

Robust Dynamic Pricing With Strategic Customers

Yiwei Chen *

Vivek F. Farias †

September 29, 2015

Abstract

We consider the canonical revenue management (RM) problem 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 resisted the design of optimal dynamic mechanisms heretofore. Optimal *pricing* schemes – an almost necessary mechanism format for practical RM considerations – have been similarly elusive.

The present paper proposes a mechanism we dub robust pricing. Robust pricing is guaranteed to achieve expected revenues that are within at least 29% of those under an optimal (not necessarily posted price) dynamic mechanism. We thus provide the first approximation algorithm for this problem. The robust pricing mechanism is practical, since it is an anonymous posted price mechanism and since the seller can compute the robust pricing policy for a problem without any knowledge of the distribution of customer discount factors and monitoring costs. The robust pricing mechanism also enjoys the simple interpretation of solving a dynamic pricing problem for myopic customers with the additional requirement of a novel ‘restricted sub-martingale constraint’ on prices that discourages rapid discounting. We believe this interpretation is attractive to practitioners. Finally, numerical experiments suggest that the robust pricing mechanism is, for all intents, near optimal.

1. Introduction

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 canonical 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

*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}

the selling horizon with the intent of purchasing a single unit of the product. Should the posted price upon a customer's 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*.

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 place strong restrictions on information asymmetry. Specifically, *no* mechanisms are available for the (typical) setting where customer discount factors or monitoring costs are private information.
3. These mechanisms frequently impute sophisticated purchase timing decisions in equilibrium that are arguably as untenable as the non-strategic 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. 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’. This restricted sub-martingale constraint effectively places an intuitive restriction on the *rate* at which the seller can discount. We dub such policies ‘robust dynamic pricing’ policies. We show how to compute optimal robust pricing policies and demonstrate that such policies admit attractive properties:

1. Computing a robust policy requires minimal data on customers beyond what is already required by the standard dynamic pricing problem assuming myopic customers. Customer discount factors and monitoring costs are private information.
2. Robust pricing policies induce customers to behave myopically under mild assumptions on customer utility.
3. Optimal robust pricing policies are in essence no harder to compute than their non-robust counterparts.

4. We exhibit a robust pricing policy that is guaranteed to garner revenues that are within 29% of those garnered under the optimal dynamic mechanism¹.

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, allowing them to be private information.

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 focuses on computation, showing that an optimal robust pricing policy can be computed by solving a three dimensional dynamic program; this section is of independent interest in that the correctness of the three dimensional DP formulation follows from a proof of a state space collapse. Section 7 complements our analysis with a brief numerical study. Section 8 concludes with thoughts on future directions.

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 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. Recent empirical work, most notably Moon et al. [2015] and Li et al. [2014] has established that this forward looking behavior is highly prevalent. Interestingly, the paper by Moon et al. [2015] directly estimates a customer utility model that is a special case of the model studied in this paper.

The antecedent literature most relevant to the present paper is associated with the area of dynamic mechanism design. Before discussing that stream of literature, however, we pause to mention that a *vast* amount of research in the operations management community has been dedicated to establishing qualitative insights on pricing strategies beneficial to a revenue manager faced with forward looking customers. This work is primarily conducted in a stylized, two period setting and

¹Note that mechanism need not be a pricing policy.

seeks to understand issues as far ranging as the ability to commit to a pricing policy to the impact of a ‘mix’ of strategic and myopic customers, to the value of ‘quick response’ policies to name just a few topics of interest. Several surveys of this literature are available, including those by Shen and Su [2007], Ho and Su [2009], Netessine and Tang [2009], and Aviv and Vulcano [2012].

Dynamic Mechanism Design: Closest 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 mechanism 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. *Only when* these deadlines are known to the seller, the authors characterize the optimal mechanism completely and show that it 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. We note that there is a sizable body of literature leading to the papers by Board and Skrzypacz [2010] and Pai and Vohra [2013], and discussed therein. We do not repeat that discussion here.

It is worth contrasting the present paper with the aforementioned mechanism design research:

1. We allow for customers’ discount factors to be 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 in essence the richest model one might consider for this canonical problem.

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 as in Gallien [2006]).
3. We provide a mechanism that enjoys a constant factor approximation guarantee relative to the optimal mechanism for our setting. The optimal dynamic mechanism is unknown (and in light of the Pai and Vohra [2013] paper, likely intractable). We provide what is the first approximation algorithm for this challenging problem.
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.

Outside of the three core points of reference discussed above, a nice variety of algorithmic work, motivated by the presence of consumer strategicity in RM flavored problems, has emerged in recent years. For instance, the paper by Borgs et al. [2014] considers a setting where a firm with time varying capacity sets prices over time to maximize revenues in the face of strategic customers. Inventory cannot be carried over from one epoch to the next. Like Pai and Vohra [2013], customers have arrival times, deadlines and valuations. However, these quantities are assumed known. The focus of the paper is thus on the dynamic optimization problem that arises in this setting and the authors contribute a surprising dynamic programming formulation. It is worth mentioning that Said [2012] considers and solves a mechanism design problem for a setting similar to Borgs et al. [2014] with the exception that customers have discount rates (as opposed to deadlines) that are homogeneous and known, and valuations remain unobserved. Finally, a recent paper by Caldentey et al. [2015] examines a very similar problem but in the absence of priors, optimizing instead a regret objective

Akan et al. [2009] consider a variant of the mechanism design problem where customers lives are the entirety of the selling horizon (i.e. all customers are present at time zero), but they become aware of their valuations at a time known to them. Customers have no preference with respect to the time of allocation and do not incur monitoring costs. The authors are able to construct the optimal mechanism in this setting; of course, the assumption that all customers are available to make a decision at the start of the horizon is restrictive and detracts from RM applications.

Although fairly distinct from the RM problem we focus on, a related interesting stream of literature has emerged focused on problems wherein a seller has *multiple* interactions with a buyer. This setting corresponds naturally to some recent applications of RM techniques, such as the allocation of impressions online to a set of advertisers, and affords the seller the opportunity to learn about customers over time. An early paper in this setting is Bergemann and Välimäki [2010];

that work focuses on efficient mechanisms. More recently, several papers have presented revenue maximizing solutions to variants around this common theme. These include Kakade et al. [2013] and Pavan et al. [2014]. The Kakade et al. [2013] paper provides a particularly elegant description of this interesting class of problems and an equally interesting, Meyersonian solution.

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 process².

Customers arrive over this period according to a Poisson process of rate λ ; an extension to non-homogenous processes is possible. 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

$$\phi \triangleq (t_\phi, v_\phi, \theta_\phi, \alpha_\phi).$$

In the sequel, we will make the dependence of each component on ϕ explicit only when needed. After making a purchase decision, customers exit the system. Assume that such a customer chooses to delay making a purchase decision to time $\tau_\phi \geq t_\phi$, 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 allocate³ 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

$$U(\phi, y_\phi) = a_\phi \left(e^{-\alpha_\phi(\tau_\phi - t_\phi)} v_\phi - p_\phi \right) - \theta_\phi(\tau_\phi - t_\phi).$$

Modeling inter-temporal preferences and monitoring costs as we have above is relatively commonplace in the literature; see for instance Aviv and Pazgal [2008], Cachon and Swinney [2009], Su and Zhang [2009], and Cachon and Feldman [2015]. In fact, Moon et al. [2015] directly estimate a special case of this model in an online RM context.

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

²More formally, we require π_t to be left continuous, and adapted to \mathcal{F}_{t-} where $\mathcal{F}_t = \sigma(\pi^t, X^t)$.

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

$f(\cdot)$. We denote $\bar{F}(\cdot) \triangleq 1 - F(\cdot)$. We assume that a customer's 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 solves⁴ the optimal stopping problem

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

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

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

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 $y^t \triangleq \{y_\phi : \tau_\phi \leq t\}$ 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, $\mathcal{H}_t = \sigma(h^t, y^{t-})$. 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}_{τ_ϕ} -measurable.

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

2. Limited Inventory: The seller cannot allocate more units of product than her initial allocation:

$$\sum_{\phi \in h^T} a_\phi \leq x_0, \text{ a.s.}$$

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

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

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 $h^T \setminus \{\phi\} \cup \{\hat{\phi}\}$.

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

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

Denote by $J^*(x_0, T)$ the optimal value obtained in the problem above. We have the following result⁵ illustrating that this constitutes an interesting benchmark (a proof may be found in the Appendix):

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

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

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); correspondingly this benchmark is no weaker than the optimal mechanisms for those respective cases.

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

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:

$$\mathbb{E} \left[(\pi_t - \pi_{t'})^+ | \mathcal{G}_t \right] \leq \underline{\theta}(t' - t) \quad (2)$$

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:

$$\hat{J}^*(x_0, T) \triangleq \sup_{\{\pi_t\} \in \Pi} \mathbb{E} \left[\int_0^{\hat{\tau} \wedge T} \pi_t dN_t \right] \quad (3)$$

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 management problem is the constraint placed on sample paths of the pricing policy via the restricted sub-martingale constraint (2).

Motivation: In the absence of the restricted sub-martingale constraint (2), the dynamic optimization problem above is identical to what one may consider the canonical RM problem studied in Gallego and van Ryzin [1994]. The second constraint implies

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

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

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

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.

Computation: The problem of computing an optimal robust pricing policy is akin to the problem of computing an optimal dynamic pricing policy, with additional requirement that the policy computed satisfy the restricted sub-martingale constraint. This additional constraint is effectively a constraint on sample paths of the pricing policy and it is ex-ante unclear that the solution of a such a problem is amenable to a traditional dynamic programming. Fortunately, in Section 6, we prove that the restricted sub-martingale constraint is, under mild technical conditions, equivalent to a *local* constraint on sample paths of the pricing policy that permits a dynamic programming solution using a so-called 'post-decision' state. In summary, this allows us to compute an optimal solution to the robust pricing problem via the solution of a three-dimensional HJB equation; the standard (myopic) dynamic pricing problem admits a two-dimensional HJB equation.

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:

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

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

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

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 $t \geq t' \geq t_\phi$:

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

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

Denote by $\hat{\pi}^*$ an optimal solution to (3). 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:

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*

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

This result is remarkable from two perspectives:

1. From a theoretical perspective, as discussed in our review of the literature, optimal mechanism design for the revenue management problem considered here is difficult. Pai and Vohra [2013] make the case that an optimal solution is likely intractable. In light of this, it is remarkable that a simple, easy to interpret anonymous pricing policy enjoys a uniform performance guarantee.
2. From a practical perspective, the mechanism we propose is easily implemented, seeing as it is precisely the same format (anonymous pricing) as the plurality of (myopic) RM policies used today, and enjoys modest information requirements.

The next two sections are dedicated to establishing Theorem 1. 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. 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.

The next two sections are dedicated to a proof of Theorem 1. Broadly, this will proceed in two steps. The first step (Section 4) entails producing an upper bound on the sellers revenue under the optimal dynamic mechanism, $J^*(x_0, T)$. It will be important that we connect this upper bound to the value under a certain dynamic pricing scheme. The second step (Section 5) entails constructing a *feasible* robust dynamic pricing policy and comparing its revenue to $J^*(x_0, T)$ via the upper bound computed in Section 4. This will conclude the proof of our performance guarantee.

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 n th largest valuation, $v_{\phi_n} \triangleq v^n$, 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:

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

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:

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

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

$$\mathbb{E}[p_\phi] \leq v_\phi \mathbb{E} \left[a_\phi e^{-\alpha_\phi(\tau_\phi - t_\phi)} \right] - \int_{v'=0}^{v_\phi} \mathbb{E} \left[a_{\phi_{v'}} e^{-\alpha_\phi(\tau_{\phi_{v'}} - t_\phi)} \right] dv'. \quad (4)$$

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

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

where the first equality is (IC'), the second equality follows from the envelope theorem (specifically, 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:

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

so that with the prior inequality, we have:

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

which is the result. ■

Now, since Lemma 3 is implied by (IC) and (IR) (noting that (IC') is implied by (IC)), we have that the following optimization problem (whose optimal value we denote by $\bar{J}^*(x_0, T)$) is a relaxation of the optimization problem for $J^*(x_0, T)$:

$$\begin{aligned} \max_{y^T \in \mathcal{Y}} \quad & \mathbb{E} \left[\Pi(y^T) \right] \\ \text{subject to} \quad & \mathbb{E}[p_\phi] \leq v_\phi \mathbb{E} \left[a_\phi e^{-\alpha_\phi(\tau_\phi - t_\phi)} \right] - \int_{v'=0}^{v_\phi} \mathbb{E} \left[a_{\phi_{v'}} e^{-\alpha_\phi(\tau_{\phi_{v'}} - t_\phi)} \right], \forall \phi \end{aligned} \quad (5)$$

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 (5) satisfies:

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

Lemma 4.

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

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 $\mathbb{E}[p_\phi]$ must be tight at an optimal solution, so that

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

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

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

so that

$$\bar{J}^*(x_0, T) = \max_{y^T \in \mathcal{Y}} \mathbb{E} \left[\sum_{\phi \in h^T} W(\phi) \right]. \quad (6)$$

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

$$\begin{aligned} & \int_{v_\phi=0}^{\infty} \int_{v'=0}^{v_\phi} \mathbb{E}_{-\phi} \left[a_{\phi_{v'}} e^{-\alpha_\phi(\tau_{\phi_{v'}} - t_\phi)} \right] f(v_\phi) dv' dv_\phi \\ &= \int_{v'=0}^{\infty} \mathbb{E}_{-\phi} \left[a_{\phi_{v'}} e^{-\alpha_\phi(\tau_{\phi_{v'}} - t_\phi)} \right] \int_{v_\phi=v'}^{\infty} f(v_\phi) dv_\phi dv' \\ &= \int_{v'=0}^{\infty} \mathbb{E}_{-\phi} \left[a_{\phi_{v'}} e^{-\alpha_\phi(\tau_{\phi_{v'}} - t_\phi)} \right] \bar{F}(v') dv' \end{aligned}$$

so that,

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

Substituting $W(\phi)$ in (6) with this identity, and applying Fubini's theorem, we have:

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

which is the result. ■

4.3. The Discounted Infinite Horizon Problem As An Upper Bound

Recall that $\hat{J}_\beta^*(x_0)$ denotes the optimal value of the discounted, infinite horizon dynamic pricing problem, with myopic customers and discount rate $\beta > 0$, i.e.

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

where $\hat{\Pi}$ is the set of left continuous pricing policies, adapted to \mathcal{G}_t , satisfying $\pi_t = \infty$ if $X_{t-} = 0$. As our final step for this section, we use the result of Lemma 4, in connection with an interesting representation of $\hat{J}_\beta^*(\cdot)$ to effectively relate our dynamic mechanism design benchmark to a simple dynamic pricing problem with myopic customers.

Lemma 5.

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

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

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

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

$$\begin{aligned} \max_{y^\infty \in \mathcal{Y}} \mathbb{E} \left[\sum_{\phi \in h^\infty} e^{-\beta t_\phi} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \right) a_\phi \right] &= \max_{\tilde{\pi} \in \tilde{\Pi}} \mathbb{E} \left[\sum_{\phi \in h^\infty} e^{-\beta t_\phi} \mathbb{E} \left[\left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \right) \mathbb{I} [v_\phi \geq \tilde{\pi}(X_{t_\phi-})] \right] \right] \\ &= \max_{\tilde{\pi} \in \tilde{\Pi}} \mathbb{E} \left[\sum_{\phi \in h^\infty} e^{-\beta t_\phi} \bar{F}(\tilde{\pi}(X_{t_\phi-})) \tilde{\pi}(X_{t_\phi-}) \right] \\ &= \hat{J}_\beta^*(x_0) \end{aligned}$$

where the second inequality used the fact that

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

This completes the proof. ■

Combining, Lemmas 5 and 4 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 6. *For any $\beta > 0$, we have:*

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

Proof. We have:

$$\begin{aligned} J^*(x_0, T) &\leq \max_{y^T \in \mathcal{Y}} \mathbb{E} \left[\sum_{\phi \in h^T} \left(v - \frac{\bar{F}(v)}{f(v)} \right) a_\phi \right] \\ &= \max_{y^\infty \in \mathcal{Y}} \mathbb{E} \left[\sum_{\phi \in h^T} \left(v - \frac{\bar{F}(v)}{f(v)} \right) a_\phi \right] \\ &= e^{\beta T} \max_{y^\infty \in \mathcal{Y}} \mathbb{E} \left[\sum_{\phi \in h^T} e^{-\beta T} \left(v - \frac{\bar{F}(v)}{f(v)} \right) a_\phi \right] \\ &\leq e^{\beta T} \max_{y^\infty \in \mathcal{Y}} \mathbb{E} \left[\sum_{\phi \in h^T} e^{-\beta t_\phi} \left(v - \frac{\bar{F}(v)}{f(v)} \right) a_\phi \right] \\ &\leq e^{\beta T} \max_{y^\infty \in \mathcal{Y}} \mathbb{E} \left[\sum_{\phi \in h^\infty} e^{-\beta t_\phi} \left(v - \frac{\bar{F}(v)}{f(v)} \right) a_\phi \right] \\ &= e^{\beta T} \hat{J}_\beta^*(x_0) \end{aligned}$$

where the first inequality is Lemma 4, the second inequality follows since $T \geq t_\phi$ for all $\phi \in h^T$. The final equality is Lemma 5. ■

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 optimal infinite horizon revenue.
3. Using this result along with the upper bound on the optimal dynamic mechanism proved in Lemma 6 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

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

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

$$p - \frac{\bar{F}(p)}{f(p)} = \hat{J}_\beta^*(x) - \hat{J}_\beta^*(x-1). \quad (7)$$

The optimal price process enjoys the following properties:

Lemma 7. *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 8. Let $T, T' > 0$, with $T > T'$. We have:

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

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 (see Brémaud [1981]), so that

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

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 9.

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

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

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

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,

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

But applying Lemma 8 to every realization of X and taking expectations yields

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

Since

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

the result follows. ■

We can now complete our proof of Theorem 1. Two inequalities established in Lemma 6 and Lemma 9 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. Computing Optimal Robust Dynamic Pricing Policies

This section presents an approach to computing an optimal robust dynamic pricing policy. The approach proceeds by characterizing the optimal robust dynamic pricing policy via an appropriate Hamilton-Jacobi-Bellman (HJB) PDE. While such characterizations are typically routine, in our case, the characterization merits careful attention. Specifically:

1. The restricted sub-martingale constraint ostensibly places a constraint on *sample path* properties of the pricing policy, and such constraints are not amenable to dynamic programming approaches in general. In our case, however, we show that this constraint is equivalent to a ‘local’ constraint on optimal prices.
2. The dynamic programming formulation appears to require the use of a so-called ‘post-decision’ state which is atypical for dynamic programming formulations employed in the context of dynamic pricing.

Before proceeding we restrict attention to a specific class of robust dynamic pricing policies that are sufficiently ‘regular’ to admit analysis. We will restrict attention to robust pricing policies within this class.

Assumption 2 (Regularity). *Denote by $\{t_n\}$ the set of times at which a sale occurs. Let π_t^c be a continuous, piecewise differentiable process, adapted to \mathcal{F}_t and potentially non-differentiable only on the set $\{t_n\}$. Moreover, let Δ_t be a bounded left continuous process adapted to \mathcal{F}_t . We require that π_t admit the decomposition:*

$$\pi_t = \pi_t^c + \sum_{t_n < t} \Delta_{t_n}$$

where both $\left| \frac{\partial^-}{\partial t} \pi_t^c \right|$ and $|\Delta_t|$ are bounded above by some constant B^7 .

We refer to the class of robust pricing policies that satisfy Assumption 2 as *regular robust pricing policies*. Observe that the decomposition above is intuitive – specifically, it decomposes the pricing policy into a ‘continuous’ component and a ‘jump’ component. Our focus in this section will be presenting a dynamic programming approach to computing an optimal regular robust pricing policy. Our first step in this task is to present a ‘local’ property on price paths that is equivalent to the restricted sub-martingale constraint for robust pricing policies.

6.1. A Local Constraint On Prices

Recall that the restricted sub-martingale constraint (2) on dynamic pricing policies took the form:

$$\mathbb{E} \left[(\pi_t - \pi_{t'})^+ \mid \mathcal{G}_t \right] \leq \underline{\theta}(t' - t) \text{ a.s.}$$

for all times $t \leq t'$. We show that for regular robust pricing policies this is equivalent to a local condition. The proof is somewhat technical in nature and deferred to the appendix.

Proposition 1. *A regular robust pricing policy π satisfies the restricted sub-martingale constraint (2) if and only if it satisfies*

$$\left(\frac{\partial^-}{\partial s} \pi_s^c \right)^- + \lambda \bar{F}(\pi_s) (\Delta_s)^- \leq \underline{\theta} \text{ a.s.} \quad (8)$$

for all s .

Beyond being computationally useful (as we shall see), the result above is illuminating. Specifically, it makes precise the notion that we do not want to allow prices to fall ‘too quickly’; an idea is at the heart of robust pricing.

6.2. The HJB Equation

Denote by \mathcal{Y} the state space $\mathbb{R} \times \mathbb{R} \times \mathbb{N}$. Denote by $y \triangleq (s, p, x)$ a generic element of \mathcal{Y} . Let $J : \mathcal{Y} \rightarrow \mathbb{R}$, be given, differentiable in its first two arguments and satisfying $J(T, \cdot, \cdot) = 0$, and $J(\cdot, \cdot, 0) = 0$. Denote by \mathcal{J} the set of all such functions. Denote $\Pi^M \triangleq \{\pi^M\}$ where $\pi^M : \mathcal{Y} \rightarrow \mathbb{R} \times \mathbb{R}$ denote a Markov ‘pricing policy’. We require

$$(\pi^M(y)_0)^- + \lambda \bar{F}(p)(\pi^M(y)_1)^- \leq \underline{\theta} \quad \forall y \in \mathcal{Y},$$

and further that $(\pi^M(y)_0)^-$ and $(\pi^M(y)_1)^-$ are both zero if $p = 0$, while $\|\pi^M(y)\|_\infty \leq B \forall y \in \mathcal{Y}$. Let Π denote the set of all such markov policies. Finally, define the operator H_{π^M} acting on \mathcal{J}

⁷ B is allowed to depend on problem data – namely λ, x_0, T and the distribution of reservation prices.

according to:

$$(H_{\pi^M} J)(y) = \frac{\partial}{\partial s} J(s, p, x) + \frac{\partial}{\partial p} J(s, p, x) \pi^M(s)_0 + \left[p + J(s, p + \pi^M(s)_1, x - 1) - J(s, p, x) \right] \lambda \bar{F}(p)$$

for all y such that $s < T$ and $x > 0$. Outside this set, we define $(H_{\pi^M} J)(y)$ to be identically 0. The HJB equation for our problem can then be stated as

$$\sup_{\pi^M \in \Pi^M} H_{\pi^M} J = 0 \tag{9}$$

While we will derive rigorous results shortly, it is worth making a few informal comments here. The HJB equation above is easy to intuit by considering an appropriate discrete time system and taking formal (non-rigorous) limits. There, one sees the role of the various state variables: s and x (time, and inventory) are usual and make an appearance in the HJB equation for the standard dynamic pricing problem. The third state variable is current price (p), and encapsulates all the information needed to guarantee a robust pricing policy. In general, this third state variable would have required us to track the entire history of the price process (leading to an intractable problem), but in our case, thanks to the result of Proposition 1, simply tracking the current price is sufficient yielding an effective state-space collapse.

We next assume that the HJB equation admits a solution:

Assumption 3. *The HJB equation admits a solution $J^* \in \mathcal{J}$.*

Now, since the supremum of a continuous function over a compact set is attained, we know that the supremum in the HJB equation is attained; denote by $\pi^*(y)$ such a point. Moreover, the set of points attaining the supremum varies continuously in s and p . We thus define the regular robust pricing policy π^* according to

$$\pi_t^* = \int_0^t \pi^{*,M}(s)_0 ds + \sum_{t_n < s} \pi^{*,M}(s)_1 \tag{10}$$

We now present the main result for this section, showing that a solution to the HJB equation yields an optimal robust pricing policy.

Theorem 2. *Let π^* be the pricing policy obtained as the solution to the HJB equation, as defined in (10). Let π_s be any regular robust pricing policy. Then,*

$$\mathbb{E} \left[\int_0^T \lambda \pi_s \bar{F}(\pi_s) ds \right] \leq \mathbb{E} \left[\int_0^T \lambda \pi_s^* \bar{F}(\pi_s^*) ds \right].$$

The proof of Theorem 2 can be found in the Appendix. Contrasting the HJB equation here with that one would encounter in a ‘vanilla’ dynamic pricing problem (such as in, say, Gallego and van Ryzin [1994]) shows that the additional computational complexity here is marginal – a

three dimensional problem as opposed to two. In essence, the HJB equation here must control the rate at which the price is ‘discounted’ between successive sales, and in doing so guarantees that the restricted sub-martingale constraint on prices is met. As mentioned previously, since this constraint depends on the current price, it then becomes necessary to track the current price as part of the state as well.

7. Numerical Experiments

The goal of the present section is to complement our theoretical analysis with several numerical experiments. Specifically, we have three goals:

1. Characterizing sub-optimality of the robust pricing mechanism relative to an optimal dynamic pricing mechanism (since our uniform performance guarantee is unlikely to be tight). Here we would like to make the point that one might expect near optimal performance in practice.
2. Characterizing the price of mis-specification: even assuming we could calculate the optimal dynamic mechanism, such a mechanism would require careful calibration of a variety of parameters (for instance, customer discount factors). We would like to make the point that the optimal mechanism is likely to be fragile with respect to mis-specification.
3. Finally, we would like to highlight the performance of a certain *non-optimal* robust pricing policy that satisfies our performance guarantees and is easier to compute than the optimal robust pricing policy.

Recall that we established a uniform performance lower bound on the robust pricing policy $\hat{\pi}_\beta^*$ (defined in (7)). In particular, 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 . In addition to performance loss for the robust mechanism, we would like to understand the risk of mis-specification in an optimal dynamic mechanism. To that end, we will 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

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

is an upper bound on $J^*(x_0, T)$ by Lemma 4. We selected the optimal β from among a set of discount factors between 0.01 and 100, examined in increments of 0.01.

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 $\lambda = 1$, $T = 10$.

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 patient 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 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.

8. 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. Computing the optimal mechanism in our setting has been intractable heretofore, and our proposed mechanism constitutes the first approximation algorithm for the problem that admits uniform guarantees. Computational experiments suggest that this policy is, for all intents, near optimal.

Looking to the future, one interesting direction for future work is data-centric. Specifically, it may be possible to collect from data some statistical information about parameters such as θ and α ; Moon et al. [2015] is a good example of empirical work in this direction. The robust pricing mechanisms we propose are effectively agnostic to much of this information, and it would be interesting to consider the design of robust mechanisms that are able to use the refined distributional information at some expense to robustness.

Acknowledgement

An extended abstract of this work appeared in the EC 2015 conference; we are grateful to the program committee and several conference participants for their excellent feedback. The first author wishes to acknowledge valuable discussions with Dragos F. Ciocan, Zizhuo Wang and Yehua Wei. The second author wishes to acknowledge valuable discussions with Itai Ashlagi, Yash Kanoria, Hamid Nazerzadeh, Nikolaos Trichakis, Ben Van Roy and Gabriel Weintraub. The second author's work was supported in part by NSF CAREER grant number CMMI-1054034.

References

- M. Akan, B. Ata, and J. Dana. Revenue management by sequential screening. *Unpublished manuscript, Carnegie Mellon University*, 2009.
- 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.
- D. Bergemann and J. Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2):771–789, 2010.
- S. Board and A. Skrzypacz. Revenue management with forward-looking buyers. *Unpublished manuscript, Stanford University. Conditionally accepted in Journal of Political Economy*, 2010.

- C. Borgs, O. Candogan, J. Chayes, I. Lobel, and H. Nazerzadeh. Optimal multiperiod pricing with service guarantees. *Management Science*, 60(7):1792–1811, 2014.
- P. Brémaud. *Point processes and queues: Martingale dynamics*. Springer, 1981.
- G. Cachon and P. Feldman. Price commitments with strategic consumers: Why it can be optimal to discount more frequently ... than optimal. *Manufacturing Service Oper. Management*, Articles in Advance:1–12, 2015.
- G. Cachon and R. Swinney. Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Sci.*, 55(3):497–511, 2009.
- R. Caldentey, Y. Liu, and I. Lobel. Intertemporal pricing under minimax regret. 2015.
- 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.
- T-H. Ho and X. Su. Strategic pricing response and optimization in operations management. In V. R. Rao, editor, *Handbook of Pricing Research in Marketing*. Edward Elgar Publishing Company, 2009.
- S. M. Kakade, I. Lobel, and H. Nazerzadeh. Optimal dynamic mechanism design and the virtual-pivot mechanism. *Operations Research*, 61(4):837–854, 2013.
- J. Li, N. Granados, and S. Netessine. Are consumers strategic? structural estimation from the air-travel industry. *Management Science*, 60(9):2114–2137, 2014.
- P. Milgrom and I. Segal. Envelop theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- K. Moon, K. Bimpikis, and H. Mendelson. Randomized markdowns and online monitoring. *Available at SSRN 2589386*, 2015.
- S. Netessine and C. S. Tang. *Consumer-Driven Demand and Operations Management Models*. Springer, Dordrecht, 2009.
- M. M. Pai and R. Vohra. Optimal dynamic auctions and simple index rules. *Math. Oper. Res.*, 38(4):682–697, 2013.
- A. Pavan, I. Segal, and J. Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014.

- M. Said. Auctions with dynamic populations: Efficiency and revenue maximization. *Journal of Economic Theory*, 147(6):2419–2438, 2012.
- Z. J. Shen and X. Su. Customer behavior modeling in revenue management and auctions: A review and new research opportunities. *Production Oper. Management*, 16(6):713–728, 2007.
- X. Su and F. Zhang. On the value of commitment and availability guarantees when selling to strategic consumers. *Management Sci.*, 55(5):713–726, 2009.
- 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.

A. Proofs for Section 2

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, $\mathcal{Y}^p \subset \mathcal{Y}$, 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 than no allocation, and if inventory is available. Now consider the optimization problem:

$$\begin{aligned} \max_{y^T \in \mathcal{Y}^p} \quad & \mathbb{E} \left[\Pi \left(y^T \right) \right] \\ \text{subject to} \quad & \mathbb{E}_{-\phi} \left[U(\phi, y_\phi) | \mathcal{P}_{t_\phi} \right] \geq \mathbb{E}_{-\phi} \left[U(\phi, y_{\hat{\phi}}) | \mathcal{P}_{t_\phi} \right], \text{ a.s., } \forall \phi, \hat{\phi}, \text{ s.t. } t_{\hat{\phi}} \geq t_\phi \quad (\text{IC}) \quad (11) \\ & \mathbb{E}_{-\phi} [U(\phi, y_\phi)] \geq 0, \quad \forall \phi. \quad (\text{IR}) \end{aligned}$$

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 $\tau_\phi = \tau_\phi^\pi$. Since by definition τ_ϕ^π is a best response to itself and the pricing policy π , (IC) is satisfied in (11). But (1) is a relaxation of (11), so $J^{\pi^*}(x_0, T) \leq J^*(x_0, T)$ completing the proof. ■

B. Proofs for Section 6

Proposition 1. A regular robust pricing policy π satisfies the restricted sub-martingale constraint (2) if and only if it satisfies

$$\left(\frac{\partial_-}{\partial s} \pi_s^c \right)^- + \lambda \bar{F}(\pi_s) (\Delta_s)^- \leq \underline{\theta} \text{ a.s.} \quad (12)$$

for all s .

Proof. We first show that (8) implies the restricted sub-martingale constraint (2). Specifically, for all times $t \leq t'$:

$$\begin{aligned}
\mathbb{E} \left[(\pi_t - \pi_{t'})^+ \mid \mathcal{G}_t \right] &= \mathbb{E} \left[\left(- \int_t^{t'} \frac{\partial_-}{\partial s} \pi_s^c ds - \int_t^{t'} \Delta_s dN_s \right)^+ \mid \mathcal{G}_t \right] \\
&\leq \mathbb{E} \left[\int_t^{t'} \left(\frac{\partial_-}{\partial s} \pi_s^c \right)^- ds + \int_t^{t'} (\Delta_s)^- dN_s \mid \mathcal{G}_t \right] \\
&= \mathbb{E} \left[\int_t^{t'} \left(\frac{\partial_-}{\partial s} \pi_s^c \right)^- + \lambda \bar{F}(\pi_s) (\Delta_s)^- ds \mid \mathcal{G}_t \right] \\
&\leq \underline{\theta}(t' - t) \text{ a.s.}
\end{aligned}$$

where the first inequality follows from a repeated application of Jensen's inequality, and the final inequality from assuming (8). We next demonstrate that the restricted sub-martingale constraint (2) implies the local condition (8). First, define the \mathcal{G}_t stopping time τ_δ according to:

$$\tau_{t,\delta} = \inf \left\{ \hat{t} : \int_t^{\hat{t}} \lambda \bar{F}(\pi_s) ds \geq \delta \right\}.$$

Then, we have from Meyer's random time-change theorem (Theorem II.T16 in Brémaud [1981]), that $N_{t,\tau_{t,\delta}} \triangleq \int_t^{\tau_{t,\delta}} dN_s$ is distributed as Poisson random variable with parameter δ . Moreover, we have by construction that

$$\lim_{\delta \rightarrow 0} \frac{\tau_{t,\delta} - t}{\delta} = 1/\lambda \bar{F}(\pi_t) \text{ a.s.}$$

Now on the set where $N_{t,\tau_{t,\delta}} = 0$, we have by Taylor's theorem and the above limit that

$$\left(\pi_t - \pi_{\tau_{t,\delta}} \right)^+ = \left(\frac{\partial_-}{\partial s} \pi_s^c \right)^- \Big|_{s=t} \frac{\delta}{\lambda \bar{F}(\pi_t)} + R_1(\delta)$$

where $|R_1(\delta)| = o(\delta)$. While on the set where $N_{t,\tau_{t,\delta}} = 1$, we have that

$$\left(\pi_t - \pi_{\tau_{t,\delta}} \right)^+ = -(\Delta_t)^- + R_2(\delta),$$

where $|R_2(\delta)| = O(\delta)$. Both remainder terms are uniformly bounded from above since prices are

uniformly bounded from above. It follows that

$$\begin{aligned}
\mathbb{E} \left[\left(\pi_t - \pi_{\tau_{t,\delta}} \right)^+ | \mathcal{G}_t \right] &\geq \mathbb{E} \left[\left(\pi_t - \pi_{\tau_{t,\delta}} \right)^+ \mathbf{1}_{N_{t,\tau_{t,\delta}}=0} | \mathcal{G}_t \right] + \mathbb{E} \left[\left(\pi_t - \pi_{\tau_{t,\delta}} \right)^+ \mathbf{1}_{N_{t,\tau_{t,\delta}}=1} | \mathcal{G}_t \right] \\
&\geq \left(\frac{\partial_-}{\partial s} \pi_s^c \right)^- \bigg|_{s=t} \frac{\delta}{\lambda \bar{F}(\pi_t)} \mathbb{E} \left[\mathbf{1}_{N_{t,\tau_{t,\delta}}=0} | \mathcal{G}_t \right] + (\Delta_t)^- \mathbb{E} \left[\mathbf{1}_{N_{t,\tau_{t,\delta}}=1} | \mathcal{G}_t \right] \\
&\quad - \mathbb{E} \left[|R_1(\delta)| \mathbf{1}_{N_{t,\tau_{t,\delta}}=0} | \mathcal{G}_t \right] - \mathbb{E} \left[|R_2(\delta)| \mathbf{1}_{N_{t,\tau_{t,\delta}}=1} | \mathcal{G}_t \right]
\end{aligned} \tag{13}$$

Applying the Cauchy-Schwartz inequality to the remainder terms, then, we note that

$$\mathbb{E} \left[|R_1(\delta)| \mathbf{1}_{N_{t,\tau_{t,\delta}}=0} | \mathcal{G}_t \right] \leq \sqrt{\mathbb{E} \left[|R_1(\delta)|^2 | \mathcal{G}_t \right] \mathbb{E} \left[\mathbf{1}_{N_{t,\tau_{t,\delta}}=0} | \mathcal{G}_t \right]} = \sqrt{\mathbb{E} \left[|R_1(\delta)|^2 | \mathcal{G}_t \right] \exp(-\delta)}$$

and

$$\mathbb{E} \left[|R_2(\delta)| \mathbf{1}_{N_{t,\tau_{t,\delta}}=1} | \mathcal{G}_t \right] \leq \sqrt{\mathbb{E} \left[|R_2(\delta)|^2 | \mathcal{G}_t \right] \mathbb{E} \left[\mathbf{1}_{N_{t,\tau_{t,\delta}}=1} | \mathcal{G}_t \right]} \leq \sqrt{\mathbb{E} \left[|R_2(\delta)|^2 | \mathcal{G}_t \right] \delta \exp(-\delta)}$$

Replacing the remainder terms in (13) with these upper bounds, dividing through by δ , and then applying the bounded convergence theorem we immediately have:

$$\liminf_{\delta \rightarrow 0} \mathbb{E} \left[\left(\pi_t - \pi_{\tau_{t,\delta}} \right)^+ | \mathcal{G}_t \right] / \delta \geq \left(\frac{\partial_-}{\partial s} \pi_s^c \right)^- \bigg|_{s=t} \frac{1}{\lambda \bar{F}(\pi_t)} + (\Delta_t)^-$$

But by Lemma 10, we have that

$$\limsup_{\delta \rightarrow 0} \mathbb{E} \left[\left(\pi_t - \pi_{\tau_{t,\delta}} \right)^+ | \mathcal{G}_t \right] / \delta \leq \limsup_{\delta \rightarrow 0} \underline{\theta} \mathbb{E} [\tau_{t,\delta} - t | \mathcal{G}_t] / \delta = \frac{\underline{\theta}}{\lambda \bar{F}(\pi_t)}$$

where the equality follows from the bounded convergence theorem. The prior two inequalities now imply that a pricing policy π_t that satisfies the restricted sub-martingale constraint must also satisfy

$$\left(\frac{\partial_-}{\partial s} \pi_s^c \right)^- + \lambda \bar{F}(\pi_s) (\Delta_s)^- \leq \underline{\theta} \text{ a.s.}$$

for all s . ■

Lemma 10. *Let τ (assumed $\geq t$) be a bounded \mathcal{G}_t -stopping time and π_t a pricing policy satisfying the restricted sub-martingale constraint. Then,*

$$\mathbb{E} \left[\left(\pi_t - \pi_\tau \right)^+ | \mathcal{G}_t \right] \leq \underline{\theta} \mathbb{E} [\tau - t | \mathcal{G}_t] \text{ a.s.}$$

Proof. Without loss of generality, we prove the result for $t = 0$. Define

$$M_s = (\pi_0 - \pi_s)^+ - \underline{\theta} s$$

Now $M_0 = 0$, and observe that M_s is a \mathcal{G}_s -supermartingale:

$$\mathbb{E}[M_{s'} | \mathcal{G}_s] \leq M_s + \mathbb{E}\left[(\pi_s - \pi_{s'})^+ | \mathcal{G}_s\right] - \underline{\theta}(s' - s) \leq M_s$$

where the second inequality follows since π_s satisfies the restricted sub-martingale constraint. Doob's optional sampling theorem now yields the result. \blacksquare

The main result of this Section was the following verification theorem showing that a solution to the HJB equation did, in fact, constitute an optimal robust pricing policy:

Theorem 2. *Let π^* be the pricing policy obtained as the solution to the HJB equation, as defined in (10). Let π_s be any regular robust pricing policy. Then,*

$$\mathbb{E}\left[\int_0^T \lambda \pi_s \bar{F}(\pi_s) ds\right] \leq \mathbb{E}\left[\int_0^T \lambda \pi_s^* \bar{F}(\pi_s^*) ds\right].$$

We prove this theorem next:

B.1. Proof of Theorem 2

For any regular robust pricing policy, π_s , we denote, with a slight abuse of notation, the *process* $(H_\pi J)_s$ according to:

$$\begin{aligned} (H_\pi J)_s &= \frac{\partial}{\partial s} J(s, \pi_s, x_s) + \frac{\partial}{\partial p} J(s, p, x_s) \Big|_{p=\pi_s} \frac{\partial}{\partial s} \pi_s \\ &\quad + [\pi_s + J(s, \pi_s + \Delta_s, x_s - 1) - J(s, \pi_s, x_s)] \lambda \bar{F}(\pi_s) \end{aligned}$$

We define $(H_\pi J)_s = 0$ on the set where $x_s = 0$. We have the following result adapted from Brémaud [1981]:

Lemma 11 (Dynkin). *For any regular robust pricing policy π_s and $J \in \mathcal{J}$, we have:*

$$\mathbb{E}\left[\int_0^T \lambda \pi_s \bar{F}(\pi_s) ds\right] = J(0, \pi_0, x_0) + \mathbb{E}\left[\int_0^T (H_\pi J)_s ds\right].$$

Proof. We have

$$\begin{aligned} J(t, \pi_t, x_t) &= J(0, \pi_0, 0) \\ &\quad + \sum_{0 < T_n \leq t} [J(T_n, \pi_{T_n}, x_{T_n}) - J(T_{n-1}, \pi_{T_{n-1}}, x_{T_{n-1}})] + J(t, \pi_t, x_t) - J(\theta_t, \pi_{\theta_t}, x_{\theta_t}) \end{aligned} \quad (14)$$

where $\theta_t = \sup\{T_n : T_n \leq t\}$. Now,

$$\begin{aligned} J(T_n, \pi_{T_n}, x_{T_n}) - J(T_{n-1}, \pi_{T_{n-1}}, x_{T_{n-1}}) &= (J(T_n, \pi_{T_n}, x_{T_n}) - J(T_n, \pi_{T_n}, x_{T_n-})) \\ &\quad + (J(T_n, \pi_{T_n}, x_{T_n-}) - J(T_{n-1}, \pi_{T_{n-1}+}, x_{T_{n-1}})) + (J(T_{n-1}, \pi_{T_{n-1}+}, x_{T_{n-1}}) - J(T_{n-1}, \pi_{T_{n-1}}, x_{T_{n-1}})) \end{aligned}$$

But,

$$\begin{aligned} & J(T_n, \pi_{T_n}, x_{T_n-}) - J(T_{n-1}, \pi_{T_{n-1}+}, x_{T_{n-1}}) \\ &= \int_{T_{n-1}}^{T_n} \left[\frac{\partial}{\partial s} J(s, p, x_s) \Big|_{p=\pi_s} + \frac{\partial}{\partial p} J(s, p, x_s) \Big|_{p=\pi_s} \frac{\partial \pi_s}{\partial s} \right] ds, \end{aligned}$$

and

$$\begin{aligned} & \sum_{0 < T_n \leq t} (J(T_n, \pi_{T_n}, x_{T_n}) - J(T_n, \pi_{T_n}, x_{T_n-})) \\ &= \int_0^t (J(s, \pi_s, x_{s-} - 1) - J(s, \pi_s, x_{s-})) dN_s \end{aligned}$$

and

$$\begin{aligned} & \sum_{0 < T_n \leq t} (J(T_n, \pi_{T_n+}, x_{T_n}) - J(T_n, \pi_{T_n}, x_{T_n})) \\ &= \int_0^t (J(s, \pi_s + \Delta_s, x_{s-} - 1) - J(s, \pi_s, x_{s-} - 1)) dN_s \end{aligned}$$

Substituting in (14), we have a.s. (since $N_t - N_{t-} = 0$ a.s.):

$$\begin{aligned} J(t, \pi_t, x_t) &= J(0, \pi_0, 0) \\ &+ \int_0^t \left[\frac{\partial}{\partial s} J(s, p, x_s) \Big|_{p=\pi_s} + \frac{\partial}{\partial p} J(s, p, x_s) \Big|_{p=\pi_s} \frac{\partial \pi_s}{\partial s} \right] ds \\ &+ \int_0^t [J(s, \pi_s + \Delta_s, x_s - 1) - J(s, \pi_s, x_s)] \lambda \bar{F}(\pi_s) ds \\ &+ \int_0^t [J(s, \pi_s + \Delta_s, x_{s-} - 1) - J(s, \pi_s, x_{s-})] (dN_s - \lambda \bar{F}(\pi_s) ds) \end{aligned}$$

We apply the above representation for $t = T$, and note that $J(T, \cdot, \cdot) = 0$. Taking expectations, and adding the expected revenue under the pricing policy π_s , namely

$$\mathbb{E} \left[\int_0^T \lambda \pi_s \bar{F}(\pi_s) ds \right]$$

to both sides, we immediately have:

$$\begin{aligned}
\mathbb{E} \left[\int_0^T \lambda \pi_s \bar{F}(\pi_s) ds \right] &= J(0, \pi_0, 0) \\
&+ \int_0^t \mathbb{E} \left[\frac{\partial}{\partial s} J(s, p, x_s) \Big|_{p=\pi_s} + \frac{\partial}{\partial p} J(s, p, x_s) \Big|_{p=\pi_s} \frac{\partial \pi_s}{\partial s} \right] ds \\
&+ \int_0^t \mathbb{E} [\pi_s + J(s, \pi_s + \Delta_s, x_s - 1) - J(s, \pi_s, x_s)] \lambda \bar{F}(\pi_s) ds \\
&= J(0, \pi_0, 0) + \mathbb{E} \left[\int_0^T (H_\pi J)_s ds \right]
\end{aligned}$$

which is the result. ■

We can now proceed with the proof of Theorem 2. Let π be any regular robust pricing policy. We have by Proposition 1, that

$$\left(\frac{\partial_-}{\partial s} \pi_s^c \right)^- + \lambda \bar{F}(\pi_s) (\Delta_s)^- \leq \underline{\theta} \text{ a.s.}$$

Consequently, for any given state $y \in \mathcal{Y}$ defining the set $P(y) \triangleq \{\pi^M(y) : \pi^M \in \Pi^M\}$, we have for all s that

$$(\pi_s^c, \Delta_s) \in P(y) \text{ a.s.}$$

so that

$$(H_\pi J^*)_s \leq \sup_{\pi^M(y_s) \in P(y_s)} (H_\pi J^*)(y_s) = \sup_{\pi^M \in \Pi^M} (H_\pi J^*)(y_s) = 0 \text{ a.s.}$$

Consequently, Lemma 11 yields

$$\mathbb{E} \left[\int_0^T \lambda \pi_s \bar{F}(\pi_s) ds \right] \leq J^*(0, \pi_0, x_0).$$

Moreover, since under pricing policy π^* , we have

$$(H_{\pi^*} J^*)_s = \sup_{\pi^M(y_s) \in P(y_s)} (H_\pi J^*)(y_s) = \sup_{\pi^M \in \Pi^M} (H_\pi J^*)(y_s) = 0 \text{ a.s.}$$

Lemma 11 applied for that pricing policy yields

$$\mathbb{E} \left[\int_0^T \lambda \pi_s^* \bar{F}(\pi_s^*) ds \right] = J^*(0, \pi_0, x_0)$$

which completes the proof.