Robust Dynamic Pricing With Strategic Customers

Yiwei Chen [*]   Vivek F. Farias [†]

May 7, 2015

Abstract

We consider the canonical problem of revenue management (RM) wherein a seller must sell an inventory of some product over a finite horizon via an anonymous, posted price mechanism. Unlike typical models in RM, we assume that customers are forward looking. In particular, customers arrive randomly over time, and strategize about their time of purchase. The private valuations of these customers decay over time and the customers incur monitoring costs; both the rate of decay and these monitoring costs are private information. Moreover, customer valuations and monitoring costs are potentially correlated. This setting has proven to be a difficult one for the design of optimal dynamic mechanisms heretofore. Optimal pricing schemes – an almost necessary mechanism format for practical RM considerations – have been similarly elusive.

We propose a class of pricing policies, and a simple to compute policy within this class, that is guaranteed to achieve expected revenues that are at least within 29% of those under an optimal (not necessarily posted price) dynamic mechanism. Moreover, the seller can compute this pricing policy without any knowledge of the distribution of customer discount factors and monitoring costs. Our scheme can be interpreted as solving a dynamic pricing problem for myopic customers with the additional requirement of a novel 'restricted sub-martingale constraint' on prices. Numerical experiments suggest that the policy is, for all intents, near optimal.

1. Introduction

The discipline of Revenue or Yield Management (RM) has, over the last two decades come to occupy a place of prominence in applied Operations Research. Today, applications of revenue management run the gamut from dynamic pricing in the airline industry to hospitality to retail. The following dynamic pricing problem is one of the central problems in revenue management: A seller is endowed with an inventory of a single product that she must sell over a finite horizon. She cannot acquire additional inventory over the course of the horizon and unsold inventory has negligible salvage value. Customers arrive randomly over the course of the selling horizon with the intent of purchasing a single unit of the product. Should the posted price upon a customers arrival exceed his valuation he leaves the system for good; otherwise he purchases a single unit of the product. The seller seeks to dynamically adjust prices with a view to maximizing expected revenue. For typical assumptions on the customer arrival process – assuming, for instance, a renewal process – this problem admits a tractable dynamic programming solution. Despite its simplicity, the canonical nature of this problem serves to highlight an important view of the role of dynamic pricing in revenue management as a *tool to hedge against uncertain demand.*

---

[*]YC is with School of Business, Renmin University of China. Email: {chenyiwei@rbs.org.cn}

[†]VF is with the Operations Research Center and Sloan School of Management, MIT. Email: {vivekf@mit.edu}

In the past decade, it has become amply clear that for a number of principal RM applications, assuming myopic customer behavior, as in the problem above, is no longer a tenable assumption. In spite of this realization, optimal dynamic mechanisms proposed for the version of this central problem that assumes *strategic* customers, face the following critique:

1. They do not admit pure pricing implementations requiring instead devices such as lotteries or end of season 'fire-sale' auctions. This typically rules out applying these mechanisms in scenarios where an anonymous posted price mechanism is the norm (unfortunately, the majority of RM applications).

2. They require the seller to calibrate a model of the customers strategicity, by learning for instance, inter-temporal preferences or search costs. This latter learning problem is non-trivial given the naturally censored data the seller has access to. The customer utility models assumed also typically place strong restrictions on inter-temporal preferences.

3. These mechanisms frequently impute sophisticated purchase timing decisions in equilibrium that are arguably as untenable as the myopic assumption given the burden they place on the customer from a computational and data standpoint.

The present paper seeks to make progress on these fronts at the expense of optimality. In particular, we propose a class of dynamic pricing policies which may be interpreted as solving the simple dynamic pricing problem for myopic customers with the additional restriction that the pricing policy satisfy what we call a 'restricted sub-martingale constraint'. We dub such policies 'robust dynamic pricing' policies and exhibit a simple to compute policy within this class with attractive properties:

1. Computing a robust policy require minimal data on customers beyond what is already required by the standard dynamic pricing problem assuming myopic customers. The specific class of robust policies we compute require *no* additional data.

2. Robust pricing policies induce customers to behave myopically under mild assumptions on customer utility.

3. We exhibit a simple to compute robust pricing policy that is guaranteed to garner revenues that are within 29% of those garnered under the optimal dynamic mechanism.

In addition to the features above, numerical results suggest that the performance of our robust pricing policy can be expected to be substantially superior to what the uniform theoretical guarantee we prove suggests. These numerical experiments also show that the loss in revenue due to an incorrectly calibrated, but otherwise optimal, dynamic mechanism can be substantial over and above the issues raised earlier.

In a nutshell, the present paper provides a tractable, provably robust approach to dynamic pricing in the face of forward looking customers. The approach is robust in that it provides revenue guarantees while making minimal assumptions of customers' inter temporal utilities and search costs, and requires minimal data about the same.

The remainder of the paper is organized as follows: Section 1.1 provides a brief literature review. Section 2 presents our model: we introduce the notion of a dynamic pricing policy, model customer utilities, and define an optimal dynamic mechanism benchmark. Section 3 introduces robust dynamic pricing policies. We state our main theoretical results on the properties of these

policies in this section. Sections 4 and 5 present our analysis and finish with a proof of our uniform performance guarantee. Section 6 complements our analysis with a brief numerical study. Section 7 concludes.

1.1. Literature Review

Revenue management is today a robust area of study with applications ranging from traditional domains such as airline and hospitality pricing to more modern ones, such as financial services. The text by Talluri and van Ryzin [2004] provides an excellent (if now, slightly dated) overview of this area. Gallego and van Ryzin [1994] is a foundational revenue management paper; the present paper effectively studies the same problem but allowing for forward looking customers.

Stylized RM Models: The last decade has seen a good amount of interest in understanding the impact forward looking or 'strategic' customers have on revenue management approaches. The majority of this work follows a stylized vain with the objective of deriving qualitative insights. This line of work typically considers a two period model, with a price change (typically a markdown) occurring in the second period, and seeks to characterize a variety of effects: Aviv and Pazgal [2008] study comparative statics with respect to heterogeneity in valuations and their decay over time; they study the efficacy of announced discounts versus contingent pricing, and finally, show that the revenue loss from incorrectly assuming customers are myopic when, in fact, they are forward looking can be large. Su and Zhang [2008] are concerned with the question of optimal ordering levels for a seller facing forward looking customers. Cachon and Swinney [2009] consider a mix of myopic, 'bargain' hunting and strategic customers, and study the impact of this mix on operational issues such as initial ordering decisions and 'quick response' replenishment. Like Aviv and Pazgal [2008] they also study the efficacy of announced discounts versus contingent pricing but find the latter to be substantially more beneficial in their model. These papers are a representative sample of this body of work; Aviv and Vulcano [2012] provide a comprehensive review of this branch of RM research.

Dynamic Mechanism: Closer to the spirit of the present paper, is research that applies dynamic mechanism design ideas to RM with forward looking customers. An early paper in this regard is Vulcano et al. [2002]; these authors consider impatient (but strategic) customers arriving sequentially over a finite horizon and propose running a modified second price auction in each period (as opposed to dynamic pricing).

An excellent paper by Gallien [2006] provides what is perhaps the first tractable dynamic pricing algorithm for a non-trivial revenue management model with forward looking customers. The model he considers is the discounted, infinite horizon variant of the canonical RM model, and he shows that the optimal dynamic mechanism can be implemented as a dynamic pricing policy in this model. A limitation in this paper is the assumption of an infinite horizon and the delicate requirement that the seller's discount rate matches that of every customer (i.e. there is no heterogeneity in buyers' inter-temporal preferences and these preferences are effectively common knowledge). More recently, Board and Skrzypacz [2010] consider a discrete time version of the same model, and assuming a *finite* horizon, compute the optimal dynamic mechanism. Board and Skrzypacz [2010] also require that all customers discount at a homogenous rate that is common knowledge. While they do solve the finite horizon RM problem, the mechanism they propose is no longer a purely dynamic pricing mechanism but requires an end-of-season 'clearing' auction.

Pai and Vohra [2013] consider a substantially more general model of (finite horizon) RM with

forward looking customers. Customers in their model have heterogenous 'deadlines' as opposed to discounting. When these deadlines are known to the seller, the authors characterize the optimal mechanism completely and show that is satisfies an elegant 'local' dependence on customer reports. On the other hand, when deadlines are private information, the authors illustrate that the optimal dynamic mechanism is substantially harder to characterize. In light of this work, it is interesting to note that both Gallien [2006] and Board and Skrzypacz [2010] compute the optimal dynamic mechanism while requiring that the customer discount rate (which one may think of as the mean of an exponentially distributed, random time until departure from the system) is common knowledge, which is restrictive. It is worth contrasting the present paper with the aforementioned mechanism design research:

1. We allow for customers' discount factors to be heterogeneous and private information, akin to the the hard 'unknown deadlines' version of the problem studied by Pai and Vohra [2013]). In addition, we assume that customers have a 'monitoring cost' and allow this cost to be correlated with their valuation. This is a rich model.

2. We consider a finite horizon problem like Board and Skrzypacz [2010] and Pai and Vohra [2013]. This makes our model relevant to the vast majority of RM applications (in contrast with the assumption of a discounted, infinite horizon case as in Gallien [2006]).

3. We provide a mechanism that enjoys a constant factor approximation guarantee relative to the optimal mechanism for our setting. We are unable to compute the optimal dynamic mechanism. Given the substantially richer model at hand, and the indication that the mechanism design problem is hard when 'deadlines' are unknown (Pai and Vohra [2013]), this is not surprising.

4. Our mechanism can be implemented as a simple anonymous posted price mechanism; it constitutes a dynamic pricing policy for the seller. In contrast, neither Board and Skrzypacz [2010] nor Pai and Vohra [2013] provide dynamic pricing mechanisms; the former requires an end of season 'clearing' auction.

2. Model

We are concerned with a seller who is endowed with x_0 units of inventory of a single product, which she must sell over the finite selling horizon $[0, T]$ via an anonymous posted price mechanism, all of which is common knowledge. We denote the price posted at time t by π_t. We denote the inventory process by X_t and the corresponding sales process by N_t; $N_t = x_0 - X_t$. We require that π_t depend only on the history of the pricing and sales process[1].

Customers arrive over this period according to a Poisson process of rate λ[2]. A customer arriving at time t is endowed with a valuation, v, a time discount factor, α, and a monitoring cost θ, all non-negative. We denote by ϕ, the 'type' of an arriving customer which we understand to be the tuple

$$\phi \triangleq (t, v, \theta, \alpha).$$

Where needed we will make the dependence of each component on ϕ explicit. After making a purchase decision, customers exit the system. Assume that such a customer chooses to delay

---

[1]More formally, we require π_t to be left continuous, and adapted to $\mathcal{F}_{t-}$ where $\mathcal{F}_t = \sigma(\pi^t, X^t)$.
[2]Our analysis extends to non-homogenous processes.

4

making a purchase decision to time $\tau_\phi \geq t$, and define the tuple $y_\phi \triangleq (\tau_\phi, a_\phi, p_\phi)$, where $p_\phi = \pi_{\tau_\phi}$. If the seller has inventory to allocate [3] and if the allocation provides the customer greater utility than no allocation then $a_\phi = 1$; otherwise $a_\phi = 0$. Such a customer garners utility

$$U(\phi, y_\phi) = a_\phi \left(e^{-\alpha(\tau_\phi - t)} v - p_\phi \right) - \theta(\tau_\phi - t).$$

We assume that a customer's type ϕ is private information, drawn from a distribution that is common knowledge. For the sorts of RM applications alluded to in the introduction, heterogeneity in α allows us to capture heterogeneity in customers' aversion to the risk of not obtaining the product while θ parameterizes the cost he incurs in monitoring prices. We denote by $\underline{\theta}$ a lower-bound on the monitoring cost of a customer; this quantity is potentially zero. We denote the marginal distribution (c.d.f.) of product valuations, v, by $F(\cdot)$ and the corresponding p.d.f. by $f(\cdot)$. We denote $\bar{F}(\cdot) \triangleq 1 - F(\cdot)$. We assume that a customer valuation v is independent of his discount factor α. We make a standard assumption on the valuation distribution:

Assumption 1. $v - \frac{\bar{F}(v)}{f(v)}$ is non-decreasing in v and has a non-negative root, v^*.

Customers are forward looking and employ (symmetric) stopping rules contingent on their type that constitute a symmetric Markov Perfect equilibrium. In particular, for customer type ϕ, τ_ϕ is a stopping rule with respect to the filtration generated by the price process, $\mathcal{P}_t$, and solves[4] the optimal stopping problem

$$\sup_{\tau \geq t_\phi} \mathsf{E} \left[U(\phi, \tau) | \mathcal{P}_{t_\phi} \right],$$

where the expectation assumes that other customers use a symmetric stopping rule.

Our goal in this paper is to construct a price process π_t, and exhibit a corresponding stopping rule τ^π to 'maximize' the seller's expected revenue

$$J_{\pi, \tau^\pi}(x_0, T) = \mathsf{E} \left[\int_0^{\hat{\tau} \wedge T} \pi_t dN_t \right],$$

where $\hat{\tau} = \inf\{t : X_t = 0\}$. We will not characterize an optimization problem to find an optimal such dynamic pricing policy, Rather, we will measure the performance of the robust dynamic pricing algorithm that is the subject of this paper via-a-vis an optimal dynamic mechanism benchmark that we discuss next.

2.1. An Optimal Dynamic Mechanism Benchmark

We denote by $h^t \triangleq \{\phi : t_\phi \leq t\}$ the set of customers (or more carefully, customer types) that arrive prior to time t. We restrict ourselves to direct mechanisms.

A mechanism specifies an allocation and payment rule that we encode as follows: customer ϕ is assigned

$$y_\phi \triangleq (\tau_\phi, a_\phi, p_\phi),$$

where $\tau_\phi \geq t_\phi$ is the time of allocation, a_ϕ is an indicator for whether or not a unit of the product is allocated and p_ϕ is the price paid by the customer. Note that y_ϕ depends on h^T. Denote by

---

[3] Multiple customers revealing themselves to the seller at the same time are allocated inventory in random order.

[4] We will later demonstrate existence of such an equilibrium stopping rule for a specific class of pricing policies. We do not prove existence in general.

$y^t \triangleq \{y_\phi : \tau_\phi \leq t\}$ the set of decisions made up to time t. Finally denote the seller's information set by $\mathcal{H}_t$, the filtration generated by the customer reports made up to time t and assignment decision prior to time t. Specifically, $\mathcal{H}_t = \sigma\left(h^t, y^{t-}\right)$. A *feasible mechanism* satisfies the following properties:

1. Causality: τ_ϕ is a stopping time with respect to the filtration $\mathcal{H}_t$. Moreover, a_ϕ and p_ϕ are $\mathcal{H}_{\tau_\phi}$-measurable.

2. Limited Inventory: The seller cannot allocate more units of product than her initial allocation: $\sum_{\phi \in h^T} a_\phi \leq x_0$, a.s.

3. No Participation Fee: $p_\phi = 0$ if $a_\phi = 0$.

We denote by $\mathcal{Y}$, the class of all such rules, y^T. The seller collects total revenue

$$\Pi\left(y^T\right) \triangleq \sum_{\phi \in h^T} p_\phi,$$

whereas the utility garnered by customer ϕ is $U(\phi, y_\phi)$. The utility garnered by customer ϕ when he reports his true type as $\hat{\phi}$ is then given by $U(\phi, y_{\hat{\phi}})$, where customer ϕ can only reveal his arrival no earlier than his true arrival (i.e., $t_{\hat{\phi}} \geq t_\phi$), and $y_{\hat{\phi}}$ depends on $h^T \backslash \{\phi\} \cup \{\hat{\phi}\}$.

The seller now faces the following optimization problem that seeks to find an optimal dynamic mechanism.

$$\begin{aligned}
\max_{y^T \in \mathcal{Y}} \quad & \mathsf{E}\left[\Pi\left(y^T\right)\right] \\
\text{subject to} \quad & \mathsf{E}_{-\phi}\left[U(\phi, y_\phi)\right] \geq \mathsf{E}_{-\phi}\left[U(\phi, y_{\hat{\phi}})\right], \forall \phi, \hat{\phi}, \text{ s.t. } t_{\hat{\phi}} \geq t_\phi \quad \text{(IC)} \\
& \mathsf{E}_{-\phi}\left[U(\phi, y_\phi)\right] \geq 0 \qquad\qquad\quad , \forall \phi. \quad\quad\quad \text{(IR)}
\end{aligned} \tag{1}$$

Denote by $J^*(x_0, T)$ the optimal value obtained in the problem above. We have the following result[5] illustrating that this constitutes an interesting benchmark:

Lemma 1. *(Valid Benchmark) For any pricing policy (π, τ^π), we have that*

$$J_{\pi, \tau^\pi}(x_0, T) \leq J^*(x_0, T).$$

Proof. Consider the class of *pricing* mechanisms that form a subset of $\mathcal{Y}^p$, where for a given pricing policy π_t, $p_\phi = \pi_{\tau_\phi}$, and $a_\phi = 1$ only if doing so yields the buyer a higher utility, and if inventory is available. Now consider the optimization problem:

$$\begin{aligned}
\max_{y^T \in \mathcal{Y}^p} \quad & \mathsf{E}\left[\Pi\left(y^T\right)\right] \\
\text{subject to} \quad & \mathsf{E}_{-\phi}\left[U(\phi, y_\phi)|\mathcal{P}_{t_\phi}\right] \geq \mathsf{E}_{-\phi}\left[U(\phi, y_{\hat{\phi}})|\mathcal{P}_{t_\phi}\right], \text{ a.s.}, \forall \phi, \hat{\phi}, \text{ s.t. } t_{\hat{\phi}} \geq t_\phi \quad \text{(IC)} \\
& \mathsf{E}_{-\phi}\left[U(\phi, y_\phi)\right] \geq 0 \qquad\qquad\qquad\quad , \forall \phi. \quad\quad\quad\quad \text{(IR)}
\end{aligned} \tag{2}$$

Denote by $J^{\pi^*}(x_0, T)$ the optimal value for this problem. Observe now that for any pricing policy (π, τ^π), we must have $J_{\pi, \tau^\pi}(x_0, T) \leq J^{\pi^*}(x_0, T)$. Specifically, given a policy (π, τ^π), consider the mechanism y^π where the seller commits to 'simulating' each customers stopping rule, i.e. use $\tau_\phi = \tau_\phi^\pi$. Given the definition of τ_ϕ, (IC) is satisfied in (2). But (1) is a relaxation of (2), so $J^{\pi^*}(x_0, T) \leq J^*(x_0, T)$ completing the proof. ∎

---

[5]The result is straightforward; we prove it since a standard revelation principle Lemma for this setting does not appear to be available in the literature.

It is worth pausing to discuss two salient facts pertinent to the formulation above:

1. The formulation allows for general mechanisms. As our objective is to produce a benchmark, this generality is desirable, as it will imply a guarantee among a much broader class of mechanisms than those that rely purely on anonymous posted prices.

2. The formulation requires truth telling be the best response in expectation over all possible customers arrival process. This is weaker than dominant strategies (as in Gallien [2006]), as well as weaker than the the requirement placed on the stopping rules assumed when a pricing mechanism is employed (which allowed customers to observe the price history).

3. Robust Dynamic Pricing

This section presents a robust dynamic pricing policy $\{\pi_t\}$ that induces customers to behave myopically, and that guarantees the seller expected revenues that are within a constant factor of the optimal mechanism benchmark, $J^*(x_0, T)$.

Specifically, we define a feasible set of pricing policies that satisfy an additional 'robustness' constraint. Let $\mathcal{F}_t = \sigma(\pi^t, X^t)$ and define by $\mathcal{G}_t = \mathcal{F}_{t-}$ the filtration yielded by the left limit of $\mathcal{F}_t$. We require:

1. π_t is left-continuous and adapted to $\mathcal{G}_t$.

2. π_t satisfies a constraint we dub the 'restricted sub-martingale' constraint. Specifically, for all t such that $X_{t-} > 0$, we require:

$$ \mathsf{E}\left[(\pi_t - \pi_{t'})^+ \, |\mathcal{G}_t \right] \leq \underline{\theta}(t' - t) \tag{3} $$

for all $t' \geq t$ where the expectation assumes that all customers behave myopically.

3. $\pi_t = \infty$ if $X_{t-} = 0.$[6]

Denote by Π the set of all processes satisfying the three constraints above. We then seek to solve the following dynamic optimization problem:

$$ \hat{J}^*(x_0, T) \triangleq \sup_{\{\pi_t\} \in \Pi} \quad \mathsf{E}\left[\int_0^{\hat{\tau} \wedge T} \pi_t dN_t \right] \tag{4} $$

where N_t is a point process with instantaneous rate $\lambda \bar{F}(\pi_t)$; see Brémaud [1981]. Notice that this optimization problem does not consider any strategic behavior on the part of customers. The only aspect in which it differs from the 'typical' revenue revenue management problem is the constraint placed on sample paths of the pricing policy via the restricted sub-martingale constraint (3).

Motivation: In the absence of the restricted sub-martingale constraint (3), the dynamic optimization problem above is identical to what one may consider the canonical RM problem Gallego and van Ryzin [1994]. The second constraint implies

$$ \mathsf{E}\left[\pi_{t'}|\mathcal{G}_t\right] \geq \pi_t - \underline{\theta}(t' - t). $$

---

[6]We adopt the convention $\infty \cdot 0 = 0$.

This allows an interesting interpretation of the constraint. If customers have *no* monitoring cost whatsoever, this constraint requires the pricing process to be a submartingale. As monitoring costs grow higher, this constraint grows weaker. In the limit of infinite monitoring costs, the constraint is vacuous (and we are back to the canonical RM problem with myopic customers as one might expect). Consequently, the constraint limits the extent of the 'price drop' a customer arriving to the system may hope to gain from waiting to purchase. The extent of this limitation grows stronger as it becomes cheaper for customers to wait.

3.1. Performance Guarantee for Robust Pricing

We present here our principle results for the robust pricing policy. First, we establish an equilibrium stopping rule for customers when the seller follows a robust pricing policy; specifically we show that customers behave myopically:

Lemma 2. *(Myopia) Assume that the seller adopts a robust dynamic pricing policy, and further, that all customers of type $\hat{\phi} \neq \phi$ behave myopically: that is they follow the stopping rule $\tau_{\hat{\phi}} = t_{\hat{\phi}}$. Then, ϕ's best response is to use the stopping rule $\tau_\phi = t_\phi$.*

Proof. Now, since the inter arrival times of customers are exponential (and so, memoryless), and moreover, since $\mathcal{F}_{t_\phi} = \mathcal{G}_{t_\phi}$ a.s. when customer ϕ chooses to not make a purchase at his time of arrival (and consequently does not reveal himself), we have that customer ϕ's best response may be calculated by solving the optimization problem:

$$\max_{\tau_\phi} \mathsf{E}\left[U(\phi, \tau_\phi) | \mathcal{G}_{t_\phi}\right].$$

We will show that $U(\phi, t)$ is a $\mathcal{G}_t$-super-martingale on $t \geq t_\phi$ when $X_{t_\phi-} > 0$; if $X_{t_\phi-} = 0$, the claim of the lemma is trivial. Doob's optional sampling theorem then immediately implies that

$$U(\phi, t_\phi) \geq \max_{\tau_\phi} \mathsf{E}\left[U(\phi, \tau_\phi) | \mathcal{G}_{t_\phi}\right]$$

which is the result. To finish the proof, we show that $U(\phi, t)$ is a $\mathcal{G}_t$-super-martingale on $t \geq t_\phi$. We have, for $t \geq t' \geq t_\phi$:

$$
\begin{aligned}
\mathsf{E}\left[U(\phi, t) | \mathcal{G}_{t'}\right] &= \mathsf{E}\left[\left(e^{-\alpha_\phi(t-t_\phi)} v_\phi - \pi_t\right)^+ \Big| \mathcal{G}_{t'}\right] - \theta_\phi(t - t') - \theta_\phi(t' - t_\phi) \\
&\leq \left(e^{-\alpha_\phi(t-t_\phi)} v_\phi - \pi_{t'}\right)^+ + \mathsf{E}\left[(\pi_{t'} - \pi_t)^+ \big| \mathcal{G}_{t'}\right] - \theta_\phi(t - t') - \theta_\phi(t' - t_\phi) \\
&\leq \left(e^{-\alpha_\phi(t-t_\phi)} v_\phi - \pi_{t'}\right)^+ - \theta_\phi(t' - t_\phi) \\
&= U(\phi, t').
\end{aligned}
$$

where the second inequality follows from the restricted sub-martingale constraint. This completes the proof. ∎

Denote by $\hat{\pi}^*$ an optimal solution to (4)[7]. The previous Lemma shows that (an) equilibrium stopping rule for customers facing such a pricing policy is the myopic rule $\tau_\phi = t_\phi$. We next present the main performance guarantee for this paper. Specifically, we show that the optimal robust pricing policy guarantees revenues that are within a constant factor of the revenue under the optimal dynamic mechanism benchmark presented in the preceding section:

---

[7]The optimal policy may not exist, in which case one could consider an ϵ-optimal policy for arbitrary ϵ.

Theorem 1. *Let $\hat{\pi}^*$ be an optimal robust pricing policy. Moreover, denote by $\tau^{\hat{\pi}^*}$ the corresponding (myopic) stopping rule $\tau_\phi^{\hat{\pi}^*} = t_\phi$. We then have that*

$$J_{\hat{\pi}^*, \tau^{\hat{\pi}^*}}(x_0, T) \geq 0.29 J^*(x_0, T).$$

The next two sections are dedicated to establishing this theorem. In anticipation of these sections, however, we find it useful to point out two salient features of our proof of this theorem:

1. We show, in fact, that the guarantee above holds for a sub-optimal robust pricing policy that can be calculated tractably. This sub-optimal policy can be interpreted as the optimal policy for an infinite horizon dynamic pricing problem with a certain 'optimized' discount rate.

2. The (sub-optimal) policy used to establish our result requires *no* knowledge of $\underline{\theta}$, so that the information requirements of this policy are *identical* to the information requirements of the dynamic pricing problem with myopic customers.

Before moving on to the proof of Theorem 1, we take a brief detour to consider the task of finding good robust pricing policies (where by good, we mean any policy satisfying Theorem 1). This is necessitated by the fact that computing the optimal robust dynamic pricing policy does not appear to be an easy task.

3.2. Computable Robust Pricing Policies Satisfying Theorem 1

Define $\hat{\pi}_\beta^* : \mathbb{N}_+ \cup \{0\} \to \mathbb{R}_+ \cup \{\infty\}$ according to

$$\hat{\pi}_\beta^*(x) - \frac{\bar{F}(\hat{\pi}_\beta^*(x))}{f(\hat{\pi}_\beta^*(x))} = \hat{J}_\beta^*(x) - \hat{J}_\beta^*(x-1)$$

where

$$\beta \hat{J}_\beta^*(x) = \begin{cases} \sup_{p \geq 0} \lambda \bar{F}(p)(p + \hat{J}_\beta^*(x-1) - \hat{J}_\beta^*(x)), & \text{if } x > 0. \\ 0, & \text{otherwise.} \end{cases}$$

Farias and Van Roy [2010] show how to compute $\hat{\pi}_\beta^*$ efficiently, and further show that $\hat{\pi}_\beta^*$ is non-decreasing, taking $\hat{\pi}_\beta^*(0) = \infty$. Now define the $\mathcal{G}_t$ random variable Δ_t according to $\Delta_t = \inf\{t - \tau : \tau < t, dN_\tau = 1\}$, and consider the class of dynamic pricing policies, parameterized by $\beta, \bar{\theta}$, given by

$$\hat{\pi}_{(\beta, \bar{\theta}), t} = \hat{\pi}_\beta^*(X_{t-}) - \bar{\theta}\Delta_t.$$

The following is an easy consequence of the fact that $\hat{\pi}_\beta^*$ is non-decreasing:

Lemma 3. *$\hat{\pi}_{(\beta, \bar{\theta}), t}$ is a feasible robust pricing policy for any $\beta > 0$ and $\bar{\theta} \leq \underline{\theta}$.*

The proof of Theorem 1 actually yields the following stronger statement:

Theorem 2. *Let $\beta_0 = 1/1.42T$ and $\bar{\theta} = 0$. Then,*

$$J_{\hat{\pi}_{(\beta_0, 0)}, \tau^{\hat{\pi}_{(\beta_0, 0)}}}(x_0, T) \geq 0.29 J^*(x_0, T).$$

where $\tau_\phi^{\hat{\pi}_{(\beta_0, 0)}} = t_\phi$ is a myopic stopping rule.

4. Analysis: An Optimal Dynamic Mechanism Upper Bound

Towards establishing Theorem 1, we find it useful to compute an upper bound on $J^*(x_0, T)$, the revenue under the optimal dynamic mechanism, in terms of the revenue under an optimal (discounted, infinite horizon) dynamic pricing policy when customers are myopic. To this end, we first prove an intuitive upper bound on $J^*(x_0, T)$ that connects this quantity to a static problem. Specifically, let us denote by ϕ_n, the customer with the nth largest valuation, v^n from among all customers arriving within the sales horizon, T. Let $\hat{x} = x_0 \wedge \max\{n : v^n \geq v^*\}$. We then show:

$$J^*(x_0, T) \leq \mathsf{E}\left[\sum_{n:n \leq \hat{x}} \left(v^n - \frac{\bar{F}(v^n)}{f(v^n)}\right)\right].$$

This upper bound enjoys a simple interpretation: specifically, it is the expected revenue under an optimal (static) auction for x_0 units of an item, where the expectation is over the number of participants in the auction. This result requires we consider a relaxation of our dynamic mechanism design problem where customers can only distort valuation (as opposed to type), and produce a further relaxation employing a suitable envelope theorem. Having proved this result, we will be able to connect this upper bound to a standard (discounted, infinite horizon) dynamic pricing problem.

4.1. A Relaxed Problem

Let us denote by $\phi_{v'}$ the report of customer ϕ when he distorts his valuation to v'. In particular:

$$\phi_{v'} \triangleq (t_\phi, v', \theta_\phi, \alpha_\phi)$$

and consider the following weakened incentive compatibility constraint:

$$\mathsf{E}_{-\phi}\left[U(\phi, y_\phi)\right] \geq \mathsf{E}_{-\phi}\left[U(\phi, y_{\phi_{v'}})\right], \ \forall \ \phi, v' \tag{IC'}$$

(IC') is a relaxation of (IC) since we only allow for distortions of valuation. In what follows, we will frequently drop the $-\phi$ subscript on the expectation where it is clear from context. We now derive an upper bound on the expected price paid by customer ϕ for any feasible mechanism that satisfies (IR) and (IC'):

Lemma 4. *If* (IC') *and* (IR) *hold, then for any* ϕ,

$$\mathsf{E}\left[p_\phi\right] \leq v_\phi \mathsf{E}\left[a_\phi e^{-\alpha_\phi\left(\tau_\phi - t_\phi\right)}\right] - \int_{v'=0}^{v_\phi} \mathsf{E}\left[a_{\phi_{v'}} e^{-\alpha_\phi\left(\tau_{\phi_{v'}} - t_\phi\right)}\right] dv'. \tag{5}$$

Proof. Denote by $u(\phi, y)$ the derivative of U with respect to v, treating y as a constant. We have:

$$\begin{aligned}
\mathsf{E}\left[U(\phi, y_\phi)\right] &= \mathsf{E}[U(\phi_v, y_{\phi_v})] \\
&= \mathsf{E}\left[\int_{v'=0}^{v_\phi} u\left(\phi_{v'}, y_{\phi_{v'}}\right) dv' + U(\phi_0, y_{\phi_0})\right] \\
&\geq \mathsf{E}\left[\int_{v'=0}^{v_\phi} u\left(\phi_{v'}, y_{\phi_{v'}}\right) dv'\right] \\
&= \mathsf{E}\left[\int_{v'=0}^{v_\phi} a_{\phi_{v'}} e^{-\alpha_\phi\left(\tau_{\phi_{v'}} - t_\phi\right)} dv'\right] \\
&= \int_{v'=0}^{v_\phi} \mathsf{E}\left[a_{\phi_{v'}} e^{-\alpha_\phi\left(\tau_{\phi_{v'}} - t_\phi\right)}\right] dv'.
\end{aligned}$$

10

where the first equality is (IC'), the second equality follows from the envelope theorem (Theorem 2) of Milgrom and Segal [2002], the first inequality is due to (IR), and the final equality is via Fubini's theorem. Further, note that:

$$\mathsf{E}[U(\phi, y_\phi)] = \mathsf{E}\left[v_\phi a_\phi e^{-\alpha_\phi(\tau_\phi - t_\phi)} - p_\phi - \theta(\tau_\phi - t_\phi)\right]$$
$$\leq \mathsf{E}\left[v_\phi a_\phi e^{-\alpha_\phi(\tau_\phi - t_\phi)} - p_\phi\right]$$
,

so that with the prior inequality, we have:

$$\mathsf{E}\left[p_\phi\right] \leq v_\phi \mathsf{E}\left[a_\phi e^{-\alpha_\phi(\tau_\phi - t_\phi)}\right] - \int_{v'=0}^{v_\phi} \mathsf{E}\left[a_{\phi_{v'}} e^{-\alpha_\phi(\tau_{\phi_{v'}} - t_\phi)}\right]$$

which is the result. ∎

Now, since Lemma 4 is implied by (IC) and (IR) (noting that (IC') is implied by (IC)), we have the the following optimization problem (whose optimal value we denote by $\bar{J}^*(x_0, T)$) is a relaxation of the optimization problem for $J^*(x_0, T)$:

$$\max_{y^T \in \mathcal{Y}} \quad \mathsf{E}\left[\Pi\left(y^T\right)\right]$$
$$\text{subject to} \quad \mathsf{E}[p_\phi] \leq v_\phi \mathsf{E}\left[a_\phi e^{-\alpha_\phi(\tau_\phi - t_\phi)}\right] - \int_{v'=0}^{v_\phi} \mathsf{E}\left[a_{\phi_{v'}} e^{-\alpha_\phi(\tau_{\phi_{v'}} - t_\phi)}\right], \, \forall \, \phi \tag{6}$$

4.2. The Relaxation And An Upper Bound

We now analyze the relaxed problem and show that $\bar{J}^*(x_0, T)$, the optimal value of the relaxed problem (6) satisfies:

$$\bar{J}^*(x_0, T) \leq \max_{y^T \in \mathcal{Y}} \mathsf{E}\left[\sum_{\phi \in h^T} \left(v - \frac{\bar{F}(v)}{f(v)}\right) a_\phi\right].$$

Lemma 5.

$$J^*(x_0, T) \leq \bar{J}^*(x_0, T) \leq \max_{y^T \in \mathcal{Y}} \mathsf{E}\left[\sum_{\phi \in h^T} \left(v - \frac{\bar{F}(v)}{f(v)}\right) a_\phi\right].$$

Proof. The first inequality is evident since the optimization problem for $\bar{J}^*(x_0, T)$ is a relaxation of that for $J^*(x_0, T)$. Now, observe that the constraint defining $\mathsf{E}[p_\phi]$ must be tight at an optimal solution, so that

$$\bar{J}^*(x_0, T) = \max_{y^T \in \mathcal{Y}} \mathsf{E}\left[\sum_{\phi \in h^T} v_\phi \mathsf{E}_{-\phi}\left[a_\phi e^{-\alpha_\phi(\tau_\phi - t_\phi)}\right] - \int_{v'=0}^{v_\phi} \mathsf{E}_{-\phi}\left[a_{\phi_{v'}} e^{-\alpha_\phi(\tau_{\phi_{v'}} - t_\phi)}\right] dv'\right]$$

where the notation $\mathsf{E}_{-\phi}$ makes explicit that an expectation is over $-\phi$. Now, denote by $W(\phi)$, the following quantity, marginalized over v_ϕ:

$$W(\phi) = \int_{v_\phi=0}^{\infty} \left(v_\phi \mathsf{E}_{-\phi}\left[a_\phi e^{-\alpha_\phi(\tau_\phi - t_\phi)}\right] - \int_{v'=0}^{v_\phi} \mathsf{E}_{-\phi}\left[a_{\phi_{v'}} e^{-\alpha_\phi(\tau_{\phi_{v'}} - t_\phi)}\right] dv'\right) f(v_\phi) dv_\phi,$$

11

so that

$$\bar{J}^*(x_0, T) = \max_{y^T \in \mathcal{Y}} \mathsf{E}\left[\sum_{\phi \in h^T} W(\phi)\right]. \tag{7}$$

Now, applying the 'standard trick' of interchanging integrals for the second term in the integrand in $W(\phi)$, we have:

$$\int_{v_\phi=0}^\infty \int_{v'=0}^{v_\phi} \mathsf{E}_{-\phi}\left[a_{\phi_{v'}} e^{-\alpha_\phi\left(\tau_{\phi_{v'}} - t_\phi\right)}\right] f(v_\phi) dv' dv_\phi$$

$$= \int_{v'=0}^\infty \mathsf{E}_{-\phi}\left[a_{\phi_{v'}} e^{-\alpha_\phi\left(\tau_{\phi_{v'}} - t_\phi\right)}\right] \int_{v_\phi=v'}^\infty f(v_\phi) dv_\phi dv'$$

$$= \int_{v'=0}^\infty \mathsf{E}_{-\phi}\left[a_{\phi_{v'}} e^{-\alpha_\phi\left(\tau_{\phi_{v'}} - t_\phi\right)}\right] \bar{F}(v') dv'$$

so that,

$$W(\phi) = \int_{v_\phi=0}^\infty \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)}\right) \mathsf{E}_{-\phi}\left[a_\phi e^{-\alpha_\phi\left(\tau_\phi - t_\phi\right)}\right] f(v_\phi) dv_\phi.$$

Substituting $W(\phi)$ in (7) with this identity, and applying Fubini's theorem, we have:

$$\bar{J}^*(x_0, T) = \max_{y^T \in \mathcal{Y}} \mathsf{E}\left[\sum_{\phi \in h^T} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)}\right) a_\phi e^{-\alpha_\phi\left(\tau_\phi - t_\phi\right)}\right]$$

$$\leq \max_{y^T \in \mathcal{Y}} \mathsf{E}\left[\sum_{\phi \in h^T} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)}\right) a_\phi\right]$$

which is the result. ∎

4.3. The Discounted Infinite Horizon Problem As An Upper Bound

As our final step for this section, we use the result of Lemma 5, in connection with a result established by Gallien [2006] to relate our dynamic mechanism design benchmark to a simple dynamic pricing problem. Recall that $\hat{J}^*_\beta(x_0)$ denote the optimal value of the discounted, infinite horizon dynamic pricing problem, with myopic customers and discount rate $\beta > 0$, i.e.

$$\hat{J}^*_\beta(x_0) = \max_{\pi \in \hat{\Pi}} \mathsf{E}\left[\int_0^{\hat{\tau}} e^{-\beta t} \pi_t dN_t\right].$$

where $\hat{\Pi}$ is the set of left continuous pricing policies, adapted to $\mathcal{G}_t$, satisfying $\pi_t = \infty$ if $X_{t-} = 0$.

Lemma 6.

$$\hat{J}^*_\beta(x_0) = \max_{y^\infty \in \mathcal{Y}} \mathsf{E}\left[\sum_{\phi \in h^\infty} e^{-\beta t_\phi} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)}\right) a_\phi\right].$$

Proof. Observe that if, on a given sample path, under the optimal policy we accept ϕ, thereby earning $v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)}$, then we would have accepted all $\phi' = (t_\phi, v_{\phi'}, \theta_\phi, \alpha_\phi)$ such that $v_{\phi'} \geq v_\phi$, since such an acceptance would earn

$$v_{\phi'} - \frac{\bar{F}(v_{\phi'})}{f(v_{\phi'})} \geq v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)}$$

12

since $v - \frac{\bar{F}(v)}{f(v)}$ was assumed to be non-decreasing. Consequently, by the optimality of stationary policies, the optimal policy takes the following form:

$$\pi(\phi, X_{t_\phi-}) = \begin{cases} 1, & \text{if } v_\phi \geq \tilde{\pi}(X_{t_\phi-}). \\ 0, & \text{otherwise.} \end{cases}$$

with $\tilde{\pi}(0) \triangleq \infty$. Call the family of all such functions $\tilde{\Pi}$. Consequently, we have

$$\max_{y^\infty \in \mathcal{Y}} \ \mathsf{E}\left[\sum_{\phi \in h^\infty} e^{-\beta t_\phi}\left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)}\right)a_\phi\right] = \max_{\tilde{\pi} \in \tilde{\Pi}} \ \mathsf{E}\left[\sum_{\phi \in h^\infty} e^{-\beta t_\phi}\mathsf{E}\left[\left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)}\right)\mathbb{I}\left[v_\phi \geq \tilde{\pi}(X_{t_\phi-})\right]\right]\right]$$

$$= \max_{\tilde{\pi} \in \tilde{\Pi}} \ \mathsf{E}\left[\sum_{\phi \in h^\infty} e^{-\beta t_\phi}\bar{F}(\tilde{\pi}(X_{t_\phi-}))\tilde{\pi}(X_{t_\phi-})\right]$$

$$= \hat{J}_\beta^*(x_0)$$

where the second inequality used the fact that

$$\int_{v=p}^\infty \left(vf(v) - \bar{F}(v)\right)dv = p\bar{F}(p).$$

This completes the proof. ∎

Combining, Lemmas 6 and 5 yield the final result for this section, an upper bound on $J^*(x_0, T)$ in terms of the optimal value of a (discounted, infinite horizon) dynamic pricing problem with myopic customers:

Lemma 7. *For any $\beta > 0$, we have:*

$$J^*(x_0, T) \leq e^{\beta T}\hat{J}_\beta^*(x_0).$$

Proof. We have:

$$J^*(x_0, T) \leq \max_{y^T \in \mathcal{Y}} \ \mathsf{E}\left[\sum_{\phi \in h^T}\left(v - \frac{\bar{F}(v)}{f(v)}\right)a_\phi\right]$$

$$= \max_{y^\infty \in \mathcal{Y}} \ \mathsf{E}\left[\sum_{\phi \in h^T}\left(v - \frac{\bar{F}(v)}{f(v)}\right)a_\phi\right]$$

$$= e^{\beta T}\max_{y^\infty \in \mathcal{Y}} \ \mathsf{E}\left[\sum_{\phi \in h^T}e^{-\beta T}\left(v - \frac{\bar{F}(v)}{f(v)}\right)a_\phi\right]$$

$$\leq e^{\beta T}\max_{y^\infty \in \mathcal{Y}} \ \mathsf{E}\left[\sum_{\phi \in h^T}e^{-\beta t_\phi}\left(v - \frac{\bar{F}(v)}{f(v)}\right)a_\phi\right]$$

$$\leq e^{\beta T}\max_{y^\infty \in \mathcal{Y}} \ \mathsf{E}\left[\sum_{\phi \in h^\infty}e^{-\beta t_\phi}\left(v - \frac{\bar{F}(v)}{f(v)}\right)a_\phi\right]$$

$$= e^{\beta T}\hat{J}_\beta^*(x_0)$$

where the first inequality is Lemma 5, the second inequality follows since $T \geq t_\phi$ for all $\phi \in h^T$. The final equality is Lemma 6. ∎

13

5. A Robust Dynamic Pricing Lower Bound and The Approximation Guarantee

Our analysis in this section will complete the proof of Theorem 1 using the upper bound on the optimal dynamic mechanism, $J^*(x_0, T)$ established in the preceding section. We will accomplish this in the following steps:

1. First, we construct a feasible robust dynamic pricing policy that is, in effect, the optimal policy for the discounted, infinite horizon problem, applied over a finite horizon.

2. We then prove that this policy accrues expected revenues that are within a constant factor of the revenue optimal infinite horizon revenue.

3. Using this result along with the upper bound on the optimal dynamic mechanism proved in Lemma 7 will yield Theorem 1.

5.1. Infinite Horizon Dynamic Pricing

Consider the infinite horizon dynamic pricing problem introduced in previous sections. Specifically, recall that we defined

$$\hat{J}_\beta^*(x_0) = \max_{\pi \in \hat{\Pi}} \mathsf{E}\left[\int_0^{\hat{\tau}} e^{-\beta t} \pi_t dN_t \right].$$

where $\hat{\tau} = \inf\{t : N_t = x_0\}$, and where $\hat{\Pi}$ is the set of left continuous pricing policies, adapted to $\mathcal{G}_t$, satisfying $\pi_t = \infty$ if $X_{t-} = 0$. We denote by $\{\hat{\pi}_{\beta,t}^*\}$ an optimal policy. From Farias and Van Roy [2010], we have that $\hat{\pi}_{\beta,t}^* \triangleq \hat{\pi}_\beta^*(X_{t-})$, where for all $x > 0$, $\hat{\pi}_\beta^*(x)$ is the root of the equation

$$p - \frac{\bar{F}(p)}{f(p)} = \hat{J}_\beta^*(x) - \hat{J}_\beta^*(x - 1). \tag{8}$$

The optimal price process enjoys the following properties:

Lemma 8. *On every sample path, $\hat{\pi}_{\beta,t}^*$ is non-decreasing in t while $\hat{\pi}_{\beta,t}^* \bar{F}(\hat{\pi}_{\beta,t}^*)$ is non-increasing in t.*

Proof. The first claim is Lemma 1 of Farias and Van Roy [2010]. For the second claim, we observe that since $\hat{J}_\beta^*(x) \geq \hat{J}_\beta^*(x-1)$, it follows from Assumption 1 that $\hat{\pi}_{\beta,t}^* \geq v^*$ for all t. Now, since $p\bar{F}(p)$ is non-increasing in p on $p \geq v^*$ by Assumption 1, it follows that $\hat{\pi}_{\beta,t}^* \bar{F}(\hat{\pi}_{\beta,t}^*)$ is also non-increasing in t. ∎

These properties of the price process yield the following simple result which will be crucial for our lower bound.

Lemma 9. *Let $T, T' > 0$, with $T > T'$. We have:*

$$\mathsf{E}\left[\int_0^{\hat{\tau} \wedge T} \hat{\pi}_{\beta,t}^* dN_t \right] \leq \frac{T}{T'} \mathsf{E}\left[\int_0^{\hat{\tau} \wedge T'} \hat{\pi}_{\beta,t}^* dN_t \right]$$

14

Proof. Since $\hat{\pi}^*_{\beta,t}\bar{F}(\hat{\pi}^*_{\beta,t})$ is non-increasing in t (as established in the preceding Lemma), we have immediately that

$$\int_0^{\hat{\tau}\wedge T} \hat{\pi}^*_{\beta,t}\bar{F}(\hat{\pi}^*_{\beta,t})dt \leq \frac{T}{T'}\int_0^{\hat{\tau}\wedge T'} \hat{\pi}^*_{\beta,t}\bar{F}(\hat{\pi}^*_{\beta,t})dt.$$

The above inequality must also therefore hold in expectation. Now $N_t - \int_0^t \hat{\pi}^*_{\beta,t'}\bar{F}(\hat{\pi}^*_{\beta,t'})dt$ is a $\mathcal{G}_t$ martingale by construction Brémaud [1981], so that

$$\mathsf{E}\left[\int_0^{\hat{\tau}\wedge T} \hat{\pi}^*_{\beta,t}dN_t\right] = \mathsf{E}\left[\int_0^{\hat{\tau}\wedge T} \hat{\pi}^*_{\beta,t}\bar{F}(\hat{\pi}^*_{\beta,t})dt\right]$$

for all $T \geq 0$, which completes the proof. ∎

5.2. A Robust Dynamic Pricing Policy And Proof Of Theorem 1

Now, consider the robust dynamic pricing policy $\{\hat{\pi}_t\}$ defined according to

$$\hat{\pi}_t = \hat{\pi}^*_\beta(X_{t-})$$

for some $\beta > 0$. We observe that this is a robust dynamic pricing policy since it is evidently left continuous and adapted to $\mathcal{G}_t$, and further, trivially satisfies the restricted sub-martingale constraint since $\hat{\pi}^*_\beta(x)$ is non-increasing in x. We show that the revenue obtained under this policy (over the finite horizon T), is lower bounded by a function of the optimal discounted infinite horizon revenue (when the discount rate is β):

Lemma 10.

$$\hat{J}^*_\beta(x_0) \leq \left(1 + \frac{e^{-\beta T}}{\beta T}\right)\hat{J}^*(x_0, T).$$

Proof. Denote by X an exponential random variable with rate β that is independent of the arrival process and customer types. Then,

$$\hat{J}^*_\beta(x_0) = \mathsf{E}\left[\int_0^{\hat{\tau}} e^{-\beta t}\hat{\pi}^*_{\beta,t}dN_t\right] = \mathsf{E}\left[\int_0^{\hat{\tau}\wedge X} \hat{\pi}^*_{\beta,t}dN_t\right]$$

where the equality follows from Fubini's theorem. Moreover, since $\hat{\pi}_t$ as defined prior to the statement of the Lemma is a feasible robust dynamic pricing policy,

$$\mathsf{E}\left[\int_0^{\hat{\tau}\wedge T} \hat{\pi}^*_{\beta,t}dN_t\right] \leq \hat{J}^*(x_0, T).$$

But applying Lemma 9 to every realization of X and taking expectations yields

$$\mathsf{E}\left[\int_0^{\hat{\tau}\wedge X} \hat{\pi}^*_{\beta,t}dN_t\right] \leq \mathsf{E}\left[\max\left\{1, \frac{X}{T}\right\}\right]\mathsf{E}\left[\int_0^{\hat{\tau}\wedge T} \hat{\pi}^*_{\beta,t}dN_t\right].$$

Since

$$\mathsf{E}\left[\max\left\{1, \frac{X}{T}\right\}\right] = 1 + \frac{e^{-\beta T}}{\beta T},$$

the result follows. ∎

15

We can now complete our proof of Theorem 1. Two inequalities established in Lemma 7 and Lemma 10 yield

$$\hat{J}^*(x_0, T) \geq \frac{1}{e^{\beta T} + 1/\beta T} J^*(x_0, T)$$

for any $\beta > 0$. Noting that $J_{\hat{\pi}^*, \tau^{\hat{\pi}^*}}(x_0, T) = \hat{J}^*(x_0, T)$, and taking $\beta = 1/1.42T$ in the preceding inequality yields Theorem 1.

6. Numerical Experiments

In the previous section we established a uniform performance lower bound on the robust pricing policy $\hat{\pi}^*$. In fact, we established this bound for the policy $\hat{\pi}^*_\beta$ taking $\beta = 1/1.42T$. This section will numerically investigate the performance of $\hat{\pi}^*_\beta$, where we will be allowed to tune β. Our objective is two fold. First, we aim to understand numerically how well $\hat{\pi}^*_\beta$ performs. Second, we would like to understand the risk of mis-specification in an optimal dynamic mechanism, and to that end will explore explore the performance loss incurred if the seller misestimates the distribution of customers' time discount factor α and monitoring cost θ, and implements the optimal dynamic mechanism under those mis-specified parameters. Throughout this section, we assume that customers' valuations are exponentially distributed with parameter 1; i.e., $F(v) = 1 - e^{-v}$ for all $v \in \mathbb{R}_+$.

First, we investigate the performance of $\hat{\pi}^*_\beta$. Table 1 reports a lower bound on relative performance. Specifically, the lower bound reported is:

$$\text{LB}(x_0, T) \triangleq \frac{\max_{\beta \in B} J_{\hat{\pi}^*_\beta}(x_0, T)}{J^{\text{UB}}(x_0, T)}.$$

where

$$J^{\text{UB}}(x_0, T) = \max_{y^T \in \mathcal{Y}} \mathsf{E}\left[\sum_{\phi \in h^T} \left(v - \frac{\bar{F}(v)}{f(v)} \right) a_\phi \right]$$

is an upper bound on $J^*(x_0, T)$ by Lemma 5. We selected the optimal β from among a set of discount factors between 0.01 and 100, examined in increments of 0.01.

We make the following two key observations from Table 1:

1. Relative Performance: For a wide range of inventory relative scarcity levels ($x_0/\lambda T$ varies from 0.1 to 1), $\hat{\pi}^*_{\beta*}$ yields revenues which are at least 79% of the optimal revenue, and in most cases, more than 90%.

2. Recall that under $\hat{\pi}^*_\beta$, customers' time discount factors and monitoring costs impact neither the policy nor customers' behavior: every customer behaves myopically, and the seller's revenue is the same as her revenue yielded in the setting in which all customers are myopic. Therefore, the results in Table 1 are *robust*, in that they hold under any type distribution of customers' time discount factors and monitoring costs.

Next, we investigate the robustness (or lack thereof) of the optimal mechanism to mis-specification of discount factor and monitoring cost. We analyze the setting in which all customers have the same time discount factor α and monitoring cost θ, but the seller *incorrectly* believes that all customers are effectively infinitely patent and do not incur such costs ($\alpha = 0$ and $\theta = 0$). The optimal mechanism under the seller's belief is simply to conduct a static revenue maximizing auction at the

Table 1: A lower bound on relative optimality (i.e., $\text{LB}(x_0, T)$).

x_0	β^*	$J_{\hat{\pi}^*_{\beta^*}}(x_0, T)$	$\text{LB}(x_0, T)$
1	0.12	1.49	0.79
2	0.15	2.36	0.84
3	0.18	2.91	0.89
4	0.22	3.26	0.93
5	0.28	3.47	0.96
6	0.34	3.58	0.98
7	0.40	3.64	0.99
8	0.59	3.66	1.00
9	0.77	3.68	1.00
10	0.77	3.68	1.00

Note. The parameters are $\lambda = 1$, $T = 10$.

end of the horizon, whereas the best response from buyers is simply to report their *appropriately discounted value* at the end of the horizon (or not participate and leave the system at t_ϕ, if this quantity turns out to be negative). Denote the revenue yielded under the misspecified 'optimal' mechanism by $J^*_{\alpha,\theta}(x_0, T)$. In this experiment, we allow α and θ to vary in a wide range (from 0.01 to 1). To understand the cost of the mis-specification of α and θ, we compare the revenue under this mis-specified policy against that under the robust dynamic pricing policy studied in the last experiment, and report the quantity:

$$\text{UB}_{\alpha,\theta} \triangleq \frac{J^*_{\alpha,\theta}(x_0, T)}{\max_{\beta \in B} J_{\hat{\pi}^*_\beta}(x_0, T)}.$$

Table 2 reports the results which are quite stark: In essentially all cases, the mis-specified 'optimal' mechanism performed worse than the robust dynamic pricing policy – in many cases *substantially* worse.

Table 2: Performance loss for an optimal but mis-specified mechanism (i.e., $\text{UB}_{\alpha,\theta}$).

$x_0 \backslash (\alpha, \theta)$	$\text{UB}_{\alpha,\theta}$								
	(.01,.01)	(.01,.1)	(.01,1)	(.1,.01)	(.1,.1)	(.1,1)	(1,.01)	(1,.1)	(1,1)
1	1.15	0.90	0.22	0.63	0.49	0.17	0.09	0.09	0.06
2	1.05	0.77	0.15	0.52	0.38	0.11	0.06	0.06	0.04
3	0.99	0.69	0.12	0.45	0.32	0.09	0.05	0.05	0.03
4	0.94	0.64	0.11	0.42	0.29	0.08	0.05	0.04	0.03
5	0.90	0.60	0.10	0.39	0.28	0.08	0.04	0.04	0.03
6	0.88	0.58	0.10	0.38	0.26	0.07	0.04	0.04	0.03
7	0.87	0.57	0.10	0.37	0.26	0.07	0.04	0.04	0.03
8	0.86	0.57	0.10	0.37	0.26	0.07	0.04	0.04	0.03
9	0.86	0.57	0.10	0.37	0.26	0.07	0.04	0.04	0.03
10	0.86	0.57	0.10	0.37	0.26	0.07	0.04	0.04	0.03

Note. The parameters are $\lambda = 1$, $T = 10$.

In conclusion, our numerical experiments suggest the following conclusions:

1. The robust dynamic pricing policy offers excellent performance relative to the optimal dynamic mechanism. This relative performance appears to far exceed the quality suggested by our uniform lower bound.

2. The performance loss incurred due to mis-specification of an optimal mechanism might easily exceed that incurred due to the use of a sub-optimal (but robust) mechanism such as our robust dynamic pricing policy.

7. Concluding Remarks

We have focused on a rich class of revenue management models. The class of models is rich in that we allowed for heterogeneity in customer discount factors and monitoring costs; these were private information in contrast to problems for which the optimal mechanism is known. We proposed a class of pricing mechanisms for this set of models inspired by two very practical requirements:

1. Pricing mechanisms are the mechanism of choice in RM – departures from such mechanism in mainstream applications are few and far-between.

2. It is unclear that calibrating a rich utility model for customers – describing how they discount or their monitoring costs – is possible given the naturally censored nature of the data available for such a task.

In the face of these requirements we have demonstrated a policy that is easy to compute and satisfies a constant factor guarantee with respect to the optimal mechanism. Computational experiments suggest that this policy is, for all intents, near optimal.

In addition to the positive results described above, we have proposed a class of dynamic optimization problems for which we *do not* find the optimal policy. Finding such an optimal policy would seem like an interesting task for future work – as opposed to solving a simple dynamic optimization problem, one must solve a problem with constraints on the trajectory of prices. A comprehensive dynamic programming characterization for this task does not appear obvious; nonetheless, there is hope to solve this problem since we conjecture that the constraints on the trajectory of prices can be enforced by 'local' constraints.

References

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

Y. Aviv and G. Vulcano. Dynamic list pricing. In Ö. Özer and R. Phillips, editors, *The Oxford Handbook of Pricing Management, USA*, chapter 23. Oxford University Press, 2012.

S. Board and A. Skrzypacz. Revenue management with forward-looking buyers. *Unpublished manuscript, Stanford University. Conditionally accepted in Journal of Political Economy*, 2010.

P. Brémaud. *Point processes and queues: Martingale dynamics*. Springer, 1981.

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

V. Farias and B. Van Roy. Dynamic pricing with a prior on market response. *Oper. Res.*, 58(1): 16–29, 2010.

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

J. Gallien. Dynamic mechanism design for online commerce. *Oper. Res.*, 54(2):291–310, 2006.

P. Milgrom and I. Segal. Envelop theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.

M. M. Pai and R. Vohra. Optimal dynamic auctions and simple index rules. *Math. Oper. Res.*, 38 (4):682–697, 2013.

X. Su and F. Zhang. Strategic customer behavior, commitment, and supply chain performance. *Management Sci.*, 54(10):1759–1773, 2008.

K. Talluri and G. van Ryzin. *The Theory and Practice of Revenue Management.* Springer Science+Business Media, 2004.

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