columbia.edu

Classical models would assume that both of these elements are known to the firm, and are used to determine the expected revenue maximizing price; the market size is relevant when supply is limited. How should the seller approach this problem if the market size and WtP distribution were unknown? What should be the form of seller's pricing policy in that case? How should that be adjusted to take advantage of partial demand information extracted, for example, from experimenting at a few price points?

There are two natural ways to specify this type of model uncertainty that lead to different formulations and different policy recommendations. The first one is stochastic, wherein the unknown WtP distribution is assumed to be drawn from a given set of possible distributions according to some known probability law, and where the firm's goal is to optimize her expected revenues - potentially risk-adjusted- over all possible market model realizations. Its main shortcoming is that it requires detailed information on the distribution of the model uncertainty, which itself may not be available. The second formulation adopts a worst-case perspective using a max-min criterion on expected revenue, wherein the WtP distribution is assumed to be selected from an appropriate set of possible distributions by an imaginary adversary ("nature") to minimize the revenue, and where the firm's objective is to select its pricing policy to maximize the worst-case revenue performance. This criterion may yield overly pessimistic results; e.g., by setting the WtP of all consumers equal to its minimum allowed value irrespective of the pricing decision. To reduce this inherent conservatism, one typically imposes constraints on the decision set of the adversary, that are either ellipsoidals (see BenTal and Nemirovski (1998), and ElGhaoui and Lebret (1997)), or polyhedra (see Bertsimas and Sim (2003), as well as Bertsimas and Thiele (2004), Perakis and Sood (2003)). In a similar vein, Lim and Shanthikumar (2004) suggested using a relative entropy constraint to bound the distance of the WtP distribution from a nominal one. Each of these extensions essentially imposes an uncertainty "budget" to the adversary so as to reduce the pessimistic nature of the associated solution. The selection of this budget is often arbitrary, and analytical solutions are often not available, leading to numerically computed results that do not provide easily interpretable structural insights.

An alternative approach to reduce the conservatism of max-min formulations while maintaining their appealing low informational requirements is through the use of the competitive ratio criterion, which measures the performance relative to a fully-informed decision maker. In broad terms, this is defined as follows: the firm first selects a policy π, and the adversary selects a worst-case distribution function for the unknown consumer attribute, $F(\cdot)$; in the above example, π could be the posted price, and $F(\cdot)$ the WtP distribution. Let $R(\pi, F)$ be the *actual expected revenue earned* for the pair of actions π and $F(\cdot)$, and $R(\pi^*(F), F)$ be the maximum expected revenue the firm *could have extracted* if she knew the selected distribution $F(\cdot)$; here $\pi^*(F)$ denotes the optimal policy if F was known. The competitive ratio is given by

$$c^* = \max_{\pi} \ \min_{F} \ \frac{R(\pi, F)}{R(\pi^*(F), F)}$$

That is, the firm strives to minimize the relative difference from the maximum revenues it could have extracted in the full-information case.[2] To contrast, the max-min criterion takes the form $\max_{\pi} \min_{F} \ R(\pi, F)$. The relative performance criterion prevents trivial choices for the adversary, such as choosing the minimal WtP for all consumers, where the firm achieves a bad outcome no matter what the policy, because the fully-informed manager would also be harmed in such instances. In that sense it allows us to distinguish between "bad market conditions" and "bad decisions".

Relative performance criteria implicitly constrain the actions of the adversary without having to impose additional constraints, and often result in intuitive policy recommendations. They have been used extensively in the computer science literature, and have recently been applied in pricing and operations management problems. Specifically, Ball and Queyranne (2004) used a competitive ratio criterion for a single-resource capacity allocation problem, while Bergemann and Schlag (2005) and Perakis and Roels (2004) adopted the regret criterion to study the monopolist pricing and the newsvendor problems, respectively. Lan et al. (2006) generalizes Ball and Queyranne's analysis and extends it to cover the regret criterion as well. Perakis and Roels (2007) applies similar techniques for network revenue management. Eren and van Ryzin (2006) apply these criteria to the problems.

Our work adopts the competitive ratio criterion to study the monopolist's pricing problem described earlier under the assumption that the underlying WtP distribution is unknown. The firm has the ability to change its price over time, and its key decision is to figure out a pricing

[2]A related criterion is to minimize the absolute regret given by $r = \ \min_{\pi} \ \max_{F} \ [R(\pi^*(F), F) - R(\pi, F)]$. Most of our results extend to that case as well, but in the interest of space and since the structural insights gleaned from the regret analysis are similar to those extracted from the competitive ratio one, we will not study it herein.

policy (how much to charge and for how long to stay at each price point) that would perform well even though the firm does not know the underlying consumers' WtP distribution. An alternate interpretation of this price skimming policy[3] is to treat the relative length of time over which a price is offered as a probability of that price point, thus interpreting the proposed price scheme as a randomizing pricing policy; this was done in Bergemann and Schlag (2005).

Our model is deterministic, disregarding the stochastic variability of the sales process, e.g., due to its Poisson nature. This allows us to emphasize the effects of "first order" uncertainty introduced by not knowing the sales rate itself at a selected price point, as opposed to "second order" fluctuation due to stochastic nature of the process. Our analytical contributions are the following: 1) The competitive ratio optimization problem is solvable in closed form, offering a detailed description of the optimal policies for the seller and the adversary, the tradeoffs faced by the seller, and a precise characterization of the resulting revenue loss (Theorem 2.1 and Proposition 2.1). 2) We extend our formulation and results to a two period setting that allows the seller to learn from the sales observations in period one (Theorems 3.1 and 3.2). 3) We address the situation where the seller only has limited price experimentation capability, or has limited past sales information. Observing the demand at a price point gives cumulative demand information above and below that price point and essentially decomposes the problem into simpler subproblems in the respective regions that are readily solvable using linear programming techniques (Proposition 4.2). As a special case of practical interest, we also study the "ex-post" problem that allows the seller to take into account actual demand observation data in her price optimization and performance analysis decisions (Proposition 4.4).

We highlight three observations from our work that are of potential interest. First, the worse case market scenarios for the firm, as captured by the WtP distribution selected by the adversary, correspond to homogeneous markets where all consumers have the same, yet unknown, valuation.

[3] Dynamic pricing is concerned with adjusting prices to regulate demand over a finite sales horizon to maximize revenue. Price skimming is a commonly used example of such a policy in many industries like airlines, hospitality and fashion. Clearing excess inventory and perishable products –rather than salvaging leftover items at low value at the end of the sales horizon– has been proposed as a possible explanation for this practice; see Talluri and van Ryzin (2004) for a review of this body of work. Another possible explanation for the use of dynamic pricing policies is as a hedging mechanism in settings where demand is uncertain; see Lazear (1986) for an analysis of this problem and Pashigan (1988), and Pashigan and Bowen (1991) for empirical evidence of this explanation. Harris and Raviv (1981) showed that a price skimming policy may emerge as the optimal mechanism when demand is uncertain. Our work shows that such a policy will optimize the firm's relative revenue performance when the demand model is unknown.

Mathematically, this corresponds to an extreme point -unit mass- distribution whose exact position is uncertain, forcing to firm hedge against opposing risks at each price point: first, the risk of foregoing revenue from pricing too low, and second, the risk of foregoing sales from pricing too high. In response, the firm's strategy tries to hedge against this exposure.

The extreme point nature of the adversary's strategy has appeared elsewhere in the literature. One example form decision theory is Smith (1995), which studies the expectation maximization problem among a set of probability distributions. He shows the equivalence of that problem to a linear program and, as a result, recovers extreme point distributions as potential solution points. Our objective function is not linear (and not even convex), and our results do not follow from Smith's observation. However, in most of the papers that such a structure emerges, it happens because the inner optimization step can be reduced to a quasi-convex maximization problem over the probability simplex, which admits an extreme point solution. In settings with learning or with partial demand information, the worst-case distribution retains some of its structural form by having point masses at distinct valuations, but becomes more dispersed. We give a complete characterization of the latter and discuss several examples in Section 4.

Second, we highlight that in settings with limited or no market information it is optimal for the firm to adopt a price skimming policy to minimize the risk of lost sales and foregone revenue that could result from mis-estimating the market characteristics. To contrast, if the firm knew the customer WtP distribution, then it would be optimal to charge a static price over the entire sales horizon. This result suggests that lack of market information could offer one possible justification for the use of dynamic pricing policies (c.f. footnote 3). Analytically, the precise form of the resulting pricing policy ensures that both the firm and the adversary are indifferent with regard to the positioning (i.e., the representative valuation) of the market.

Third, the effect of learning is both significant and quick in the sense that even a few observations at different price points can provide considerable lift in the revenue performance of the proposed policies. Both the resulting competitive ratio, which is a worse case bound, and the actual performance relative to some underlying WtP distribution unknown to the seller, improve considerably. In the case where the seller is not restricted in the number of price points that she can experiment at, we show that it is optimal to use a price skimming policy during a "learning" period. This achieves full learning of the demand model and allows the seller to price at the optimal (full-information) price in the remainder of the sales horizon. Often, there may be practical constraints that link the firm's pricing decision over time, e.g. retailers hesitate to increase prices after an early mark down. We show that in such settings it is still optimal to adopt a price skimming policy, but in this case the seller is willing to sacrifice performance due to the learning phase so as to retain adequate pricing flexibility in the remainder of the sales horizon.

Incorporation of partial information is typically done in a Bayesian setting under some parametric assumptions for the WtP distribution and using conjugate pairs of distributions to maintain tractability; see, e.g., Lobo and Boyd (2003), Aviv and Pazgal (2005), Araman and Caldentey (2005), and Farias and Van Roy (2006). Assuming a parametric family of distributions for the unknown demand runs the risk of model misspecification due to the arbitrariness of that assumption. Similar to this paper, another subset of literature uses non-parametric approaches, which make minimal distributional assumptions and often involve some form of an adaptive learning algorithm; see, e.g., van Ryzin and McGill (2000), Huh and Rusmevichientong (2006), and Eren and Maglaras (2006). An interesting recent paper in the latter set is Besbes and Zeevi (2006) that studies a prototypical dynamic pricing problem in a stochastic environment. Two important insights from their work for purposes of our paper is that they show that: a) in settings with long sales horizons and large market sizes, an asymptotically optimal policy in terms of its relative regret is to divide the sales horizon in two phases that are dedicated to learning and revenue optimization, respectively; and b) the uncertainty due to the stochastic nature of the demand arrival process is indeed negligible in such settings.

As a closing remark we note that a potential practical shortcoming of an approach based on the relative performance criteria is that it may loose its analytic and computational tractability as one tries to incorporate partial information about the unknown demand model primitives. In that respect, most papers (and ours) that adopt this framework strive for the derivation of insights about the structure of good policies and the effect of ambiguity on system performance, as opposed to the computation of implementable policies. One exemption is Perakis and Roels (2007) which incorporates partial demand information using the probabilistic tight bounds for mean and variance specifications derived by Bertsimas and Popescu (2005). Another is Bergemann and Schlag

(2005) which allows for the unknown distribution to be within a distance of a nominal distribution that could encapsulate prior information. Both papers work with the regret criterion and do not generalize easily to the competitive ratio criterion, and both frameworks do not seem to allow for a tractable way in which to incorporate intuitive information extracted from past sales that usually translate to fractiles of the WtP distribution. Our paper shows how to incorporate the latter type of demand information in a tractable way.

The remainder of the paper is structured as follows. Section 2 introduces the prototypical dynamic pricing problem with no market information and no learning. Sections 3 and 4 study two period extensions that allow for different degrees of learning.

2 Dynamic Pricing with No Market Information

2.1 Problem formulation

We consider a monopolist selling a homogeneous good over a sales horizon that is normalized to have length one. The firm's is assumed to have ample capacity. Potential customers arrive at the firm according to a deterministic arrival process with rate Λ, each with a WtP for one unit of that product, denoted by v, which is an independent draw from a common discrete distribution F on the set $\{p_1, \ldots, p_K\}$ where $p_1 = \underline{v}$ and $p_K = \bar{v}$. That is the support of the WtP distribution is an appropriate discretization of the range $[\underline{v}, \bar{v}]$, e.g., in \$1 or 5% increments.[4] Assuming that the price at time t is equal to $p(t)$, then the sale rate at that instant is given by $\lambda(t) = \Lambda \mathrm{P}(v \geq p(t)) = \Lambda \bar{F}(p(t))$, where $\bar{F}(\cdot) = 1 - F(\cdot)$, and the corresponding revenue rate is $p(t)\Lambda\bar{F}(p(t))$.

The firm's goal is to maximize the total revenues accrued in $[0, 1]$. When the WtP distribution F is known, this problem reduces to a special case of the deterministic relaxation of the single-product dynamic pricing problem studied by Gallego and van Ryzin (1994) (that paper considered the capacity constrained case), for which it is optimal to charge a constant price $p^* = \operatorname{argmax}_i p_i \bar{F}(p_i)$ throughout the sales horizon. That is, the dynamic nature of the pricing decisions is inconsequential, and the above problem reduces to the classical monopolist pricing problem.

This paper considers the problem of selecting a pricing policy when the firm only knows the

[4] It is common to assume that the customer arrival process is Poisson, but in the sequel we will restrict attention to a deterministic model where, in addition, customers are assumed to arrive continuously as opposed to in unit increments. The rate Λ can also be interpreted as the "market size."

support $[\underline{v}, \bar{v}]$ of the distribution function F, but not F itself. Sections 3 and 4 will consider extensions that incorporate demand learning from early sales. Estimating (or bounding) the support rather than the distribution itself is much easier in practice; for example, $\underline{v}$ might represent the "cost of goods sold" below which the firm is not willing to engage in trade, and $\bar{v}$ might represent the price of a superior substitute in the market.

The firm's pricing strategy is a vector $t \in \mathbb{R}^K$, where t_i is the length of time over which the firm will use price p_i. Note that the labeling of the price points and the assumption that $p_1 = \underline{v}$ and $p_K = \bar{v}$ are innocuous since it is always possible to decide not to offer some particular price p_j by setting the corresponding $t_j = 0$. Given a policy t and a distribution F, the revenue accrued by the firm is given by $R(t, F) := \Lambda \sum_{j=1}^{K} t_j \, p_j \, \mathrm{P}(v \geq p_j) = \Lambda \sum_{j=1}^{K} t_j \, p_j \, \bar{F}(p_j)$.

The firm selects a strategy t, and then an imaginary adversary selects a distribution F after he observes the firm's policy t. The goal of the firm is to optimize its relative performance when compared to that of a fully informed player, i.e. one that could maximize its revenues with full knowledge of the distribution F; this is the so called "competitive ratio" criterion.

Specifically, let $t^*(F) \in \operatorname{argmax}_t R(t, F)$, be the policy that maximizes the total revenue with full information about $F(\cdot)$, which is given by $t^*_j(F) = 1$ for $j = \operatorname{argmax}_i \; p_i \bar{F}(p_i)$ and $t^*_i(F) = 0$ for all $i \neq j$. The competitive ratio (CR) problem is given by

$$c^* = \max_t \quad \min_F \left\{ \frac{R(t, F)}{R(t^*(F), F)} \;:\; \sum_{j=1}^{K} t_j = 1, \quad t \geq 0 \right\}. \tag{1}$$

2.2 Characterization of the optimal pricing policy

For any distribution F, let $f_j := \mathrm{P}(p_{j+1} > v \geq p_j)$ for $j = 1 \dots K-1$, $\;f_K := \mathrm{P}(v = \bar{v})$, and $\bar{f}_j := \sum_{j \leq k} f_k = \mathrm{P}(v \geq p_j)$. This allows us to rewrite the revenue function as $R(t, F) = \sum_k f_k \sum_{j \leq k} p_j \; t_j = \sum_j t_j p_j \bar{f}_j$, and (1) as:

$$c^* = \max_t \;\; \min_{\bar{f}} \;\; \left\{ \frac{\sum_j t_j p_j \bar{f}_j}{\max_j \{p_j \bar{f}_j\}} \;\; : \;\; 1 = \bar{f}_1 \geq \bar{f}_2 \geq \cdots \bar{f}_K \geq 0, \quad \sum_j t_j = 1, \quad t \geq 0 \right\}, \tag{2}$$

where the denominator, $\max_j \{p_j \bar{f}_j\} = R(t^*(F), F)$, is the maximum revenue that the firm could extract, if $F(\cdot)$ was known, by charging the revenue maximizing price throughout the sales horizon.

The key observation that underlies the solution of (2) is that the objective function is quasi-concave in $\bar{f}$, and as a result, the adversary's problem admits an extreme point optimal solution which is easy to characterize and exploit.

Theorem 2.1 *Consider the dynamic pricing problem with no market information specified in (1), or equivalently in (2). The firm's optimal policy is the following price skimming rule:*

$$t_1 = \left(K - \sum_{j=1}^{K-1} \frac{p_j}{p_{j+1}} \right)^{-1} \quad \textit{and} \quad t_j = \frac{p_j - p_{j-1}}{p_j}\, t_1 \quad \textit{for } j = 2, \ldots, K.$$

and the resulting competitive ratio is $c^* = t_1$.

Proof: The denominator in the objective in (2) is the maximum of K linear functions in $\bar{f}$, and is therefore convex in $\bar{f}$. The numerator in (2) is linear, and thus concave in $\bar{f}$. Thus, for fixed t, the adversary's problem is one of minimizing a quasi-concave function over a polyhedron, which admits an extreme point optimal solution. The polyhedron defined by $1 \geq \bar{f}_2 \geq \cdots \bar{f}_K \geq 0$ has K extreme points, all of which correspond to vectors of the form $(1, 0, 0, \ldots, 0), (1, 1, 0, \ldots, 0), \ldots (1, 1, 1, \ldots, 1)$. Since, for every fixed t the optimal value for the inner minimization occurs at one of these extreme points, (2) can be reduced to the problem:

$$c^* = \max_t \left\{ \min_{j=1,\ldots,K} \frac{\sum_{i \leq j} p_i\; t_i}{p_j} \quad : \; \textstyle\sum_j t_j = 1, \quad t \geq 0 \right\},$$

which, in turn, is equivalent to the linear program

$$c^* = \max_{t,c} \left\{ \; c \; : \; c \leq \textstyle\sum_{i \leq j} \frac{p_i}{p_j}\; t_i \;\; \forall j, \quad \sum_j t_j = 1, \quad t \geq 0 \right\}.$$

This LP can be solved in closed-form as follows. Consider its dual:

$$c^* = \min_{x,y} \left\{ \; y \; : \; y \geq p_j \textstyle\sum_{i \geq j} \frac{x_i}{p_i} \;\; \forall j, \quad \sum_j x_j = 1, \quad x \geq 0 \right\}.$$

The first step is to construct a dual feasible solution that satisfies the first set of inequality constraints with equalities. Solving $y = p_j \sum_{i \geq j} \frac{x_i}{p_i}$ $\forall j$, then $x_K = y$ and $x_j = \frac{p_{j+1} - p_j}{p_{j+1}}\; y$ for $j = 1, \ldots, K-1$. Substituting these into the normalizing constraint $\sum_j x_j = 1$, yields $y =$

$\left(K - \sum_{j=1}^{K-1} \frac{p_j}{p_{j+1}}\right)^{-1}$. This dual objective value is equal to the primal objective value that corresponds to the primal feasible solution given in the body of Theorem 2.1. By strong duality, we conclude that the solution in the proposition is optimal for the primal problem. □

Consequently, if the firm does not know the WtP distribution F, it no longer charges a constant price, but it adopts a price skimming policy that charges each price point for an appropriate amount of time. As mentioned earlier, an alternative interpretation is to treat the t_j's as probabilities and t as a randomized pricing policy; see Bergemann and Schlag (2005). To gain some intuition behind this result, recall that the worst case scenario for the firm occurs when the market is homogeneous and every potential customer shares the same WtP. This setting raises two types of opposing risk for the firm at each price. First, if the firm prices too high for a significant portion of its sales horizon, it may suffer low sales when the market's WtP is low. Second, if the firm prices too low for a significant portion of its sales horizon, it may forego a significant revenue opportunity when the market's WtP is high. In both cases, the resulting competitive ratio would be low. Our analysis specifies how to balance these two effects in constructing the optimal pricing policy, which essentially makes the adversary indifferent between the extreme market scenarios that is optimal for him to choose. It is also worth comparing the above behavior against the solution to the maxmin formulation with objective $\max_t \ \min_F R(t, F)$. In this case, the optimal strategy for the adversary is to put all of the probability at $\underline{v}$, while the firm would also price at $\underline{v}$ for the entire sales horizon, making the result too conservative. Actually, the revenue performance of the resulting policy is typically much higher than the competitive ratio as illustrated by the numerical examples in Sections 3 and 4.

The competitive ratio for (2) depends on the discretization of the grid $\{p_1, \ldots, p_K\}$. The next result derives a lower bound for the competitive ratio that is independent of that grid.

Proposition 2.1 *For any price grid $\{p_1, \ldots, p_K\}$ used of any size K, the optimal competitive ratio c^* for (2) satisfies the following bound:*

$$c^* \geq (1 + \ln(\bar{v}/\underline{v}))^{-1} =: c^{LB}.$$

Proof: Note that

$$
\begin{aligned}
K - \sum_{j=1}^{K-1} \frac{p_j}{p_{j+1}} &= 1 + \sum_{j=1}^{K-1} \frac{p_{j+1} - p_j}{p_{j+1}} = 1 + \sum_{j=1}^{K-1} \int_{p_j}^{p_{j+1}} \frac{1}{p_{j+1}}\, dx \leq 1 + \sum_{j=1}^{K-1} \int_{p_j}^{p_{j+1}} \frac{1}{x}\, dx \\
&= 1 + \int_{\underline{v}}^{\bar{v}} \frac{1}{x}\, dx = 1 + \ln(\bar{v}/\underline{v})
\end{aligned}
$$

Then, c^* satisfies

$$
c^* = \left(K - \sum_{j=1}^{K-1} \frac{p_j}{p_{j+1}} \right)^{-1} \geq (1 + \ln(\bar{v}/\underline{v}))^{-1} = c^{LB}. \qquad \square
$$

The lower bound is achieved as the number of prices grows large and $\{p_1, \ldots, p_K\}$ becomes a dense covering of the range $[\underline{v}, \bar{v}]$. As intuition would suggest, as the relative difference between the lowest and the highest valuation decreases, i.e. as $\underline{v}/\bar{v} \to 1$, $c^{LB} \uparrow 1$. That is, as the aggregate uncertainty about the market preferences decreases, the risk from pricing too low becomes negligible, and the firm's revenue approaches the one that is achieved under full knowledge of F.

We note that Ball and Queyranne (2004) recover similar results in their study of the single-resource (airline) capacity allocation problem under a competitive ratio criterion. Similarly, Bergemann and Schlag (2005) obtained analogous results using the maximum regret criterion.

3 The Effect of Learning

We consider a version of the dynamic pricing problem with passive learning where: a) the firm splits the sales horizon into two periods of length τ^1 and $\tau^2 = 1 - \tau^1$; b) selects a pricing strategy in $[0, \tau^1]$ without any information about F; and c) subsequently selects a pricing strategy in $(\tau^1, 1]$ that uses fractile information extracted in $[0, \tau^1]$ for all price points charged in that period. Note that the adversary need only commit to the underlying WtP distribution at the price points where the firm choose to experiment in period one, and can select the remaining information, i.e. unobserved specifications of the distribution, in period two. We assume that the demand measurements are noiseless, i.e. the firm observes $\bar{F}(p_j)$ instead of a random variable with that mean. The term "passive" learning indicates that the firm cannot update its information set incrementally during $[0, \tau^1]$ or $(\tau^1, 1]$; instead, it only updates its information at time τ^1, and makes use of this new information in $(\tau^1, 1]$. In that sense, the first period has a dual role of learning and revenue optimization, while the second period is solely dedicated to revenue optimization. This structure is suggested due to its simplicity and potential practical appeal, which was recently proposed and analyzed in an asymptotic setting with large market size and large sales horizon by Besbes and Zeevi (2006).

An alternative interpretation of our model is one where the firm sells one product through many potential stores, and where rather than dynamic pricing over time, the fractions t_j^i represent the fraction of the stores that price at p_j in period i for $j = 1, \ldots, K$ and $i = 1, 2$. [This is also an alternate interpretation for the randomized policy of Bergemann and Schlag (2005).] In that setting, the firm selects at how many stores to apply each potential price point, and then combines the information extracted from all these stores to update its demand information and optimize its downstream pricing decisions. Gaur and Fisher (2005) have studied this problem, although the emphasis in their paper was the issue of how to combine the demand information from each store taking into account the differences between the local market conditions faced by each store; this feature is not considered in our paper.

An important consideration is whether the prices used during the first period constrain the firm's pricing options in the second period. We study two variants:

a) *Unconstrained learning*: no downstream pricing constraints in period two; i.e., both markups and markdowns are allowed.

b) *Constrained learning*: the firm can only apply price markdowns in period two. This is easiest understood through the interpretation of a firm selecting how to price a product across many stores, in which case in a constrained setting, if a particular store priced at p_j in period one, then the same store can only price at p_j or lower in period two. This constraint -or slight variation thereof- is plausible from a practical viewpoint, and creates a clear trade-off between the pricing decisions in the two periods. Algebraically, this constraint translates into $\sum_{j \geq k} t_j^1 \geq \sum_{j \geq k} t_j^2$ for all k, where t_j^1 and t_j^2 are the fractions of the intervals $[0, \tau^1]$ and $(\tau^1, 1]$ respectively dedicated to price p_j and $\sum_j t_j^i = 1$ for $i = 1, 2$.

These two variants are analyzed in the following subsections. As it will be apparent, an important feature of the emerging solution for both variants is that the firm chooses to experiment on all possible price points in $\{p_1, \ldots, p_K\}$ in period one, so as to price under full information in period two. Section 4 will study more restricted settings where the firm can only experiment on a few price points in period one.

3.1 Unconstrained learning

Adopting our previous notation, we will normalize the length of each period to one but scale the market size that corresponds to each period to Λ^1 and Λ^2, where $\Lambda^i = \Lambda\tau^i$ for $i = 1, 2$. This problem can be formulated as the following two-stage dynamic game:

$$\begin{aligned} c^* \;=\; \max_{t^1}\min_{\bar{f}^1}\max_{t^2}\min_{\bar{f}^2} \quad & \frac{\Lambda^1 \sum t_j^1 p_j \bar{f}_j^1 + \Lambda^2 \sum t_j^2 p_j \bar{f}_j^2}{(\Lambda^1 + \Lambda^2)\max_j\{p_j \bar{f}_j^2\}} && \\ s.t. \quad & \sum_j t_j^i = 1, \quad t^i \geq 0 && i = 1,\ 2 \\ & 1 = \bar{f}_1^i \geq \bar{f}_2^i \geq \cdots \bar{f}_K^i \geq 0 && i = 1,\ 2 \\ & t_j^1(\bar{f}_j^2 - \bar{f}_j^1) = 0, && \forall j \end{aligned} \tag{3}$$

where t_j^i is the proportion of time spent at price j during period i for $j = 1, \ldots, K$ and $i = 1, 2$. $\bar{f}_j^2,\ j = 1, \ldots, K$ are the fractiles of the WtP distribution chosen by the adversary in period two, while $\bar{f}_j^1$ represent the information revealed to the firm at the end of period one. The constraints $t_j^1(\bar{f}_j^2 - \bar{f}_j^1) = 0$ force the adversary to be consistent in the choice of the WtP distribution and the information revealed to the firm in the first period. There are no constraints linking the pricing decisions in periods one and two.

We first prove that the firm tests all prices in period one. (The proof is relegated to the appendix.)

Proposition 3.1 *Let (t_*^1, t_*^2) denote the solution of (3). Then, $t_{*,j}^1 > 0$ for $j = 1, \ldots, K$.*

This allows the firm to price under full information in period two such that

$$\max_{t^2}\left\{\Lambda^2 \sum t_j^2 p_j \bar{f}_j^1 \;:\; \sum_j t_j^2 = 1, \quad t^2 \geq 0\right\} = \Lambda^2 \max_j\{p_j \bar{f}_j^1\},$$

i.e. the firm extracts the maximum possible revenue in period two, same as the adversary. Given this observation the problem reduces to

$$\max_{t^1}\min_{\bar{f}^1}\left\{\frac{\Lambda^1\sum t_j^1 p_j \bar{f}_j^1}{(\Lambda^1+\Lambda^2)\max_j\{p_j\bar{f}_j^1\}}+\frac{\Lambda^2}{(\Lambda^1+\Lambda^2)}\ :\sum_j t_j^1=1,\ t^1\geq 0,\ 1=\bar{f}_1^1\geq\cdots\bar{f}_K^1\geq 0\right\}, \quad (4)$$

which is equivalent to the single period problem studied in Section 2. Using Theorem 2.1, we conclude the following:

Theorem 3.1 *For the two period dynamic pricing problem with "unconstrained learning" described in (3), the firm's optimal decision is the following*

- *Period one: adopt a price skimming policy for which*

$$t_1=\left(K-\sum_{j=1}^{K-1}\frac{p_j}{p_{j+1}}\right)^{-1} \quad \text{and} \quad t_j=\frac{p_j-p_{j-1}}{p_j}\,t_1 \quad \text{for } j=2,\ldots,K.$$

- *Period two: price at p_{j^*}, where $j^*=\operatorname{argmax}_j\{p_j\bar{f}_j^1\}$.*

Let $\lambda^i:=\Lambda^i/(\Lambda^1+\Lambda^2),\ i=1,2$. The competitive ratio is

$$c^*=\lambda^1\left(K-\sum_{j=1}^{K-1}\frac{p_j}{p_{j+1}}\right)^{-1}+\ \lambda^2. \quad (5)$$

That is the competitive ratio is a weighted average of the one for the single period problem identified in Theorem 2.1 and 1 with respective weights of λ^1 and λ^2. As we shrink the length of the learning period $\lambda^1\to 0,\ \lambda^2\to 1$ and $c\to 1$. This follows from the assumption that allows for perfect learning of the demand at each tested price point irrespective of the time spent on it. In reality, as the length of the learning phase decreases, the accuracy of the firm's observations diminishes due to the inherent uncertainty of the demand realization, which may lead to an error in the estimation of the underlying WtP distribution and of the firm's downstream pricing decisions.

As mentioned earlier, Besbes and Zeevi (2006) studied a stochastic variant of what we call here the "unconstrained" problem in settings where the market size is large, and showed that the firm can select the learning period to be short and still ensure that the estimation error due to the stochastic nature of the demand is asymptotically negligible. As $\lambda_1\to 0$, the pricing policy in the first period becomes irrelevant as long as the firm does try all K price points; indeed Besbes and Zeevi (2006) prescribes a uniform pricing policy in period one, since its effect on the overall performance becomes negligible.

3.2 Constrained learning

As explained earlier the downstream pricing constraints can be succinctly summarized by the set of conditions $\sum_{j\geq k} t^1_j \geq \sum_{j\geq k} t^2_j$ for all k. The resulting formulation becomes:

$$\begin{aligned} c^* \;=\; \max_{t^1}\min_{\bar{f}^1}\max_{t^2}\min_{\bar{f}^2} \quad & \frac{\Lambda^1 \sum t^1_j p_j \bar{f}^1_j + \Lambda^2 \sum t^2_j p_j \bar{f}^2_j}{(\Lambda^1+\Lambda^2)\max_j\{p_j \bar{f}^2_j\}} && (6)\\ s.t. \quad & \sum_j t^i_j = 1, \quad t^i \geq 0 && i=1,\ 2\\ & 1 = \bar{f}^i_1 \geq \bar{f}^i_2 \geq \cdots \bar{f}^i_K \geq 0 && i=1,\ 2\\ & t^1_j(\bar{f}^2_j - \bar{f}^1_j) = 0, && \forall j\\ & \sum_{j\geq k} t^1_j \geq \sum_{j\geq k} t^2_j && \forall k. \end{aligned}$$

We state the following result without proof (which follows the steps of Proposition 3.1).

Proposition 3.2 *Let (t^1_*, t^2_*) denote the solution to (6). Then, $t^1_{*,j} > 0$ for $j = 1, \ldots, K$.*

Using this result, the problem can be reduced to the following formulation:

$$\begin{aligned} \max_{t^1}\min_{\bar{f}^1} & \frac{\Lambda^1 \sum t^1_j p_j \bar{f}^1_j + \max_{t^2}\left\{\Lambda^2 \sum t^2_j p_j \bar{f}^1_j \;\; : \sum_{j\geq k} t^1_j \geq \sum_{j\geq k} t^2_j \;\; \forall k, \;\; \sum_j t^2_j = 1, \; t^2 \geq 0\right\}}{(\Lambda^1+\Lambda^2)\max_j\{p_j \bar{f}^1_j\}} \qquad (7)\\ s.t. \qquad & \sum_j t^1_j = 1, \quad t^1 \geq 0, \quad 1 = \bar{f}^1_1 \geq \bar{f}^1_2 \geq \cdots \bar{f}^1_K \geq 0 \; . \end{aligned}$$

The optimization in the second term of the numerator reflects the revenue maximization problem of the firm in the second period under full information but with the downstream pricing constraints. It is easy to show that in period two the firm adopts the revenue maximizing price $p_{j^*} = \operatorname{argmax}\{p_j \; \bar{f}^1_j\}$ for as long as possible, while marginally satisfying the downstream pricing constraints for all prices below p_{j^*}; at the optimal solution, $t^2_j = t^1_j$ for $j < j^*$ and $t^2_{j^*} =$

$\left(1 - \sum_{i<j^*} t_i^1\right)$. As a result, we can rewrite (7) as

$$
\begin{aligned}
c^* \;=\; \max_{t^1} \min_{\bar{f}} \quad & \frac{\Lambda^1 \sum t_j^1 p_j \bar{f}_j + \Lambda^2 \;\left[\sum_{i<j^*} t_i^1 p_i \bar{f}_i + \left(1 - \sum_{i<j^*} t_i^1\right)\; p_{j^*}\; \bar{f}_{j^*}\right]}{(\Lambda^1 + \Lambda^2) \max_j \{p_j \bar{f}_j\}} \\
s.t. \quad & \sum_j t_j^1 = 1, \quad t^1 \geq 0, \quad 1 = \bar{f}_1^1 \geq \bar{f}_2^1 \geq \cdots \bar{f}_K^1 \geq 0 \;.
\end{aligned}
\tag{8}
$$

Theorem 3.2 *For the two period dynamic pricing problem with "constrained learning" described in (6), the firm's optimal decision is the following*

- *Period one: adopt a price skimming policy for which*

$$
t_1^1 = \left(1 + \frac{p_2 - p_1}{p_2} + \sum_{j=3}^{K} \frac{(p_j - p_{j-1}) \prod_{i=2}^{j-1} [p_i - \lambda^2 p_{i-1}]}{\prod_{i=2}^{j} p_i \quad (\lambda^1)^{j-1}}\right)^{-1} \quad \textit{and} \quad t_2^1 = \frac{p_2 - p_1}{p_2},
$$

$$
t_j^1 = \frac{(p_j - p_{j-1})[p_{j-1} - \lambda^2 p_{j-2}]}{(p_{j-1} - p_{j-2}) p_j \lambda^1} \; t_{j-1}^1 \quad \textit{for } j = 3, \ldots, K.
$$

- *Period two: price at p_{j^*} for $t_{j^*}^2 = \left(1 - \sum_{i<j^*} t_i^1\right)$ and at p_j for t_j^1 for $j < j^*$, where $j^* = \text{argmax}_j \{p_j \bar{f}_j^1\}$.*

Let $\lambda^i = \Lambda^i / (\Lambda^1 + \Lambda^2), \; i = 1, 2$. The competitive ratio is

$$
c^* \;=\; \lambda^1 \; \left(1 + \frac{p_2 - p_1}{p_2} + \sum_{j=3}^{K} \frac{(p_j - p_{j-1}) \prod_{i=2}^{j-1} [p_i - \lambda^2 p_{i-1}]}{\prod_{i=2}^{j} p_i \quad (\lambda^1)^{j-1}}\right)^{-1} \;+\; \lambda^2 \;. \tag{9}
$$

Proof: The proof makes use of the following result, which we prove in the appendix.

Lemma 3.1 *For any feasible t^1, the optimal solution for the inner minimization in (8) occurs at an extreme point of the simplex $1 \geq \bar{f}_2 \geq \cdots \bar{f}_K \geq 0$.*

There are K extreme points to consider for the inner minimization in problem (8) above corresponding to vectors of the form $(1, 0, 0, \ldots, 0), (1, 1, 0, \ldots, 0), \ldots (1, 1, 1, \ldots, 1)$. Using Lemma 3.1, we can rewrite (8) as follows:

$$
c^* = \; \max_{t} \min_{j=1\cdots K} \left\{ \frac{\Lambda^1 \sum_{i \leq j} p_i t_i^1 \; + \; \Lambda^2 \left[\sum_{i<j} p_i t_i^1 + \left(1 - \sum_{i<j} t_i^1\right) \; p_j\right]}{(\Lambda^1 + \Lambda^2) p_j} \quad : \; \sum_j t_j^1 = 1, \;\; t^1 \geq 0 \right\},
$$

which, in turn, is equivalent to:

$$
\begin{aligned}
c^* = \max_{t,c} \quad & c \\
s.t. \quad & c \leq \frac{\Lambda^1 p_1 t_1^1 \; + \; \Lambda^2 p_1}{(\Lambda^1 + \Lambda^2) p_1} \\
& c \leq \frac{\Lambda^1 \sum_{i \leq j} p_i t_i^1 \; + \; \Lambda^2 \left[\sum_{i<j} p_i t_i^1 + \left(1 - \sum_{i<j} t_i^1\right) \; p_j \right]}{(\Lambda^1 + \Lambda^2) p_j} \qquad j = 2 \cdots K \\
& \sum_j t_j^1 = 1, \quad t^1 \geq 0 \; .
\end{aligned}
$$

This LP can be solved in closed form using its dual. This completes the proof. □

We make three observations. First, irrespective of whether the firm has downstream pricing constraints, it is optimal to adopt a price skimming policy which charges all prices for a positive amount of time in period one. Consequently, the optimal decision of the firm again decomposes into two parts: in the first period, the firm tests all prices in an optimal manner to learn the WtP distribution. In the second period, the firm maximizes its revenue under full information by charging p_{j^*} as long as possible and meeting the downstream pricing constraints marginally.

Second, the effect of downstream pricing constraints makes the firm charge higher prices in the first period, compared to the unconstrained case, to hedge against foregone revenues in the second period from not being able to charge higher prices.

Third, the constrained formulation offers a natural extension to the model studied in Besbes and Zeevi (2006) in the sense that one could adopt their style of analysis to prove the asymptotic optimality of our proposed policy. In contrast to our comments after the analysis of the unconstrained learning case, the pricing policy adopted during the learning phase has a crucial effect on the overall system performance even if the length of the learning phase is shrunk to zero.

3.3 Discussion and numerical results

The numerical results reported next give a rough indication of the theoretical performance improvement under these two learning schemes when compared to the results in Section 2 with no learning. Specifically, for a fixed price grid with K price points that uniformly span the support $[\underline{v}, \bar{v}]$ and given Λ^1, Λ^2, the solution of the single period problem with $\Lambda = \Lambda^1 + \Lambda^2$, identified by Theorem

2.1, is compared to the solutions of the two learning schemes identified in Theorems 3.1 and 3.2.

The main observation from these results is that the effect of learning is most pronounced in settings with higher ambiguity as measured by $\bar{v}/\underline{v}$. This is intuitive, as in these cases the risks associated with worst-case pricing are accentuated[5]. The same conclusions hold for the constrained learning case. The second period problem reduces to the monopolist's revenue maximization problem for both the unconstrained and constrained learning formulations. But while the effect of the first period revenue diminishes as $\lambda^1 \to 0$, its impact on the second period revenue *does not* in the case with downstream pricing constraints. Table 1 below provides a numerical example of these gains for different parameters.

Table 1: Competitive ratios of Single Period, Constrained Learning, and Unconstrained Learning cases, and respective gains due to learning. $K = 20$ prices uniformly spanning the support.

$(\lambda^1,\ \bar{v}/\underline{v})$	**Single Period CR**	**Const. CR**	**Gain%**	**Unconst. CR**	**Gain%**
(0.1 , 2)	0.595	0.900	51%	0.960	61%
(0.1 , 6)	0.372	0.900	142%	0.937	152%
(0.1 , 10)	0.322	0.900	180%	0.932	190%
(0.4 , 2)	0.595	0.706	19%	0.838	41%
(0.4 , 6)	0.372	0.624	68%	0.749	101%
(0.4 , 10)	0.322	0.614	91%	0.729	126%
(0.7 , 2)	0.595	0.634	7%	0.717	20%
(0.7 , 6)	0.372	0.461	24%	0.560	51%
(0.7 , 10)	0.322	0.427	33%	0.525	63%

4 Learning with Limited Price Experimentation

From a practical viewpoint, firms typically have a limited time and budget for learning, and as a result try to gauge the WtP distribution only at certain price points. This section extends our analysis to cover settings where the firm can experiment with only a small number of price points in the first period.

[5]Analytically, the competitive ratio for the unconstrained learning case is $c_l := \lambda^1\ c\ +\ \lambda^2$, and the relative gain is $(c_l/c - 1) = \lambda^2(c^{-1} - 1)$. Using the lower bound $c^{LB} = (1 + \ln(\bar{v}/\underline{v}))^{-1}$ derived in Proposition 2.1, we see that the gain can be close to $\lambda^2 \ln(\bar{v}/\underline{v})$, which is increasing in $\bar{v}/\underline{v}$.

4.1 Single price in period one

The first case we study is one where the firm can only experiment with one price in period one, which we denote by p_n. It will consequently, observe the fractile of the WtP distribution at that point, denoted by $\bar{f}_n$. The adversary needs to commit to only $\bar{f}_n$ in period one and is free to choose the remainder of the distribution in period two. The firm can use its knowledge of $\bar{f}_n$ in its pricing decision for period two. This is formulated as follows:

$$\begin{aligned} c^* \;=\; \max_{n\in\{1,\ldots,K\}} \min_{\bar{f}_n\in[0,1]} \max_{t} \min_{\bar{f}} \quad & \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2 \sum t_j p_j \bar{f}_j}{(\Lambda^1+\Lambda^2)\max_j\{p_j \bar{f}_j\}} \\ s.t. \quad & \sum_j t_j = 1, \quad t \geq 0 \\ & 1 = \bar{f}_1 \geq \bar{f}_2 \geq \cdots \geq \bar{f}_n \\ & \bar{f}_n \geq \bar{f}_{n+1} \geq \cdots \geq \bar{f}_K \geq 0, \end{aligned} \tag{10}$$

where t_j is the proportion of time spent at p_j in the second period. Note that for the inner maxmin problem, $\bar{f}_n$ is a given constant rather than an optimization variable. The solution to the inner subproblem is of independent interest as it demonstrates how to price and what is the worst-case WtP distribution in settings where the firm has partial demand information in the form of a sales observation at one price point. This is extended later in Section 4.2 to allow for multiple such observations.

The blueprint of our analysis is to show that given $(p_n, \bar{f}_n)$, the problem of the firm decouples into two related subproblems similar to that of Section 2: one on the grid $\{p_1, \ldots, p_{n-1}\}$ with a probability mass of $1-\bar{f}_n$, and the other on the grid $\{p_n, \ldots, p_K\}$ with probability mass of $\bar{f}_n$. For each subproblem, results of Section 2 such as the extreme point optimality for adversary's decision continue to hold. The strategy of the firm is again a price skimming policy for each subinterval $[\underline{v}, p_n)$ and $[p_n, \bar{v}]$ and then p_n to balance the potential revenue loss due to each subinterval.

We start our analysis by noting that the adversary's inner problem in (10) is one of minimizing a quasi-concave function in $\bar{f}$ over a polyhedron, as in Section 2.2, which admits an extreme point optimal solution. However, instead of $C(K,1) = K$ extreme points, the partitioned constraints above admits $C(n-1,1) * C(K-n+1,1) = (n-1)(K-n+1)$ extreme points. For example, one such extreme point corresponds to the $\bar{f}$ vector of the form $((1, 1, \bar{f}_n, \bar{f}_n, \ldots, \bar{f}_n), (\bar{f}_n, \bar{f}_n, \bar{f}_n, 0, 0, \ldots, 0))$, which corresponds to a point mass of size $1-\bar{f}_n$ at price p_{j_1} and a point mass of size $\bar{f}_n$ at price p_{j_2} for $1 \le j_1 < n \le j_2 \le K$. Exploiting this concave minimization structure, and defining

$$c_{j_1,j_2} = \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2(\sum_{j=1}^{j_1} t_j p_j + \sum_{j=j_1+1}^{j_2} t_j p_j \bar{f}_n)}{(\Lambda^1+\Lambda^2)\max\{p_{j_1}, p_{j_2}\bar{f}_n\}}, \quad 1 \le j_1 < n \le j_2 \le K,$$

the problem in (10) can be rewritten as

$$c^* = \max_{n\in\{1,\dots,K\}} \min_{\bar{f}_n\in[0,1]} \max_{\sum t_j=1,\, t\ge 0} \min_{1\le j_1<n\le j_2\le K} c_{j_1,j_2}. \tag{11}$$

That is, the adversary now selects two point masses, the first of size $(1-\bar{f}_n)$ that is positioned in $[\underline{v}, p_n)$, and the second of size $\bar{f}_n$ that is positioned in $[p_n, \bar{v}]$.

Proposition 4.1 *For given p_n, $\bar{f}_n$, and t, the extreme points of the inner minimization problem $\min_{1\le j_1<n\le j_2\le K} c_{j_1,j_2}$ in (11) is characterized by a pair of indices (j_1, j_2) that correspond to positioning of the point masses $1-\bar{f}_n$ at p_{j_1} with $1 \le j_1 < n$ and $\bar{f}_n$ at p_{j_2} with $n \le j_2 \le K$ respectively. There exists an optimal solution to (11) that places $1-\bar{f}_n$ probability at p_1 or $\bar{f}_n$ probability at p_n. Consequently, the optimal choice of (j_1, j_2) is of the form (j_1, n) or $(1, j_2)$.*

Proof: The proof is divided into two cases. Let us first suppose that $p_{j_1} < p_{j_2}\bar{f}_n$ and $j_1 > 1$ at the optimal solution. Then, the optimal ratio for fixed n, $\bar{f}_n$, and t, denoted by $c(n, \bar{f}_n, t)$, is

$$\begin{aligned}
c(n,\bar{f}_n,t) = c_{j_1,j_2} &= \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2(\sum_{j=1}^{j_1} t_j p_j + \sum_{j=j_1+1}^{j_2} t_j p_j \bar{f}_n)}{(\Lambda^1+\Lambda^2)p_{j_2}\bar{f}_n} \\
&\ge \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2(t_1 p_1 + \sum_{j=2}^{j_1} t_j p_j \bar{f}_n + \sum_{j=j_1+1}^{j_2} t_j p_j \bar{f}_n)}{(\Lambda^1+\Lambda^2)p_{j_2}\bar{f}_n} \\
&= \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2(t_1 p_1 + \sum_{j=2}^{j_2} t_j p_j \bar{f}_n)}{(\Lambda^1+\Lambda^2)p_{j_2}\bar{f}_n} \\
&= c_{1,j_2},
\end{aligned} \tag{12}$$

so c_{1,j_2} is also optimal whenever $p_{j_1} < p_{j_2}\bar{f}_n$.

Second, suppose that $p_{j_1} \ge p_{j_2}\bar{f}_n$ and $j_2 > n$ at the optimal solution. Consequently, the optimal ratio is

$$
\begin{aligned}
c(n, \bar{f}_n, t) = c_{j_1,j_2} =& \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2 (\sum_{j=1}^{j_1} t_j p_j + \sum_{j=j_1+1}^{j_2} t_j p_j \bar{f}_n)}{(\Lambda^1 + \Lambda^2) p_{j_1}} \\
\geq& \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2 (\sum_{j=1}^{j_1} t_j p_j + \sum_{j=j_1+1}^{n} t_j p_j \bar{f}_n)}{(\Lambda^1 + \Lambda^2) p_{j_1}} \\
=& c_{j_1,n},
\end{aligned} \tag{13}
$$

and $c_{j_1,n}$ is also optimal whenever $p_{j_1} \geq p_{j_2} \bar{f}_n$. $\square$

As a result the problem can be reduced to

$$
\begin{aligned}
c^* \;=\; \max_{n \in \{1,\dots,K\}} \min_{\bar{f}_n \in [0,1]} \max_{t,c} \quad & c \\
s.t. \quad & c \leq c_{j_1,n} \quad 1 \leq j_1 < n \\
& c \leq c_{1,j_2}, \quad n \leq j_2 \leq K \\
& \sum t_j = 1, \quad t \geq 0 \;.
\end{aligned} \tag{14}
$$

The next proposition characterizes the worst-case $\bar{f}_n$ that the firm can observe and Corollary 4.1 formulates the corresponding problem of choosing p_n to experiment at according to this.

Proposition 4.2 *For fixed p_n, there exists an optimal solution of the outer minimization in (10) with $\bar{f}_n = 0$ or 1. Hence, it is sufficient to restrict attention to K extreme points where the unit probability mass is allocated to a single price.*

Propositions 4.1 and 4.2 together imply that for a given price p_n in period one, the adversary's problem decomposes into two subproblems: a) an upper problem where the unit probability mass is placed at some price $p_j \in [p_n, \bar{v}]$, and b) a lower problem where the unit mass is placed at some price $p_j \in [\underline{v}, p_{n-1}]$. The adversary selects the solution that yields a smaller ratio. In return, the firms' problem in period one is to choose the price point p_n that balance the ratios in subproblems a) and b).

Corollary 4.1 *Let*

$$
c^u(n) = \max_{\sum t_j = 1,\, t \geq 0} \; \min_{n \leq j_2 \leq K} \frac{\Lambda^1 p_n + \Lambda^2 \sum_{j=1}^{j_2} t_j p_j}{(\Lambda^1 + \Lambda^2) p_{j_2}} \quad \textit{and} \quad c^l(n) = \max_{\sum t_j = 1,\, t \geq 0} \; \min_{1 \leq j_1 < n} \frac{\Lambda^2 \sum_{j=1}^{j_1} t_j p_j}{(\Lambda^1 + \Lambda^2) p_{j_1}}
$$

the competitive ratio problem in (10) reduces to

$$c^* \;=\; \max_{n \in \{1,\dots,K\}} \min\{c^u(n), c^l(n)\}\;. \tag{15}$$

Observe that both $c^u(n)$ and $c^l(n)$ are equivalent to simple linear programs as demonstrated in Section 2.2. Furthermore, $c^l(n)$ is directly equivalent to the single period problem in Section 2.2 using only prices $p_1 = \underline{v}$ to p_{n-1}, and has its solution readily available.

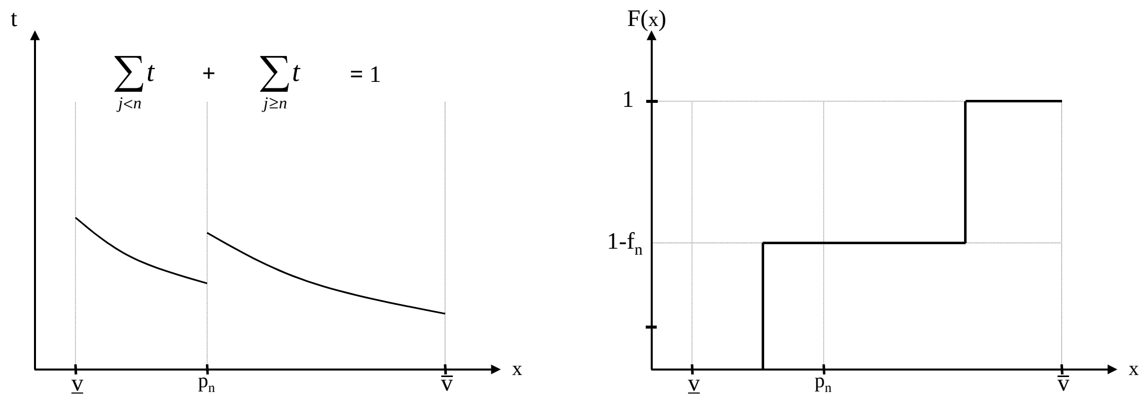

(a) A representative policy (b) A candidate WtP distribution

Figure 1: Structure of the worst-case distribution with one test price

The choice of p_n balances two types of risk: a) the firm faces the risks of lost sales and foregone revenue in period one; b) the exposure faced by the seller in each of the subintervals $[\underline{v}, p_n)$ and $[p_n, \bar{v}]$, which increases with the ambiguity measured by $\underline{v}/p_n$ and $p_n/\bar{v}$. As Λ^1 decreases, i.e., the emphasis shifts on balancing these risks in period two. Hence, as $\Lambda^1 \to 0$ and as $\{p_1, \dots, p_K\}$ becomes a dense covering of $[\underline{v}, \bar{v}]$ with $K \uparrow \infty$, the optimal price at which to experiment is such that the relative ambiguity of the two subproblem becomes equal: specifically optimal p_n is satisfies $\underline{v}/p_n = p_n/\bar{v}$, i.e. optimal p_n is the geometric mean of $\underline{v}$ and $\bar{v}$. This result directly follows from Proposition 2.1.

4.2 Multiple prices and incorporating partial demand information

Let N be the number of price points used in the first period with $1 \le N < K$, and label these prices by $p_{i_1},\ p_{i_2}, \dots, p_{i_N}$. In this subsection we will focus on the following practically important problem: given a testing schedule $\{p_{i_1}, \dots, p_{i_N}\}$ and the associated fractiles $\{\bar{f}_{i_1}, \dots, \bar{f}_{i_N}\}$, how should the firm exploit and incorporate information into its pricing decision? As a byproduct, we will characterize the worst-case market condition based on the observed demand information, which is of interest on its own. Specifically, we will not carry through the full analysis as we did in Section 4.1. We illustrate below that analogues of Propositions 4.1 and 4.2 can be derived and the structure of Corollary 4.1 still holds. The solution, however, to the problem of selecting the optimal set of prices $\{p_{i_1}, \ldots, p_{i_N}\}$ for testing is combinatorial and does not seem to simplify significantly.

For a given set of indices of test prices, $\mathcal{N} = \{i_1, \ldots, i_N\}$, we formulate the problem as:

$$
\begin{aligned}
c_{\mathcal{N}} = \quad \max_{t^1} \min_{\bar{f}_{i_n} \in [0,1]} \max_{t^2} \min_{\bar{f}} \quad & \frac{\Lambda^1 \sum_{i_n \in \mathcal{N}} p_{i_n} \bar{f}_{i_n} t^1_{i_n} + \Lambda^2 \sum_{j=1}^{K} t^2_j p_j \bar{f}_j}{(\Lambda^1 + \Lambda^2) \max_j \{p_j \bar{f}_j\}} \\
s.t. \quad & \sum_{i_n \in \mathcal{N}} t^1_{i_n} = 1, \quad t^1 \geq 0 \\
& \sum_{j=1}^{K} t^2_j = 1, \quad t^2 \geq 0 \\
& 1 = \bar{f}_1 \geq \cdots \geq \bar{f}_K \geq 0 \\
& \bar{f}_j = \bar{f}_{i_n} \quad \forall\, j = i_n \in \mathcal{N}\ .
\end{aligned}
\tag{16}
$$

Note that problem (16) mimics the unconstrained learning problem previously studied, which is indeed a special case of (16) with $N = K$.

The essence of the single price analysis carries through in the following sense. First, the inner maxmin problem that pertains to the second period pricing problem decomposes into $N + 1$ subproblems in intervals $[\underline{v}, p_{i_1}),\ [p_{i_1}, p_{i_2}), \ldots, [p_{i_N}, p_K]$ that can be studied using the results of Section 2: the firm uses a price skimming policy within that interval, and the adversary selects a point mass distribution for the probability that belongs to the respective interval. Second, one can "piece" together the above results to characterize the first period behavior which only requires comparing these $N + 1$ subproblems. Furthermore, complexity-wise, this overall method requires concentrating at only K extreme points in total. Together these results yield the solution to (16).

The main idea in the proof of Proposition 4.1 for the single price analysis is that there exists a price point yielding a maximum revenue rate for a given distribution and information constraint, and allocating the probability mass at all other intervals to the lowest possible price can only improve the objective function of the adversary (i.e. reduce the ratio), because it potentially reduces the numerator with the denominator unchanged (see equations (12) and (13)). The same argument goes through when we have more fractile observations with the same steps for each interval. We state this in the following, the proof of which follows similar steps to Proposition 4.1 and is therefore omitted.

Proposition 4.3 *If fractile information at $N < K$ points are given, one can still restrict attention to K extreme points for the inner minimization problem. For $\eta = 1 \dots N-1$, each extreme point is of the form $f_j = \bar{f}_{i_\eta} - \bar{f}_{i_{\eta+1}}$ for some $i_\eta \leq j < i_{\eta+1}$ and $f_{i_n} = \bar{f}_{i_n} - \bar{f}_{i_{n+1}}$ for all other $n \neq \eta$.*

That is we consider a price interval $[p_{i_\eta}, p_{i_{\eta+1}})$ one at a time and fix the probability mass for all other intervals at the lowest price possible, ie. $f_{i_n} = \bar{f}_{i_n} - \bar{f}_{i_{n+1}}$, $n \neq \eta$, while assuming the mass at this interval is at one of the price points within the interval. Consequently, it is sufficient again to concentrate on a total of K extreme points.

This seemingly simple result is important because the number of extreme points can in general grow exponentially with additional information as explained before. For example, while adding a constraint on the mean of the WtP distribution will increase the number of extreme points to $O(K^2)$, fractile information can be incorporated without effectively increasing the number of extreme points to consider. Furthermore, given the fractile information the resulting competitive ratio problem for the remaining sales horizon can be solved using a simple LP formulation which will be illustrated later in this section.

Proposition 4.4 *For any feasible t^1, there exists an optimal solution for the outer minimization in (16) which occurs at an extreme point of the simplex $1 \geq \bar{f}_2 \geq \cdots \bar{f}_K \geq 0$. In other words, the adversary chooses a distribution which allocates the unit probability mass to a single price in the first period problem.*

Proof: Fix some t^1, let the optimal solution to the inner maximization be t^2 and to the inner minimization be $\bar{f}$, and denote the resulting optimal ratio by $c(t^1) := n(t^1)/d(t^1)$, where $n(t^1)$ and $d(t^1)$ denote the corresponding values of the numerator and the denominator respectively in (16) at the optimal solution. Also, let $j^* := \text{argmax}_j\{p_j \bar{f}_j\}$ be the index of the revenue maximizing price.

First, observe that there exists an optimal distribution with $\bar{f}_j = 0$ for $j > j^*$ for the inner problem. To see this, consider the constraints $1 = \bar{f}_1 \geq \bar{f}_2 \geq \cdots \bar{f}_{j^*} \geq \bar{f}_{j^*+1} \geq \cdots \bar{f}_K \geq 0$. Suppose that for any fixed values of $\bar{f}_1, \dots, \bar{f}_{j^*}$, some of the variables $\bar{f}_{j^*+1}, \dots, \bar{f}_K$ have positive values.

Then, by reducing them to zero, we do not change the value of $p_{j^*}\ \bar{f}_{j^*} = \max_j\{p_j \bar{f}_j\}$, and hence the the value of the denominator in (16), while potentially reducing the value of the numerator. This would yield a lower competitive ratio. It follows that $\bar{f}_{j^*+1} = \ldots \bar{f}_K = 0$ for some optimal solution.

Now, we also show that there exists an optimal distribution with $\bar{f}_{j^*} = 1$ for the inner problem. Suppose that the optimal solution has $\bar{f}_{j^*} < 1$. Then increasing $\bar{f}_{j^*}$ by $\epsilon := 1 - \bar{f}_{j^*}$ would increase the numerator $n(t^1)$ at most by $\epsilon \left(\Lambda^1 \sum_{i_n \le j^*, i_n \in \mathcal{N}} p_{i_n} t^1_{i_n} + \Lambda^2 \sum_{j \le j^*} t^2_j p_j\right)$ while increasing the denominator exactly by $\epsilon\ (\Lambda^1 + \Lambda^2) p_{j^*}$. The new competitive ratio, denoted by c_ϵ, with $\bar{f}_{j^*} = 1$ and $\bar{f}_j = 0$ for $j > j^*$, satisfies

$$c_\epsilon \le \frac{n(t^1)\ +\ \epsilon \left(\Lambda^1 \sum_{i_n \le j^*, i_n \in \mathcal{N}} p_{i_n} t^1_{i_n} + \Lambda^2 \sum_{j \le j^*} t^2_j p_j\right)}{d(t^1)\ +\ \epsilon\ (\Lambda^1 + \Lambda^2) p_{j^*}} \le \frac{n(t^1)}{d(t^1)} = c(t^1)\ , \tag{17}$$

which shows that setting $\bar{f}_{j^*} = 1$ is also optimal. The second inequality above follows from

$$\frac{\epsilon \left(\Lambda^1 \sum_{i_n \le j^*, i_n \in \mathcal{N}} p_{i_n} t^1_{i_n} + \Lambda^2 \sum_{j \le j^*} t^2_j p_j\right)}{\epsilon\ (\Lambda^1 + \Lambda^2) p_{j^*}} \le \frac{\Lambda^1 \sum_{i_n \le j^*, i_n \in \mathcal{N}} p_{i_n} \bar{f}_{i_n} t^1_{i_n} + \Lambda^2 \sum_{j \le j^*} t^2_j p_j \bar{f}_j}{(\Lambda^1 + \Lambda^2) p_{j^*} \bar{f}_{j^*}} = \frac{n(t^1)}{d(t^1)}\ ,$$

as $1 = \bar{f}_1 \ge \bar{f}_2 \ge \cdots \ge \bar{f}_{j^*}$ and $\bar{f}_j = 0$ for $j > j^*$; the inequality holds with equality if and only if $1 = \bar{f}_1 = \bar{f}_2 = \cdots = \bar{f}_{j^*}$. $\square$

Propositions 4.3 and 4.4 together imply that for a given set $\mathcal{N}$ of first period prices chosen by the firm, the problem of the adversary again decomposes into $N+1$ subproblems corresponding to each of the intervals $[\underline{v}, p_{i_1}),\ [p_{i_1}, p_{i_2}), \ldots, [p_{i_N}, p_K]$. In each subproblem, the adversary positions the corresponding probability mass at a single price point that belongs the respective interval in a way that minimizes the ratio.

Corollary 4.2

Let

$$c^0 = \max_{t^1,\, t^2} \ \min_{1 \le j0 < 1_n} \frac{\Lambda^2 \sum_{j \le j0} t^2_j p_j}{(\Lambda^1 + \Lambda^2) p_{j0}}\ ,$$

$$c^n = \max_{t^1,\, t^2} \ \min_{i_n \le j_n < i_{n+1}} \frac{\Lambda^1 \sum_{i_n \le j_n, i_n \in \mathcal{N}} p_{i_n} t^1_{i_n} + \Lambda^2 \sum_{j \le j_n} t^2_j p_j}{(\Lambda^1 + \Lambda^2) p_{j_n}}, \quad n = 1 \ldots N\ ,$$

the competitive ratio problem in (16) reduces to

$$c_{\mathcal{N}} = \min_{n\in\{0,\dots,N\}} \{c^n\} \ . \tag{18}$$

Once again each subproblem c^n is equivalent to a linear program. The next step would ideally be optimizing the set of test prices $\mathcal{N} \subset \{1, \dots, K\}$ in period one. However, this is purely a combinatorial problem that requires numerical techniques, and therefore, is left out of our discussion.

A practically important special case to the problem (16) above can be used to incorporate additional information available to the dynamic pricing problem of Section 2. The overall setting is the same, but instead of a learning period, the fractile information is assumed to be readily available for a subset $\mathcal{N}$ of prices. This limited information could represent an expert opinion, an industry forecast, past experience, or the result of price testing. Mathematically, this problem, which is equivalent to the inner minimax formulation of (16) with $\Lambda^1 = 0$, is given by:

$$\begin{aligned} c_{\mathcal{N}} = \ \max_t \ \min_{\bar{f}} \quad & \frac{\sum_{j=1}^K t_j p_j \bar{f}_j}{\max_j\{p_j \bar{f}_j\}} \\ s.t. \quad & \sum_{j=1}^K t_j = 1, \quad t \geq 0 \\ & 1 = \bar{f}_1 \geq \cdots \geq \bar{f}_K \geq 0 \\ & \bar{f}_j = \bar{f}_{i_n} \quad \forall\ j = i_n \in \mathcal{N} \ . \end{aligned} \tag{19}$$

Fractile information for a subset of prices $\{p_{i_1}, \dots, p_{i_N}\}$ can be incorporated without increasing the complexity of the problem, as explained by Proposition 4.3. The resulting problem can be reduced to an LP with K constraints. Each extreme point identified in Proposition 4.3 corresponds to a linear upper bound constraint on the objective function of the following equivalent LP formulation:

$$\begin{aligned} \max_{t,\, c} \quad & c \\ s.t. \quad & c \leq \frac{t_j p_j \bar{f}_{i_\eta} + \sum_{n\neq\eta,\, i_n\in\mathcal{N}} t_{i_n} p_{i_n} \bar{f}_{i_n}}{\max\{p_j \bar{f}_{i_\eta},\ \max_{n\neq\eta,\, i_n\in\mathcal{N}}\{p_{i_n}\bar{f}_{i_n}\}\}} \qquad \eta = 1\dots N-1, \quad i_\eta \leq j < i_{\eta+1} \\ & \sum_{j=1}^K t_j = 1, \quad t \geq 0 \ . \end{aligned} \tag{20}$$

The solution of this LP provides both a pricing policy recommendation and the corresponding optimal competitive ratio. The actual revenue performance of the policy t is quite good across many common demand functions as illustrated with the numerical examples reported below, even when the fractile/sales information is available at only a few price points.

4.3 Numerical examples

We conclude this section with a set of numerical results that highlight the revenue improvement that is achieved through partial learning under our policy. Our experiments contrast the "no information" policy of Sections 2 to the partial information policy extracted via (20) if one is given fractile information at a set of price points. The fractile data was generated using four common WtP distributions, each of which corresponds to common demand model listed in Table 2. We restricted the WtP to the range $[\underline{v}, \bar{v}]$ for each distribution. For the Normal and Gumbel distributions we extracted the mean as the midpoint of the range and that standard deviation by assuming that the range is equal to $\pm 3\sigma$. For the exponential distribution we assumed that the WtP of a typical consumer is given by $\underline{v} + w$, where w is an exponentially distributed in $[0, \bar{v} - \underline{v}]$ and its rate parameter is selected so that the probability that w lies in that range is 99.5% (this is consistent with the $\pm 3\sigma$ assumption of the Normal distribution).

In each test case, we also compared against a policy that tries to make use of the observed fractile information by first fitting an exponential demand model to this data, and then use this model to compute a static price to be used throughout the sales horizon. The latter heuristic will, of course, turn out to be optimal in test cases that correspond to an underlying exponential WtP distribution, but its performance on the other three test cases will give a rough idea of the performance loss due to the wrong parametrization of the demand model.

WtP distribution	**Demand model**
Uniform	Linear demand
Exponential	Exponential demand
Normal	Probit demand
Gumbel	Logit demand

Table 2: WtP distributions and corresponding demand models.

The three sets of results summarized in Tables 3 - 5 illustrate the performance of the policy derived using the competitive ratio analysis in a variety of settings as we varied the range of the WtP distribution, the number of test prices for which the seller has observed information, and also as we varied the ambiguity of the range information, which is captured by the ratio $\bar{v}/\underline{v}$. The last example reported in Table 6 illustrates the performance of the pricing policy extracted form the competitive ratio analysis when only one price point is tested in period one and as we vary the positioning of the price point within the predefined range $[\underline{v}, \bar{v}]$. Note that the range of WtP distribution for that last example coincides with that in Table 4.

There are several observations to be made. First, although the competitive ratio is conservative, the actual revenue performance of the policy across different distributions is significantly higher. Actually, as the ambiguity ratio $\bar{v}/\underline{v}$ gets smaller the performance of the policy derived from the competitive ratio analysis increases significantly, e.g. from 19.3% to 48.3% between the second and third examples . Second, partial information (and learning) significantly increases both the performance guarantee of the competitive ratio policy and the actual revenue performance across distributions. Third, the revenue performance of the competitive ratio policy with partial information is very good across all distributions. Significant gains are achieved even when experimenting at a small number of price points: sampling at just 3 prices out of 500 achieves in excess of 85% of the maximum achievable revenues under full information across all distributions tested even when the ambiguity ratio is very high. Decreasing the ambiguity ratio, from 500 to 100, and increasing the number of test prices slightly, from 3 to 5 points, increases the guaranteed performance to 95% across all test cases. Fourth, using partial information to fit an incorrect parametric model to the unknown distribution can lead to substantial revenue loss. For example, in Table 3 above, fitting an exponential distribution for the underlying Gumbel distribution results in only 42.4% of the maximum achievable revenues, whereas the competitive ratio policy can use the partial information to capture 85.7%. In fact, competitive ratio policy with partial information outperforms parametric fitting across all distributions, except for the exponential case, for which exponential fitting is optimal.

Finally, in all of our experiments we observed that the performance of the policy extracted via the competitive ratio analysis performed very well even when the test prices did not happen to fall close to the optimal price for the underlying WtP distribution. The results in Table 6 provide

WtP Dist.	CR w/o Info	Rev Perf CR Policy	CR with Fractile Info	Rev Perf CR Pol w Frac.	Rev Perf under Exponential Fit
Uniform	14.7	29.4	71.0	96.0	89.1
Exponential	14.7	39.6	47.1	90.9	100
Normal	14.7	25.2	65.0	87.4	47.7
Gumbel	14.7	24.4	64.8	85.7	42.4

Table 3: Competitive ratio with and without fractile information, and corresponding revenue performance under common WtP distributions. $K = 500$ prices, price grid $[1, 500]$ in increments of 1. Sales (fractile) information at 3 price points: 125, 250, and 375.

WtP Dist.	CR w/o Info	Rev Perf CR Policy	CR with Fractile Info	Rev Perf CR Pol w Frac.	Rev Perf under Exponential Fit
Uniform	19.3	38.2	80.3	98.1	91.8
Exponential	19.3	38.4	80.1	98.1	100
Normal	19.3	32.6	77.5	95.8	54.0
Gumbel	19.3	31.6	77.6	95.3	50.8

Table 4: $K = 100$ prices, price grid $[1, 100]$ in increments of 1. Sales (fractile) information at 5 price points: 16, 33, 50, 66, and 83.

WtP Dist.	CR w/o Info	Rev Perf CR Policy	CR with Fractile Info	Rev Perf CR Pol w Frac.	Rev Perf under Exponential Fit
Uniform	48.3	85.1	84.7	98.0	98.0
Exponential	48.3	65.6	82.5	92.2	100
Normal	48.3	67.5	85.7	96.6	84.7
Gumbel	48.3	66.3	86.0	96.8	81.8

Table 5: $K = 100$ prices, price grid $[51, 150]$ in increments of 1. Sales (fractile) information at 5 price points: 66, 83, 100, 116, and 133.

Test Price (in \$)	20	40	60	80
Uniform	(38, 73)	(50, 87)	(57, 95)	(50, 84)
Exponential	(35, 77)	(35, 75)	(30, 70)	(25, 62)
Normal	(39, 65)	(51, 76)	(51, 74)	(30, 49)
Gumbel	(39, 63)	(52, 75)	(49, 71)	(29, 47)

Table 6: Competitive Ratio and actual revenue performance as % from optimal when the firm only experiments with one price. $K = 100$ prices, price grid $[1, 100]$ in increments of 1.

an illustration of this last comment by focusing at the performance of the policy extracted via the competitive ratio analysis as we vary the position of a single price point for which the seller has fractile information. As a benchmark we note that the competitive ratio and actual revenue performance of the policy based on the competitive ratio analysis without the fractile information was reported in the first and second columns of Table 4. In addition, the optimal prices that would correspond to each of the WtP distributions that we tried were 50 for the Uniform, 19 for the

Exponential, 39 for the Normal and the Gumbel distributions. This observation is robust in the sense that it holds even if the WtP distribution is such that the revenue function is not unimodal, which a structural property that could, of course, help in narrowing down the region in which the optimal may be.

References

V. F. Araman and R. A. Caldentey. Dynamic pricing for non-perishable products with demand learning. *Working Paper, New York University, New York, NY 10012*, 2005.

Y. Aviv and A. Pazgal. Pricing of short life-cycle products through active learning. *Working Paper, Washington University, St. Louis, MO 63130*, 2005.

M. Ball and M. Queyranne. Competitive analysis of on-line booking: Dynamic policies. *Working Paper, University of Maryland, College Park, MD*, 2004.

A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Math. Oper. Res.*, 23(4):769–805, 1998.

D. Bergemann and K. Schlag. Robust monopoly pricing: The case of regret. *Working Paper, Yale University, New Haven, Connecticut 06520-8281*, 2005.

D. Bertsimas and I. Popescu. Optimal inequalities in probability theory: A convex optimization approach. *SIAM Journal of Optimization*, 15(3):780–804, 2005.

D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52:35–53, 2003.

D. Bertsimas and A. Thiele. A robust optimization approach to supply chain manage- ment. *Proceedings of the 14th IPCO*, pages 86–100, 2004.

O. Besbes and A. Zeevi. Blind nonparametric revenue management: Asymptotic optimality of a joint learning and pricing method. *Working Paper, Columbia University, New York, NY 10027*, 2006.

L. El Ghaoui and H. Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM J. on Matrix Analysis and Applications*, 18(4):1035–1064, 1997.

S. Eren and C. Maglaras. Revenue management heuristics under limited market information: A maximum entropy approach. *Working Paper, Columbia University, New York, NY 10027*, 2006.

S. Eren and G. van Ryzin. Product line positioning without marked information. *Working Paper, Columbia University, New York, NY 10027*, 2006.

V. F. Farias and B. Van Roy. Dynamic pricing with a prior on market response. *Working Paper, Stanford University, Stanford, CA 94305*, 2006.

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

V. Gaur and M. L. Fisher. In-store experiments to determine the impact of price on sales. *Production and Operations Management*, 14(4):377–387, 2005.

M. Harris and A. Raviv. A theory of monopoly pricing schemes with demand uncertainty. *American Economic Review*, 71(3):347–365, 1981.

W. T. Huh and P. Rusmevichientong. An adaptive algorithm for multiple-fare-class capacity control problems. *Working Paper, Columbia University, New York, NY 10027*, 2006.

Lan, Gao, M. Ball, and I Karaesmen. Revenue management with limited demand information. *Working Paper, University of Maryland, College Park, MD*, 2006.

E.P. Lazear. Retail pricing and clearance sales. *American Economic Review*, 76:14–32, 1986.

A.E.B. Lim and J.G. Shanthikumar. Relative entropy, exponential utility, and robust dynamic pricing. *Working Paper, University of California Berkeley, CA*, 2004.

M. S. Lobo and S. Boyd. Pricing and learning with uncertain demand. *Working Paper, Duke University, Durham, NC 27708*, 2003.

B.P. Pashigan. Demand uncertainty and sales: a study of fashion and markdown pricing. *American Economic Review*, 78(5):936–953, 1988.

B.P. Pashigan and B. Bowen. Why are products sold on sale? explanations of pricing regularities. *Quarterly Journal of Economics*, pages 1014–38, 1991.

G. Perakis and G. Roels. Regret in the newsvendor model with partial information. *Working Paper, MIT, Cambridge, MA,*, 2004.

G. Perakis and G. Roels. Robust controls for network revenue management. *Working Paper, UCLA, Los Angeles, CA 90095*, 2007.

G. Perakis and A. Sood. Competitive multi-period pricing with fixed inventories. *Working Paper, MIT, Cambridge, MA,*, 2003.

J. Smith. Generalized chebyshev inequalities: Theory and applications in decision analysis. *Operations Research*, 43(5):807825, 1995.

K.T. Talluri and G. van Ryzin. The theory and practice of revenue management. *Kluwer Academic Publishers, Norwell, Mass.*, 2004.

G. van Ryzin and J. McGill. Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels. *Management Science*, 46(6):760–775, 2000.

5 Appendix

Proof of Proposition 3.1: We first establish an upper-bound by interchanging the order of the inner maximization and minimization, and using the minimax inequality. That is, we have that $c \leq \hat{c}$, where

$$\begin{aligned}
\hat{c} \;=\; \max_{t^1} \min_{\bar{f}^1} \min_{\bar{f}^2} \max_{t^2} \quad & \frac{\Lambda^1 \sum t_j^1 p_j \bar{f}_j^1 + \Lambda^2 \sum t_j^2 p_j \bar{f}_j^2}{(\Lambda^1 + \Lambda^2) \max_j p_j \bar{f}_j^2} & \\
s.t. \quad & \sum t_j^i = 1, \quad t^i \geq 0 & i = 1,\ 2 \\
& 1 = \bar{f}_1^i \geq \bar{f}_2^i \geq \cdots \bar{f}_K^i \geq 0 & i = 1,\ 2 \\
& t_j^1 (\bar{f}_j^2 - \bar{f}_j^1) = 0, & \forall j
\end{aligned}$$

However, the information about $\bar{f}^1$ and $\bar{f}^2$ is revealed successively, and $\bar{f}^1$ can be different from $\bar{f}^2$ only at points j where $t_j^1 = 0$, i.e. even if $\bar{f}^1$ can be different from $\bar{f}^2$, this difference does not change the objective function value. Therefore, we can assume $\bar{f}^2$ contains the information of $\bar{f}^1$. Consequently, letting $\bar{f} := \bar{f}^2 = \bar{f}^1$, the formulation of $\hat{c}$ can be written as follows:

$$\begin{aligned} \hat{c} \;=\; & \max_{t^1} \min_{\bar{f}} \max_{t^2} \quad \frac{\Lambda^1 \sum t_j^1 p_j \bar{f}_j + \Lambda^2 \sum t_j^2 p_j \bar{f}_j}{(\Lambda^1 + \Lambda^2) \max_j p_j \bar{f}_j} \\ & \qquad\qquad s.t. \quad \sum t_j^i = 1,\ t^i \geq 0, \quad i = 1,\ 2, \quad 1 = \bar{f}_1 \geq \cdots \bar{f}_K \geq 0. \end{aligned}$$

Now, fix t^1 and $\bar{f}$, and consider the inner maximization. The denominator and the first part of the numerator is fixed, and therefore, the inner maximization is simply equivalent to maximizing the second part of the numerator, i.e. maximizing the second period revenue. Therefore, the upper bound problem is equivalent to:

$$\begin{aligned} \hat{c} \;=\; & \max_{t^1} \min_{\bar{f}} \quad \frac{\Lambda^1 \sum t_j^1 p_j \bar{f}_j + \max_{t^2} \left\{ \Lambda^2 \sum t_j^2 p_j \bar{f}_j \;:\; \sum_j t_j^2 = 1, \quad t^2 \geq 0 \right\}}{(\Lambda^1 + \Lambda^2) \max_j p_j \bar{f}_j} \\ & \qquad s.t. \quad \sum_j t_j^1 = 1,\ t^1 \geq 0, \quad 1 = \bar{f}_1 \geq \cdots \bar{f}_K \geq 0, \end{aligned}$$

As, $\max_{t^2} \left\{ \Lambda^2 \sum t_j^2 p_j \bar{f}_j \;: \sum_j t_j^2 = 1,\ t^2 \geq 0 \right\} = \Lambda^2 \max_j \{p_j \bar{f}_j\}$, the upper bound problem reduces to

$$\hat{c} \;=\; \max_{t^1} \min_{\bar{f}} \left\{ \frac{\Lambda^1 \sum t_j^1 p_j \bar{f}_j}{(\Lambda^1 + \Lambda^2) \max_j \{p_j \bar{f}_j\}} + \frac{\Lambda^2}{(\Lambda^1 + \Lambda^2)} \;: \sum_j t_j^1 = 1,\ t^1 \geq 0,\ 1 = \bar{f}_1 \geq \cdots \bar{f}_K \geq 0 \right\},$$

which is exactly the problem (4), whose solution is also the optimal solution of c and satisfies $t_*^1 > 0$ componentwise as identified by Theorem 3.1 above. In other words, the optimal solution of $\hat{c}$ is (t_*^1, t_*^2) identified by Theorem 3.1, at this point the value c achieves its upper-bound $\hat{c}$, i.e. $c = \hat{c}$. Therefore this point is also optimal for c and $t_*^1 > 0$ componentwise. □

Proof of Lemma 3.1: Fix some t^1, let the optimal solution to the inner minimization be $\bar{f}^*$, and the resulting optimal ratio be $c(t^1) := n(t^1)/d(t^1)$ where $n(t^1)$ and $d(t^1)$ denote the corresponding values of the numerator and the denominator respectively in (8) at the optimal solution.

First observe that $\bar{f}_j = 0$ for $j > j^*$ must hold at the optimal solution for the inner problem. To see this, consider the constraints $1 = \bar{f}_1 \geq \bar{f}_2 \geq \cdots \bar{f}_{j^*} \geq \bar{f}_{j^*+1} \geq \cdots \bar{f}_K \geq 0$. For any fixed values of $\bar{f}_1, \cdots,\ \bar{f}_{j^*}$, if some of the variables $\bar{f}_{j^*+1}, \cdots,\ \bar{f}_K$ have positive values, reducing them does not change the value of $p_{j^*} \ \bar{f}_{j^*} = \max_j\{p_j \bar{f}_j\}$, and hence the the value of the denominator in (8), while strictly reducing the value of the numerator, which would yield a lower competitive ratio than the optimal one, yielding a contradiction.

Now, we show that $\bar{f}_{j^*} = 1$ must hold at the optimal solution of the inner problem, which requires a little more work. Again, assume by contradiction that $\bar{f}_{j^*} < 1$, then increasing $\bar{f}_{j^*}$ by ϵ increases $n(t^1)$ at most by $\epsilon \left(\Lambda^1 \sum_{i \le j^*} t_i^1 p_i + \Lambda^2 \ \left[\sum_{i<j^*} t_i^1 p_i + \left(1 - \sum_{i<j^*} t_i^1 \right) \ p_{j^*} \right] \right)$, while increasing the denominator exactly by $\epsilon \ (\Lambda^1 + \Lambda^2) p_{j^*}$. Therefore, the resulting new ratio, denoted by c_ϵ, satisfies

$$c_\epsilon \le \frac{n(t^1) \ + \ \epsilon \left(\Lambda^1 \sum_{i \le j^*} t_i^1 p_i + \Lambda^2 \ \left[\sum_{i<j^*} t_i^1 p_i + \left(1 - \sum_{i<j^*} t_i^1 \right) \ p_{j^*} \right] \right)}{d(t^1) \ + \ \epsilon \ (\Lambda^1 + \Lambda^2) p_{j^*}} < \frac{n(t^1)}{d(t^1)} = c(t^1) \ , \quad (21)$$

which results in a contradiction. The strict inequality above holds because, in general we have,

$$\begin{aligned} & \frac{\epsilon \left(\Lambda^1 \sum_{i \le j^*} t_i^1 p_i + \Lambda^2 \ \left[\sum_{i<j^*} t_i^1 p_i + \left(1 - \sum_{i<j^*} t_i^1 \right) \ p_{j^*} \right] \right)}{\epsilon \ (\Lambda^1 + \Lambda^2) p_{j^*}} \\ & \le \frac{n(t^1)}{d(t^1)} = c(t^1) = \frac{\Lambda^1 \sum_{i \le j^*} t_i^1 p_i \bar{f}_i + \Lambda^2 \ \left[\sum_{i<j^*} t_i^1 p_i \bar{f}_i + \left(1 - \sum_{i<j^*} t_i^1 \right) \ p_{j^*} \ \bar{f}_{j^*} \right]}{(\Lambda^1 + \Lambda^2) p_{j^*} \ \bar{f}_{j^*}} \ , \end{aligned}$$

as $1 = \bar{f}_1 \ge \bar{f}_2 \ge \cdots \ge \bar{f}_{j^*}$ and the inequality holds with equality iff $1 = \bar{f}_1 = \bar{f}_2 = \cdots = \bar{f}_{j^*}$ which is not possible by the contradictory assumption. $\square$

Proof of Proposition 4.2: Suppose that for fixed n, the adversary selects $\bar{f}_n \in (0, 1)$.Using the result of Proposition 4.1, if $p_{j_1} < p_{j_2} \bar{f}_n$, then for any vector t, the competitive ratio is of the form

$$c(n, \bar{f}_n, t) = c_{1,j_2} = \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2 (t_1 p_1 + \sum_{j=2}^{j_2} t_j p_j \bar{f}_n)}{(\Lambda^1 + \Lambda^2) p_{j_2} \bar{f}_n},$$

for some $j_2,\ n \le j_2 \le K$. But, for fixed n and $\bar{f}_n \in (0,1)$,

$$
\begin{aligned}
c(n,\bar{f}_n,t) =& \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2(t_1 p_1 + \sum_{j=2}^{j_2} t_j p_j \bar{f}_n)}{(\Lambda^1+\Lambda^2) p_{j_2} \bar{f}_n} \\
=& \frac{\Lambda^1 p_n}{(\Lambda^1+\Lambda^2) p_{j_2}} + \frac{\Lambda^2}{(\Lambda^1+\Lambda^2)} \frac{t_1 p_1}{p_{j_2}\bar{f}_n} + \frac{\Lambda^2}{(\Lambda^1+\Lambda^2)} \frac{\sum_{j=2}^{j_2} t_j p_j}{p_{j_2}} \\
\ge& \frac{\Lambda^1 p_n}{(\Lambda^1+\Lambda^2) p_{j_2}} + \frac{\Lambda^2}{(\Lambda^1+\Lambda^2)} \frac{t_1 p_1}{p_{j_2}} + \frac{\Lambda^2}{(\Lambda^1+\Lambda^2)} \frac{\sum_{j=2}^{j_2} t_j p_j}{p_{j_2}} \\
=& c(n,1,t) \quad \forall t,\ j_2\ .
\end{aligned} \tag{22}
$$

That is, whenever it is optimal for the adversary to place $\bar{f}_n \in (0,1)$ at price p_{j_2} satisfying $p_{j_1} < p_{j_2}\bar{f}_n$, it is also optimal to place the whole unit mass at p_{j_2}, considering $p_{j_1} < p_{j_2}$ for all $j_1,\ j_2$ as $1 \le j_1 < n \le j_2 \le K$.

If $p_{j_1} \ge p_{j_2}\bar{f}_n$, for any vector t the competitive ratio is of the form

$$
c(n,\bar{f}_n,t) = c_{j_1,n} = \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2(\sum_{j=1}^{j_1} t_j p_j + \sum_{j=j_1+1}^{n} t_j p_j \bar{f}_n)}{(\Lambda^1+\Lambda^2) p_{j_1}},
$$

for some $j_1,\ 1 \le j_1 \le n-1$. And, for fixed n and $\bar{f}_n \in (0,1)$,

$$
\begin{aligned}
c(n,\bar{f}_n,t) =& \frac{\Lambda^1 p_n \bar{f}_n + \Lambda^2(\sum_{j=1}^{j_1} t_j p_j + \sum_{j=j_1+1}^{n} t_j p_j \bar{f}_n)}{(\Lambda^1+\Lambda^2) p_{j_1}} \\
\ge& \frac{\Lambda^2 \sum_{j=1}^{j_1} t_j p_j}{(\Lambda^1+\Lambda^2) p_{j_1}} \\
=& c(n,0,t) \quad \forall t,\ j_1\ .
\end{aligned} \tag{23}
$$

That is, whenever it is optimal for the adversary to place $\bar{f}_n \in (0,1)$ at price p_{j_2} satisfying $p_{j_1} \ge p_{j_2}\bar{f}_n$, it is also optimal to place the whole unit mass at p_{j_1}. That is, (22) and (23) show that irrespective of the value of t and the indices $j_1,\ j_2$ in the inner problems, there exists a weakly dominant strategy for the adversary that allocates the whole probability mass to a single price. $\square$