

MIT Open Access Articles

On the Efficacy of Static Prices for Revenue Management in the Face of Strategic Customers

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Chen, Yiwei and Vivek F. Farias. "On the Efficacy of Static Prices for Revenue Management in the Face of Strategic Customers." Proceedings of the 2016 ACM Conference on Economics and Computation, Maastricht, The Netherlands, ACM Press, July 2019.

As Published: <http://dx.doi.org/10.1145/2940716.2940737>

Publisher: ACM Press

Persistent URL: <https://hdl.handle.net/1721.1/128486>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



On the Efficacy of Static Prices for Revenue Management in the Face of Strategic Customers

April 11, 2017

Abstract

The present paper considers a canonical revenue management problem wherein a monopolist seller seeks to maximize revenues from selling a fixed inventory of a product to customers who arrive over time. We assume that customers are forward looking and strategize on the timing of their purchase, an empirically confirmed aspect of modern customer behavior. In the event that customers were myopic, foundational work by Gallego and Van Ryzin [1994] established that *static* prices are asymptotically optimal for this problem. In stark contrast, for the case where customers are forward looking, available results in mechanism design and dynamic pricing suggest a substantially more complicated prescription, and are often constrained by restrictive assumptions on customer type.

The present paper studies this revenue management problem while assuming that customers are forward looking and strategic. We demonstrate that for a broad class of customer utility models, static prices surprisingly continue to remain asymptotically optimal in the regime where inventory and demand grow large. We further show that irrespective of regime, an optimally set static price guarantees the seller revenues that are within at least 63.2% of that under an optimal dynamic mechanism. The class of customer utility models we consider is parsimonious and enjoys empirical support. It subsumes many of the utility models considered for this problem in existing mechanism design research; we allow for multi-dimensional customer types. We also allow for a customer's disutility from waiting to be positively correlated with his valuation. Our conclusions are thus robust and provide a simple prescription for a canonical RM problem that is near-optimal across a broad set of modeling assumptions.

1. Introduction

Consider the following canonical revenue management (RM) problem: a monopolist seller is endowed with an inventory of a product that she must sell over some fixed horizon via an anonymous,

posted price mechanism. Customers arrive over time and consider purchasing this product. Should a customer choose to make a purchase, he must pay a price equal to that posted at the time of his purchase. The goal of the seller is simply to maximize expected revenue from sales of the product.

The problem above is incredibly well understood in the setting where customers are *myopic*. Myopic customers either choose to make a purchase immediately upon arrival or else forego the opportunity to purchase the product and ‘leave the system’. Indeed, given the appropriate assumptions on the customer arrival process, this problem can be solved as a simple dynamic program. In fact, the seller can get away with doing something even simpler. Specifically, in an early foundational paper on revenue management, Gallego and Van Ryzin [1994], established that if the seller chose to maintain prices fixed at an appropriate level over the selling horizon, she was guaranteed to earn revenues that were close to those under an optimal dynamic pricing policy. Specifically, they showed that such a policy was asymptotically optimal in a regime where the seller had a large inventory and faced commensurately large demand. Gallego and van Ryzin’s insight is not merely theoretical; it has been borne out in a host of practical applications of revenue management ranging from the problem discussed above to substantially more complicated RM problems.

Technological changes have brought into question the basic assumptions made by Gallego and Van Ryzin [1994]. Specifically, thanks to the internet the search costs associated with ‘finding a deal’ have reduced dramatically. Consumers have the ability to monitor prices, obtain historical prices, and in certain cases, have access to tools that recommend the optimal timing of a purchase. A burgeoning body of empirical work has established that *forward looking* customers are fast becoming the norm. In summary, it is not uncommon for customers in the digital realm to strategize on the timing of their purchase. Forward looking customers must trade off the cost of delaying a purchase against the potential value of securing a discount. Any heterogeneity across customers in the nature of this trade-off introduces the potential for inter-temporal price discrimination on the part of the seller; a feature essentially absent from the model with myopic customers. Substantially less is known about the RM problem studied by Gallego and Van Ryzin [1994] in the setting where customers are forward looking. Indeed, if one were to take the path of assuming that customers were strategic¹, research on this problem divides broadly into two threads:

- One thread adopts the lens of dynamic mechanism design. In a nutshell, the research here has led to optimal (but complicated) mechanisms in models that require that all customers are homogenous in how they value their time. Loosening this assumption appears to lead to intractable mechanism design problems with multi-dimensional types. Approximation

¹as opposed to some boundedly-rational model of behavior

algorithms for the problem have also recently been proposed, but the mechanisms remain sophisticated with restrictive assumptions on customer utility models.

- A second thread of research forgoes optimal mechanisms, focusing instead on dynamic pricing policies with pre-announced price schedules. In the absence of inventory constraints, there has been progress in computing and characterizing the structure of optimal price schedules for several classes of customer utility models.

Against this backdrop, the present paper makes the following contribution:

We demonstrate that for a broad class of customer utility models, the fixed price policy is asymptotically optimal in the regime where inventory and demand grow large. We further show that irrespective of regime, an optimally set fixed price guarantees the seller revenues that are within at least 63.2% of that under an optimal dynamic mechanism.

This work thus bears a strikingly simple economic message: the asymptotic optimality of static prices established by Gallego and Van Ryzin [1994] extends to a general setting where customers are forward looking. For a broad class of customer utility models, the seller can only expect to gain a vanishingly small amount from dynamic policies and/or mechanisms that attempt to exploit the fact that customers may strategize on the timing of their purchase. As we will see, the class of customer utility models we consider is parsimonious. It subsumes a plurality of models considered in earlier research, some of which enjoy empirical support. At the outset we note two important features for the class of models we study: First, we permit the disutility incurred by a customer from delaying a purchase to be positively correlated with his valuation. Second, we allow for multi-dimensional customer types which permits for heterogeneity in both valuation as well as the cost of a delay. As we shall see, these features lend robustness to our conclusions.

We delay a detailed discussion of our contributions vis-a-vis the extant literature to Section 1.1. We note here that relative to mechanism design research on this problem, we not only derive near-optimal mechanisms for what has been thought to be a very difficult model, but in addition show that this can surprisingly be achieved with a simple fixed price policy. Relative to the dynamic pricing research on the special case of this problem that ignores inventory, we provide a crisp understanding of when static pricing policies suffice.

The paper proceeds as follows. Section 2 presents our model. That section discusses the key features of the utility model assumed for customers and establishes the parsimony of the model assumed. Section 2 also makes precise the problem solved by the customer and seller respectively.

Section 3 formulates the more general problem of designing a dynamic mechanism for the problem at hand and states our main performance guarantees for the fixed price policy. In establishing these results, our first step is to relax this mechanism design problem. This relaxation is presented in Section 4. In Section 5 we establish our performance guarantees by demonstrating lower bounds on the fixed price policy that are directly comparable to the upper bound derived from our relaxation. Section 6 provides a problem instance showing that our uniform performance guarantee is, in fact, tight. That section also seeks to illustrate the limits of our analysis by demonstrating a class of utility models that do not satisfy our assumptions and for which static prices can indeed be improved upon, even in the fluid regime. Finally, Section 7 concludes with comments on extensions to our results, as well as research questions that remain.

1.1. Related Literature

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. Among others, the text by Talluri and Van Ryzin [2006] and Özer and Phillips [2012] provide excellent overviews of this area.

As already discussed, Gallego and Van Ryzin [1994] is a foundational revenue management paper that is particularly pertinent to the present paper. Those authors introduce a model akin to the one we study here, except with myopic customers. The main insight in this foundational paper is that appropriately set static prices are asymptotically optimal in a setting where available inventory and demand grow large. It has become amply clear that the assumption of myopia is fast becoming untenable in revenue management. Specifically, empirical work, most notably by 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 present paper can thus be seen as extending the conclusions of Gallego and Van Ryzin [1994] to the setting where customers are forward looking, for a broad class of customer utility functions.

As described earlier, the antecedent literature most relevant to the present paper divides roughly into two groups. The first of these studies the problem from a mechanism design perspective, whereas the second focuses attention on the design of optimal price schedules.

Dynamic Mechanism Design: The problem we study can naturally be seen as one of dynamic mechanism design. An early paper by Vulcano et al. [2002] considers short-lived but strategic

customers arriving in sequential batches over a finite horizon and proposes running a modified second price auction in each period (as opposed to dynamic pricing). Gallien [2006] provides what is perhaps the first tractable dynamic mechanism for a classical 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. This work assumes that a customer’s value for the product depreciates exponentially at a constant rate that is common knowledge. Board and Skrzypacz [2013] consider a discrete time version of the same model, and assuming a *finite* horizon, compute the optimal dynamic mechanism. Board and Skrzypacz [2013] also require that all customers discount at a homogenous rate that is common knowledge. The mechanism they propose consists of a ‘hybrid’ of a dynamic pricing mechanism with an end-of-season ‘clearing’ auction. The homogeneity required for discount rates in these models is limiting. Besbes and Lobel [2015] make the excellent point that not permitting heterogeneity in customers’ sensitivity to a delay might artificially limit the impact of inter temporal price discrimination, and consequently artificially mitigate the need for dynamic pricing.

Pai and Vohra [2013] consider a substantially more general model of (finite horizon) RM with forward looking customers. Customers in their model have heterogeneous ‘deadlines’ as opposed to discounting. When these deadlines are known to the seller (a strong assumption), 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 recent work, Chen and Farias [2015] consider a model that allows for heterogeneity in customers’ disutility from delaying a purchase. The authors introduce a class of ‘robust’ dynamic pricing policies which they show are guaranteed to garner expected revenue that is at least 29% of the revenue yielded under the seller’s optimal direct dynamic mechanism. The class of utility models we study subsumes those studied by Chen and Farias [2015], and as already discussed we allow for a customer’s valuation and his disutility from a delay to be positively correlated which is something Chen and Farias [2015] do not permit. Another relevant mechanism design paper is by Haghpanah and Hartline [2015]. Their work can be seen as an elegant generalization to the celebrated result of Stokey [1979]. One (coarse) interpretation of their result in the RM context is as follows: They establish in the setting where inventory is infinite that myopic behavior is optimal on the customer’s part with the corresponding optimal mechanism for the seller being an anonymous posted price set to the static revenue maximizing price. They do so while assuming that

the customers’ loss in value from a delay is private information. There is of course no competition among customers in this setting – a fact that is essential to the result. In our setting, inventory is finite and this makes for a fundamental change to the problem. A customer must now compete with other customers (as opposed to just future versions of himself). And he must do so with asymmetric information.

Relative to this past work that takes a mechanism design approach, and ignoring distinctions such as discrete time modeling vs. continuous time modeling etc., we consider a general setting. Specifically, we allow for a rich class of customer utility models. We allow for heterogeneity in the sensitivity to delay and we allow for inventory to be limited. Despite this generality, we show that a simple policy – a fixed price – is asymptotically optimal, while simultaneously providing a uniform performance guarantee.

Setting Price Schedules in the Presence of Strategic Customers: A second stream of relevant literature foregoes optimal mechanisms to focus exclusively on committing to (potentially time varying) price schedules. Among the first papers in this vain is Stokey [1979]. She considered a class of customer utility functions subsumed by the model we study wherein the functional form prescribing a customers sensitivity to delay is common knowledge. Her paper arrives at ‘the unexpected conclusion’ that the seller will forego the opportunity to price discriminate entirely, setting prices at the static revenue maximizing price. As discussed above, those conclusions have been strengthened substantially by Haghpanah and Hartline [2015] using cutting edge techniques from dynamic mechanism design. Our work can be seen as taking this insight further to the harder revenue management setting (where inventory is a constraint) while simultaneously allowing for a very general class of utility models and customer heterogeneity along multiple dimensions.

Borgs et al. [2014] is among the first RM papers that consider a monopolist with the power to commit to a price schedule. The authors consider 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 (modeling a service system). Customers have arrival times, deadlines and valuations; valuations are assumed independent of the arrival time and deadline. In addition, the seller knows the fraction of customers corresponding to each arrival time-deadline pair. Borgs et al. [2014] show how to compute the optimal price schedule for this setting. 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.

Continuing on this theme, Besbes and Lobel [2015] consider an infinite horizon model wherein customers arrive to the system over time and strategize on their time of purchase. Inventory constraints are not considered. Customers have valuations and a willingness to wait that may be correlated with their valuation. The authors establish an elegant result – they characterize the optimal price schedule as being cyclic and also provide an efficient algorithm for its computation. In our lexicon, the disutility model considered by the authors is effectively a step function – a customer incurs no disutility if he makes the purchasing decision before the deadline, otherwise, his disutility is equal to his valuation. Consequently, viewed as function of valuation for some fixed allocative decision, the disutility function contains ‘jumps’. A key requirement for our result will be that for a given allocative decision, disutility not increase ‘too quickly’ with valuation; a requirement that such a function clearly cannot fulfill. In fact, we will later show that the sufficiency of static prices rests precisely on the rate of at which disutility increases with valuation. Loosely speaking, as long as this increase is sub-linear (a condition we will see is implied by a large number of models considered in the theoretical and empirical literature), static prices suffice. If on the other hand, the increase can be rapid (such as a step function) we show that in fact static prices do not suffice, even in an asymptotic regime.

There are a number of additional examples of this theme in recent RM literature. Caldentey et al. [2015] take a novel view of uncertainty and consider the dynamic pricing problem in a minimax setting that allows those authors to capture uncertainty in customer valuation as well as arrival times, thereby taking a ‘robust’ view of custom type as opposed to the prior driven view taken by all of the other literature we have discussed, as well as the present paper. Liu and Cooper [2015] and Lobel [2017] both study settings where as opposed to being strategic, customers are ‘patient’, a behavioral model in the mould of satisficing. Both those papers identify and show how to compute optimal cyclic pricing policies. It is interesting to note that other researchers have motivated cyclic pricing policies by considering price reference effects; see Hu et al. [2016], Wang [2016].

Vis-a-vis the work above on setting optimal price schedules, our work sheds light on the conundrum of when to use ‘promotions’ (or non-static price paths) versus ‘everyday low prices’ (or static prices). We provide a crisp understanding of when the latter suffices for RM problems. The conditions we identify for the sufficiency of static prices are evidently fairly general inasmuch as they capture modeling assumptions in antecedent literature: We allow for customers to be heterogenous in both their valuation as well as parameters impacting their sensitivity to a delay. We assume inventory is limited. We also permit a customer’s disutility from waiting to be positively correlated with his valuation. In summary, we establish that a simple fixed price policy is, to a first order,

optimal for a broad set of assumptions around the canonical RM problem with strategic customers.

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 . With a slight abuse of notation², we denote the inventory process by X_t and the corresponding sales process by N_t . Of course, $N_t = x_0 - X_t$. We require that π_t depend only on the history of the pricing and sales process³. In addition, we require $\pi_t = \infty$ if $X_{t-} = 0$, and $\pi_t < \infty$ otherwise.

Customers arrive over time according to a Poisson process of rate λ ; an extension to non-homogenous processes is presented later. A customer arriving at time t is endowed with a valuation, $v \in \mathbb{R}_+$, and a collection of K attributes, $\theta \in \mathbb{R}_+^K$. As we will see shortly, θ and v will jointly parameterize the customer's disutility from 'staying in the system'. We denote by ϕ , the 'type' of an arriving customer which we understand to be the tuple

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

Denote by Φ the set of all types ϕ . 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 a customer of type ϕ chooses to delay making a purchase decision to time $\tau_\phi \geq t_\phi$. We let a_ϕ indicate whether the customer leaves having purchased a unit of the product. Specifically, if the seller has inventory to allocate⁴ and if a purchase provides the customer greater utility than no purchase then $a_\phi = 1$; otherwise $a_\phi = 0$. If $a_\phi = 1$, the customer pays the seller the amount $p_\phi = \pi_{\tau_\phi}$. The customer then garners utility

$$U(\phi, y_\phi) = (v_\phi - p_\phi) a_\phi - M(\phi, y_\phi),$$

where we define the tuple $y_\phi \triangleq (\tau_\phi, a_\phi, p_\phi)$. The function $M(\cdot, \cdot)$ captures the customer's disutility from delaying his purchase to time τ_ϕ . The structure of $M(\cdot, \cdot)$ will play a significant role in the sequel as it encodes the dependence of the customer's cost to delaying a purchase on his type. We

²The abuse of notation here is that we choose to omit the dependence of X_t and N_t on π_t ; this will always be clear from context.

³More formally, we require π_t to be left continuous, and adapted to \mathcal{F}_{t-} where $\mathcal{F}_t = \sigma(\pi^t, X^t)$.

⁴Multiple customers revealing themselves to the seller at the same time are allocated inventory in random order.

will discuss our assumptions on this structure shortly.

We assume that a customer's type ϕ is private information. Recall that a customer's type is specified by the tuple $(t_\phi, v_\phi, \theta_\phi)$. This is in contrast with the typical model that specifies type based only on time of arrival and valuation, i.e. (t_ϕ, v_ϕ) ; see, for instance, Aviv and Pazgal [2008], Board and Skrzypacz [2013], Caldentey and Vulcano [2007], Gallien [2006], Yin et al. [2009]. Putting aside the technical challenge this creates, doing so is important from a modeling perspective. For instance, as we shall see it lets us model the fact that customer type is determined not just by valuation but also sensitivity to delays, something that cannot be modeled via the more restrictive type specification.

Recall that the arrival times t_ϕ are the points of a Poisson process. We assume that the valuation v_ϕ is independent of the arrival time t_ϕ . This assumption is in analogy with a large body of the research on revenue management for strategic customers. Aviv and Pazgal [2008], Besbes and Lobel [2015], Board and Skrzypacz [2013], Gallien [2006], Vulcano et al. [2002], Yin et al. [2009] all make such an assumption and point out that a primary motivation for dynamic pricing is inter-temporal price discrimination which remains relevant despite the assumption.

Since the only quantity dependent on θ_ϕ is the disutility function, $M(\cdot, \cdot)$, we assume that v_ϕ and θ_ϕ are independent. Specifically, we may exchange any assumptions on correlation between v_ϕ and θ_ϕ with assumptions on the structure of $M(\cdot, \cdot)$. To see how, notice that if indeed these random variables were dependent, we could always construct a common probability space on which we write θ_ϕ as some function, say $h(\cdot, \cdot)$, of v_ϕ and $\hat{\theta}_\phi$ where $\hat{\theta}_\phi$ is indeed independent of v_ϕ . We can then obtain an equivalent problem by employing the disutility function \hat{M} defined according to:

$$\hat{M}\left((t_\phi, v_\phi, \hat{\theta}_\phi), y\right) = M\left((t_\phi, v_\phi, h(v_\phi, \hat{\theta}_\phi)), y\right).$$

Put a different way, any restrictions to the nature of the dependence can simply be captured by structural assumptions on the disutility function $M(\cdot, \cdot)$ which we discuss shortly. We prefer the latter approach as it leads to making the assumptions on the nature of such a dependence concrete. We make no assumptions on the correlation between t_ϕ and θ_ϕ .

We assume that customer valuations have a cumulative density function given by $F(\cdot)$, and have a density function, denoted by $f(\cdot)$. We denote $\bar{F}(\cdot) \triangleq 1 - F(\cdot)$. We make a standard assumption on the valuation distribution:

Assumption 1. *The virtual value function of the valuation distribution, $v - \frac{\bar{F}(v)}{f(v)}$, is non-decreasing in v and has a non-negative root v^* .*

In the remainder of this section, we first discuss the assumptions we place on the disutility model. We will then move on to presenting the problems faced by a customer in timing his decision whether and when to purchase as well as that faced by the revenue manager who must dynamically adjust prices knowing only the history of prices and of purchases made thus far.

2.1. The Disutility Model

The structure of the disutility function $M(\cdot, \cdot)$ captures precisely the dependence of the customer's cost to delaying a purchase on the customer's type. We will place a set of structural restrictions on $M(\cdot, \cdot)$ that are general enough to capture a variety of realistic models. Specifically, we assume:

Assumption 2. *For any type $\phi \in \Phi$, and any $y \triangleq (\tau_y, a_y, p_y)$ with $\tau_y \geq t_\phi$, we have:*

1. $M(\phi, y) \geq 0$.
2. If $\tau_y = t_\phi$, then $M(\phi, y) = 0$.
3. $M(\phi, y)$ is differentiable with respect to v_ϕ ; denote $m(\phi, y) \triangleq \frac{\partial}{\partial v_\phi} M(\phi, y)$.
4. $M(\phi, y)$ is non-decreasing and concave in v_ϕ .

Let us interpret the conditions imposed by the assumptions on $M(\cdot, \cdot)$: The first assumption simply formalizes our interpretation of $M(\cdot, \cdot)$ as a *disutility*. The second assumption effectively normalizes the disutility function, requiring it to be zero for a delay of zero. Together with the first assumption this implies that all else being the same (i.e. for a given allocative decision a_y , and price p_y), the customer would prefer no delay ($\tau_y = t_\phi$) over a positive delay ($\tau_y > t_\phi$). The third assumption is made for analytical convenience and we do not believe it fundamental for the conclusions in this paper. The assumption simplifies our analysis and lends itself to notational clarity. The fourth assumption captures the essence of the structure we impose on the disutility incurred due to a delay, and consists of two components. The first, is that this disutility is increasing in the customer's valuation so that high value customers incur a larger cost to delaying a purchase than those that place a lower value on the product. This assumption is natural and has widespread support in both theoretical and empirical literature. [Stokey, 1979] is a foundational paper on intertemporal price discrimination to make such an assumption. Modern papers in RM and service operations more generally, also make such an assumption; see, for instance, Afèche and Pavlin [2016], Aviv and Pazgal [2008], Board and Skrzypacz [2013], Doroudi et al. [2013], Gallien [2006], Gurvich et al. [2016], Katta and Sethuraman [2005], Kilcioglu and Maglaras [2015], Moon et al. [2015], Nazerzadeh and Randhawa [2015]. The second part of the assumption can be interpreted as

controlling the *rate* at which this disutility can grow with the customer’s value. Our requirement of concavity implies that this growth must be sub-linear. We will see shortly that this assumption again finds widespread support in the literature.

As it turns out, a number of concrete examples of disutility functions considered in the literature fit the assumptions above. We discuss these families of disutility functions next:

Monitoring Cost: Starting with the classical work of Diamond [1971], a common assumption in the economics literature on pricing that results in the ability to violate the so-called law of one price, has been the presence of a ‘search’ or monitoring cost. The notion of search cost here could correspond to any effort the customer might expend in monitoring prices. It is further worth noting that a search cost model has been empirically verified to provide a good fit in an empirical study of customer purchasing decisions at a clothing retailer that practices dynamic pricing [Moon et al., 2015]. A natural model for search cost would simply assume that it grows linearly in the time the customer monitors prices. Specifically:

$$M(\phi, y) = \theta_\phi (\tau_y - t_\phi)^+$$

where $\theta_\phi > 0$ is the unit-time search cost incurred by a customer of type ϕ . This is a canonical model in the economics literature; see for example Anderson and Renault [1999], Rob [1985] or Ellison and Wolitzky [2012]. Clearly, this model satisfies the requirements of Assumption 2.

Recall that we require that θ_ϕ be independent of v_ϕ so the unit time search cost above is independent of valuation. But one may easily go further and specify an explicit dependence of search cost on v_ϕ ; for instance:

$$M(\phi, y) = \theta_\phi h(v_\phi) (\tau_y - t_\phi)^+.$$

If $h(\cdot)$ were a non-negative, non-decreasing, concave function, then again, this more general disutility function satisfies the requirements of Assumption 2. A number of recent pieces of research that attempt to model customer’s disutility from a delay in service systems, including Afèche and Pavlin [2016], Doroudi et al. [2013], Gurvich et al. [2016], assume such a model taking $h(\cdot)$ to be a linear function. Nazerzadeh and Randhawa [2015] and Katta and Sethuraman [2005] assume that $h(\cdot)$ is sub-linear; a closely related but slightly more general function class than the concave functions we permit.

Finally, we could generalize the model further, specifically by taking

$$M(\phi, y) = \theta_\phi h(v_\phi) g(\tau_y - t_\phi).$$

If in addition to the earlier requirement on $h(\cdot)$, $g(\cdot)$ were a non-negative function with $g(0) = 0$, we would still satisfy the requirements of Assumption 2. This would allow us in turn to capture models of disutility with more general dependences on delay, such as those in Ata and Olsen [2009] or Afèche and Mendelson [2004].

Exponential Discounting: In addition to monitoring costs, disutility could also arise because the product is ‘perishable’ so that its value to the customer decays over time. A canonical model for this sort of disutility arises as follows: One assumes that the useful lifetime of a perishable product to a customer of type ϕ following his arrival is exponentially distributed with parameter θ_ϕ . If the customer actually received the product at a time $\tau_y > t_\phi$, his expected disutility from the delay (due to the loss in the usable lifetime of the product) is then simply

$$M(\phi, y) = v_\phi a_y (1 - \exp(-\theta_\phi (\tau_y - t_\phi))).$$

Put a different way, this equivalently states that

$$U(\phi, y) = a_y (v_\phi \exp(-\theta_\phi (\tau_y - t_\phi)) - p_y)$$

which in turn is a canonical model both in the economics oriented literature on dynamic pricing for perishable products such as Board and Skrzypacz [2013] but also the revenue management literature eg. Aviv and Pazgal [2008], Gallien [2006]. Of course, it is easy to see that this model of disutility also satisfies the requirements of Assumption 2.

The above are merely examples of disutility functions that satisfy Assumption 2. They serve to illustrate that while we do indeed need to place *some* restrictions on the nature of the disutility function, the assumptions we have placed are capable of capturing important phenomena. We next discuss the problems faced by the customer and the seller respectively.

2.2. The Customer and Seller Problems

The dynamic pricing policy π utilized by the seller is assumed to be common knowledge. Recall that this policy can depend only on the sales process and historical prices. In particular, the seller

does not have the ability to observe customers who have delayed their purchase and remain in the system nor customers that left without making a purchase, either immediately upon arrival or after some delay. The seller is assumed to have the power to commit to the pricing policy. This assumption now enjoys excellent support in the revenue management setting thanks to antecedent research. See Liu and Van Ryzin [2008] for a comprehensive justification from an RM perspective, or Board and Skrzypacz [2013] for one from an economic perspective.

Now consider a customer of type ϕ who decides to reveal himself to the seller and make a purchase decision at some time $t \geq t_\phi$. Of course, if no inventory is available at that time, the price posted by the seller is formally infinite, so that the customer will choose to leave without making a purchase (so that $a_\phi = 0$). If inventory is available, so that $X_{t-} > 0$ and no other customers present themselves at time t , the customer chooses to make a purchase (setting $a_\phi = 1$) if and only if doing so yields at least as much utility as not making a purchase. Finally, if multiple customers present themselves to the seller at time t (an unlikely event but one that cannot be ruled out), the seller allocates them inventory in random order. If the number of customers exceeds the remaining inventory, then clearly some customers will not be allocated inventory; denote by A_t^ϕ the random indicator that the seller allocates inventory to customer ϕ if he presents himself to make a purchase at time t . Of course, $A_t^\phi = 1$ is $X_{t-} > 0$ and ϕ is the only customer to request a unit of the product at time t and $A_t^\phi = 0$ if $X_{t-} = 0$. In summary the maximum utility that customer ϕ can garner should he decide to reveal himself to the seller and make a purchasing decision at time t is given by:

$$U^*(\phi, t) \triangleq \begin{cases} -M(\phi, (t, 0, 0)) \vee (v_\phi - \pi_t - M(\phi, (t, 1, \pi_t))) & \text{if } X_{t-} > 0 \text{ and } A_t^\phi = 1 \\ -M(\phi, (t, 0, 0)) & \text{otherwise} \end{cases}$$

Customers strategize about the time of their purchase and employ stopping rules contingent on their type that constitute a symmetric Bayes Perfect Nash equilibrium (BPNE). Such an equilibrium can be formally defined by a map τ^π from types to stopping rules⁵. In particular, at each point of time t , all customers can observe historical prices up to time t . For customer type ϕ , $\tau^\pi(\phi) \triangleq \tau_\phi$ is a stopping rule with respect to the filtration generated by the historical price process $\mathcal{P}_t = \sigma(\{\pi_s : s \in [0, t]\})$. The stopping rule is derived as a solution to the optimal stopping problem

$$\sup_{\tau_\phi \geq t_\phi} \mathbf{E}_{-\phi} [U^*(\phi, \tau_\phi) | \mathcal{P}_{t_\phi}],$$

⁵In the sequel we will at times, with an abuse of notion, use this map and the corresponding stopping rules interchangeably

where the expectation assumes that other customers also employ type dependent stopping rules given by τ^π . We will later demonstrate existence of such an equilibrium stopping rule for a specific pricing policy. We do not prove existence in general.

Now consider that the seller employs the pricing policy (or process) π , and let τ^π be an equilibrium stopping rule for such a policy. The seller's revenue is then given by

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

where $\hat{\tau} = \inf\{t : X_t = 0\}$. N_t is the sales process where, as noted earlier, we have suppressed the dependence of N_t on π and τ^π . The task of finding an 'optimal' policy is an apparently challenging one. In fact, simpler problems than this are already intractable: First, the customer stopping rule τ^π (upon which N_t depends) is for general pricing policies, a potentially complicated and hard to characterize function of π . That is, even having fixed a policy π , characterizing an equilibrium stopping rule is in general a challenging task. Second, the potential presence of customers in the system over an extended period of time (as they contemplate a purchase) induce long range dependencies in the pricing process, so that even given a fixed stopping rule (i.e. fixing customer behavior), finding an optimal pricing policy may not be a simple task in that traditional dynamic programming approaches fail. In summary, the seller's problem of finding an optimal pricing policy (assuming such a policy exists) is intractable for the model we have described so far. Even assuming we could surmount these challenges, other issues remain. For instance, it may be difficult to calibrate such a policy to data given that type distributions would need to be inferred from transactions. If the pricing policy chosen by the seller induced complex equilibrium stopping rules, the predictive power of the model might be an issue.

So motivated, we will in the next Section, take the approach of computing an upper bound on *any* pricing policy, and illustrate the power of a simple, fixed price policy by comparing the revenues the seller can hope to earn under that policy to our upper bound. Our approach to computing an upper bound will be driven by viewing the seller's problem through the lens of dynamic mechanism design. Since the class of dynamic mechanisms subsumes the class of dynamic pricing policies, we can construct a dynamic mechanism design problem that yields an upper bound to the seller's revenue under any dynamic pricing policy. We illustrate that fixed prices continue to remain powerful in the setting where customers are forward looking.

3. Static Prices and an Optimal Dynamic Mechanism Benchmark

As discussed in the introduction, Gallego and Van Ryzin [1994] proposed, in a seminal piece of work, the use of a simple *static price policy* for revenue management problems of the type we have just discussed. They showed that in the setting where inventory and the customer arrival rate grow large, such a policy is asymptotically optimal. The paper is considered seminal for its simple message to practitioners: static prices are to a first order, optimal; dynamic pricing can only hope to capture second order benefits. Of course, in settings where customers strategize on the timing of their purchase – the topic of this paper – it is no longer clear that static prices retain this desirable property. In fact, the *raison d'être* for ‘promotional pricing’ is inter-temporal price-discrimination that seeks to arbitrage differences in the disutility incurred by customers from a delayed purchase.

The primary economic insight of this paper is that, in fact, the value of this inter-temporal price discrimination is limited for the class of utility models we consider. Specifically, it continues to be the case that the static pricing policy is, to a first order, optimal – it approaches the revenues earned under an optimal dynamic mechanism in the very regime studied by Gallego and Van Ryzin [1994]. In addition, we will show that such a policy is also guaranteed to earn a constant factor (roughly 63%) of the revenue of an optimal dynamic mechanism irrespective of regime.

In what follows in this Section, we first recall the static price policy proposed by Gallego and Van Ryzin [1994]. Under this policy, it is a dominant strategy for customers to not delay a purchase decision – an attractive property from the perspective of the seller. We will then turn to producing upper bounds on performance under *any* pricing policy and to that end consider the still more general task of producing an optimal dynamic mechanism. Finally, we will state our main results. Subsequent sections are devoted to establishing these results.

3.1. A Candidate Static Price Policy

We define and briefly motivate static price policies: Let $\hat{\pi}$ be an arbitrary measurable function from $[0, T]$ to \mathbb{R}_+ , and consider the following ‘fluid’ optimization problem:

$$\begin{aligned} \max_{\hat{\pi}} \quad & \int_0^T \lambda \hat{\pi}_t \bar{F}(\hat{\pi}_t) dt \\ \text{s.t.} \quad & \int_0^T \lambda \bar{F}(\hat{\pi}_t) dt \leq x_0 \end{aligned} \tag{1}$$

This problem treats customers as myopic and infinitesimal (hence ‘fluid’). It is easy to show (see Gallego and Van Ryzin [1994]) that the optimal solution to this program is given precisely by the

static price policy $\{\pi_t^{\text{FP}} : t \in [0, T]\}$ defined according to⁶

$$\pi_t^{\text{FP}} = \bar{F}^{-1} \left(\min \left\{ \frac{x_0}{\lambda T}, \bar{F}(v^*) \right\} \right) \triangleq \pi^{\text{FP}}$$

so that the optimal value of the program is precisely $\lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}})$. Now observe that if the seller implements π^{FP} , it is a (weakly) dominant strategy for customers to not delay a potential purchase (or leave immediately if no purchase is made):

Lemma 1. *For the static pricing policy, π^{FP} , the myopic stopping rule, $\tau_\phi^{\text{FP}} = t_\phi$ is weakly dominant.*

The proof of this fact is immediate from the definition of $U^*(\phi, t)$: under any fixed price policy, $U^*(\phi, t)$ is non-increasing on $t \geq t_\phi$ on every sample path. In fact, if the disutility function $M(\cdot, \cdot)$ were strictly positive for positive delays, myopic behavior is a strongly dominant strategy.

Now observe that if customers behaved myopically (so that $\tau_\phi = t_\phi$), then the event that two customers present themselves simultaneously to the seller has measure zero, and $a_\phi^{\text{FP}} = \mathbf{1} \{v_\phi \geq \pi^{\text{FP}}\}$. Consequently, the sales process is a Poisson process with intensity $\lambda \bar{F}(\pi^{\text{FP}})$ so that

$$J_{\pi^{\text{FP}}, \tau^{\text{FP}}} (x_0, T) = \pi^{\text{FP}} \mathbb{E} \left[\min \left(N_{\pi^{\text{FP}}}^{\lambda T}, x_0 \right) \right]$$

where $N_{\pi^{\text{FP}}}^{\lambda T}$ is a Poisson random variable with mean $\lambda T \bar{F}(\pi^{\text{FP}})$. We next set out to construct a benchmark policy with which to compare the revenue under this fixed price policy.

3.2. A Dynamic Mechanism Design Benchmark

As discussed earlier, the task of optimizing over pricing policies is a non-trivial one and even characterizing the optimization problem appears to be a challenging task. As such, our goal in this section is to produce an upper bound, which we will denote $J^*(x_0, T)$, on the revenue under *any* pricing policy. Specifically, we require that for any pricing policy π , and an associated stopping rule τ^π ,

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

We will produce this upper bound by allowing the seller to employ a general dynamic mechanism for the problem at hand. Specifically, dynamic mechanisms subsume dynamic pricing policies (in the sense of strategic equivalence), so that the seller's revenue under the optimal dynamic mechanism serves as an upper bound on the revenue the seller can earn under any dynamic pricing policy. We care about the dynamic mechanism design problem only in so much as it yields a useful upper

⁶We will abuse notation slightly by also using π^{FP} to denote the static price policy itself.

bound so that issues concerning the practical relevance of a general dynamic mechanism are not relevant to our discussion.

To set up the dynamic mechanism design problem, we begin by introducing some relevant notion. We denote by $h^t \triangleq \{\phi : t_\phi \leq t\}$ the set of 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 unlike the dynamic pricing setting the customer must *explicitly report his type* ϕ to the seller in this setup (although he may potentially lie). The seller then determines whether the customer is allocated a good, when he is allocated the good, and at what price according to y_ϕ . Note that y_ϕ may depend on the reports of some subset of customers, but the structure of this dependence must be causal and satisfy other constraints that we now formalize.

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 allocation decisions 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.
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. Customers pay nothing to participate, so that $p_\phi = 0$ 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., $t_{\hat{\phi}} \geq t_\phi$).

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

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

Denote by $J^*(x_0, T)$ the optimal value obtained in the problem above. Chen and Farias [2015] establish that for any dynamic pricing policy there exists a direct dynamic mechanism that satisfies the constraints of (2) and has objective value equal to the seller's revenue under the dynamic pricing policy, thereby establishing the following lemma:

Lemma 2. (*Valid Benchmark*) *For any pricing policy and corresponding stopping rule, (π, τ^π) , we have that*

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

The upshot of this result is that we now have an upper bound on what the seller can hope to attain under any dynamic pricing policy that we can characterize as the optimal value to a more familiar – but still challenging – optimization problem, namely (2). In a subsequent section, we will further analyze this upper bounding optimization problem to facilitate a comparison with the revenues under the static pricing policy described in the previous section. The second salient point worth discussing here is that the upper bound we have set up is with respect to a substantially broader class of mechanisms than simply those that correspond with anonymous dynamic pricing. As revenue management evolves, it stands to reason that sellers may want to experiment with approaches to selling that transcend the traditional anonymous posted price approach; in practice we see experiments with rebates, auction formats, and the like. Assuming, we are able to show that static price revenues compare favorably with our upper bound, we will have established that such pricing policies are desirable not just in comparison with general dynamic pricing policies, but with respect to any (reasonable) mechanism the seller might hope to concoct.

3.3. Principle Results

Our principle results establish that static pricing policies offer surprisingly strong performance, even in the face of strategic customers. Specifically, we compare the revenue the seller may hope to earn under the static pricing policy, namely the quantity $J_{\pi^{\text{FP}}, \tau^{\pi^{\text{FP}}}}(x_0, T)$, with an upper bound on the revenue she may hope to earn under essentially any reasonable selling mechanism, namely

$J^*(x_0, T)$. We will consider two types of performance bounds. The first, will be in the ‘fluid’ regime studied originally by Gallego and Van Ryzin [1994], where inventory and the scale of demand grow large simultaneously. Now our setting mimics that of Gallego and Van Ryzin [1994] with the obvious exception that we allow customers to be forward looking (as opposed to myopic and short lived). Consequently, given the broad impact of the original performance guarantee provided by Gallego and Van Ryzin [1994], a performance guarantee in the fluid regime has obvious value. Other authors have already noted that this sort of scaling preserves the potential relevance of inter-temporal price discrimination and mechanism design more generally eg. [Besbes and Lobel, 2015, Liu and Cooper, 2015]. Our second performance guarantee will be uniform and non-asymptotic, i.e. it will be relevant over all parameter regimes. We will show that this latter guarantee is also tight in the sense that a specific problem instance achieves the bound implicit in the guarantee.

The Fluid Regime: Following Gallego and Van Ryzin [1994], we consider a sequence of problems, parameterized by n . In the n th problem, we have initial inventory $x_0^{(n)} = nx_0$, and customers arrive at the rate $\lambda^{(n)} = n\lambda$. We denote by a superscript (n) quantities relevant to the n th model in this scaling. So, for instance, $J_{\pi_{\text{FP}}, \tau_{\pi_{\text{FP}}}}^{(n)}(nx_0, T)$ denotes the revenue under the static price policy in the n th model. Colloquially, as n grows, we are scaling the inventory and volume of demand in the problem instance. All other aspects of the model – namely the customer utility model, and the horizon T , stay unchanged. Our main result is a guarantee that shows that in the ‘fluid regime’, the static price policy is asymptotically optimal. Specifically:

Theorem 1. *Provided Assumptions 1 and 2 are satisfied, we have:*

$$\frac{J_{\pi_{\text{FP}}, \tau_{\pi_{\text{FP}}}}^{(n)}(nx_0, T)}{J^{*,(n)}(nx_0, T)} = 1 - O\left(\frac{1}{\sqrt{n}}\right).$$

This result makes a strikingly simple economic statement. Static pricing policies constitute, to a first order, an optimal selling mechanism; any gains one may hope to make from dynamic pricing and/or sophisticated selling mechanisms must necessarily contribute a vanishingly small incremental revenue to the seller. Theorem 1 provides a significant generalization to the conclusions drawn by Gallego and Van Ryzin [1994]. Whereas their conclusion rested heavily on the assumption that customers were myopic and short-lived, Theorem 1 shows that those conclusions are robust to potentially long lived customers that strategize on the timing of their purchase.

Our second result, relaxes the requirement of a fluid regime and applies uniformly across all parameter regimes:

Theorem 2. *Provided Assumptions 1 and 2 are satisfied, we have:*

$$\frac{J_{\pi^{\text{FP}}, \tau^{\pi^{\text{FP}}}}(x_0, T)}{J^*(x_0, T)} \geq 1 - \frac{1}{e}.$$

In addition, there exists a problem instance for which this lower bound on performance is achieved.

This result complements our fluid regime result by stating that irrespective of regime or parameter settings, the static price policy will always achieve at least $\sim 63.2\%$ of the revenue the seller can hope to earn under *any* dynamic mechanism. This analysis is, in fact, tight so that we can construct a problem instance that achieves this bound. Constant factor guarantees of this nature have assumed a place of prominence in a number of operational problems ranging from revenue management to inventory and supply chain management. We interpret this guarantee as a strong indicator of the robustness of static prices across parameter regimes.

The remainder of this paper is dedicated predominantly to establishing the two principle results above. We proceed in two broad steps. First, in Section 4 we analyze the dynamic mechanism design problem, (2) that provides us with our benchmark. Then in Section 5, we explicitly compare the revenue under the static price policy to this benchmark using the tools developed in Section 4 among other techniques. Finally, in a concluding Section we discuss the limitations of our analysis (or equivalently the family of disutility functions that our analysis is restricted to), and present a simple computational experiment to illustrate the spectrum of performance possible between our constant factor lower bound, and our asymptotic optimality result.

4. Analyzing the Dynamic Mechanism Design Problem

The optimal dynamic mechanism design problem that serves to yield the upper bound for our setting, (2) is challenging and has resisted optimal solution as discussed in the literature review. Here we find it convenient to relax problem (2) with the goal of computing tractable upper bounds. We obtain our upper bound as follows: we consider a simpler, upper bounding, one-dimensional mechanism design problem where customers can only misrepresent their valuation. This mechanism design problem serves as a relaxation to the optimal mechanism design problem defining $J^*(x_0, T)$. We derive an upper bound on the optimal value of this simpler mechanism design problem using a Myersonian approach. Put very loosely, our upper bounding problem is stated in terms of a ‘virtual allocation’ rule \bar{a}_ϕ that is a function of the allocation rule a_ϕ and the disutility of customer ϕ . Our task will then be one of finding a dynamic ‘virtual allocation’ policy that maximizes the expected

sum of virtual values that are ‘virtually’ allocated, subject to an inventory constraint that must be met in expectation. Put more precisely, let us define the ‘virtual allocation’ rule

$$\bar{a}_\phi \triangleq \mathbb{E}_{-\phi} [a_\phi - m(\phi, y_\phi)]$$

where we recall that $m(\phi, y) \triangleq \frac{\partial}{\partial v_\phi} M(\phi, y)$. Consider the following problem, whose optimal value we denote by $\bar{J}^*(x_0, T)$:

$$\begin{aligned} \max_{y^T \in \mathcal{Y}} \quad & \mathbb{E} \left[\sum_{\phi \in h^T} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \right) \bar{a}_\phi \right] \\ \text{s.t.} \quad & \mathbb{E} \left[\sum_{\phi \in h^T} \bar{a}_\phi \right] \leq x_0 \\ & \bar{a}_\phi \in [0, 1], \quad \forall \phi \text{ with } v_\phi > 0 \end{aligned} \tag{3}$$

Our main result in this section is that the optimal value of this program, $\bar{J}^*(x_0, T)$ is an upper bound on $J^*(x_0, T)$. The value of this result lies in the structure of the program (3) which is substantially more tractable than the program defining the optimal dynamic mechanism. Specifically by appropriately ‘dualizing’ the inventory constraint in this program in the next Section, we will be able to directly compare $\bar{J}^*(x_0, T)$ with the revenue under the static price policy.

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, let

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

and consider the following weaker 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. We now derive an upper bound on the expected price paid by customer ϕ for any feasible mechanism that satisfies (IR) and (IC'). The result lies on an appropriately general envelope theorem and uses our assumption on the concavity of $M(\cdot, \cdot)$ in v_ϕ .

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

$$\mathbb{E}_{-\phi}[p_\phi] \leq v_\phi \bar{a}_\phi - \int_{v'=0}^{v_\phi} \bar{a}_{\phi_{v'}} dv'.$$

Proof. First, we show that (IR) implies that

$$\mathbb{E}_{-\phi}[U(\phi_0, y_{\phi_0})] = 0. \quad (4)$$

To see this notice that by definition and Assumption 2, Part (1),

$$\mathbb{E}_{-\phi}[U(\phi_0, y_{\phi_0})] = 0 \cdot \mathbb{E}_{-\phi}[a_{\phi_0}] - \mathbb{E}_{-\phi}[p_{\phi_0}] - \mathbb{E}_{-\phi}[M(\phi_0, y_{\phi_0})] \leq 0.$$

But since (IR) requires $\mathbb{E}_{-\phi}[U(\phi_0, y_{\phi_0})] \geq 0$, we must have (4).

Now, define $u(\phi, y) \triangleq \frac{\partial}{\partial v_\phi} U(\phi, y)$. Applying the envelope theorem, we have:

$$\begin{aligned} \mathbb{E}_{-\phi}[U(\phi, y_\phi)] &= \int_{v'=0}^{v_\phi} \mathbb{E}_{-\phi}[u(\phi_{v'}, y_{\phi_{v'}})] dv' + \mathbb{E}_{-\phi}[U(\phi_0, y_{\phi_0})] \\ &= \int_{v'=0}^{v_\phi} \mathbb{E}_{-\phi}[a_{\phi_{v'}} - m(\phi_{v'}, y_{\phi_{v'}})] dv' + \mathbb{E}_{-\phi}[U(\phi_0, y_{\phi_0})] \\ &= \int_{v'=0}^{v_\phi} \mathbb{E}_{-\phi}[a_{\phi_{v'}} - m(\phi_{v'}, y_{\phi_{v'}})] dv' \end{aligned} \quad (5)$$

The first equality follows from Fubini's theorem and the envelope theorem (specifically, Theorem 2 of Milgrom and Segal [2002]). The second equality follows the definition of $u(\cdot)$, and the final equality follows from (4). Consequently,

$$\begin{aligned} \mathbb{E}_{-\phi}[p_\phi] &= v_\phi \mathbb{E}_{-\phi}[a_\phi] - \mathbb{E}_{-\phi}[U(\phi, y_\phi)] - \mathbb{E}_{-\phi}[M(\phi, y_\phi)] \\ &= v_\phi \mathbb{E}_{-\phi}[a_\phi] - \int_{v'=0}^{v_\phi} \mathbb{E}_{-\phi}[a_{\phi_{v'}} - m(\phi_{v'}, y_{\phi_{v'}})] dv' - \mathbb{E}_{-\phi}[M(\phi, y_\phi)] \\ &\leq v_\phi \mathbb{E}_{-\phi}[a_\phi - m(\phi, y_\phi)] - \int_{v'=0}^{v_\phi} \mathbb{E}_{-\phi}[a_{\phi_{v'}} - m(\phi_{v'}, y_{\phi_{v'}})] dv' \end{aligned}$$

The first equality follows from the definition of $U(\cdot)$ and the fact that by the definition of feasible policies (i.e. since $y \in \mathcal{Y}$), $a_\phi p_\phi = p_\phi$. The second equality follows from our application of the envelope theorem above. Finally, by the assumed concavity of $M(\cdot)$ in v_ϕ , we have $M(\phi, y) \geq v_\phi m(\phi, y) + M(\phi_0, y)$, which, with the fact that $M(\phi_0, y) \geq 0$, yields first inequality. This completes the proof. ■

We next prove a corollary to this Lemma that allows us to replace the objective in the optimal

mechanism design problem, (2), with an analytically tractable quantity. Specifically, we have:

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

$$\mathbb{E} \left[\sum_{\phi \in h^T} p_\phi \right] \leq \mathbb{E} \left[\sum_{\phi \in h^T} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \right) \bar{a}_\phi \right].$$

Proof. We observe that Lemma 3 implies:

$$\begin{aligned} \mathbb{E} \left[\sum_{\phi \in h^T} p_\phi \right] &= \mathbb{E} \left[\sum_{\phi \in h^T} \mathbb{E}_{-\phi} [p_\phi] \right] \\ &\leq \mathbb{E} \left[\sum_{\phi \in h^T} v_\phi \bar{a}_\phi - \int_{v'=0}^{v_\phi} \bar{a}_{\phi_{v'}} dv' \right]. \end{aligned}$$

We now prove that the right hand side is the required quantity by changing the order of integration:

$$\begin{aligned} \mathbb{E} \left[\sum_{\phi \in h^T} v_\phi \bar{a}_\phi - \int_{v'=0}^{v_\phi} \bar{a}_{\phi_{v'}} dv' \right] &= \mathbb{E} \left[\sum_{\phi \in h^T} \mathbb{E}_{v_\phi} \left[v_\phi \bar{a}_\phi - \int_{v'=0}^{v_\phi} \bar{a}_{\phi_{v'}} dv' \right] \right] \\ &= \mathbb{E} \left[\sum_{\phi \in h^T} \int_{v_\phi=0}^{\infty} \left(v_\phi \bar{a}_\phi - \int_{v'=0}^{v_\phi} \bar{a}_{\phi_{v'}} dv' \right) f(v_\phi) dv_\phi \right] \\ &= \mathbb{E} \left[\sum_{\phi \in h^T} \int_{v_\phi=0}^{\infty} v_\phi \bar{a}_\phi f(v_\phi) dv_\phi - \int_{v'=0}^{\infty} \bar{a}_{\phi_{v'}} \int_{v_\phi=0}^{v'} f(v_\phi) dv_\phi dv' \right] \\ &= \mathbb{E} \left[\sum_{\phi \in h^T} \int_{v_\phi=0}^{\infty} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \right) \bar{a}_\phi f(v_\phi) dv_\phi \right] \\ &= \mathbb{E} \left[\sum_{\phi \in h^T} \mathbb{E}_{v_\phi} \left[\left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \right) \bar{a}_\phi \right] \right] \\ &= \mathbb{E} \left[\sum_{\phi \in h^T} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \right) \bar{a}_\phi \right]. \end{aligned}$$

Here the second equality follows from the fact that v_ϕ is independent of θ_ϕ and t_ϕ , the third equality follows from an exchange in the order of integration, and the fifth equality again employs the fact that v_ϕ is independent of θ_ϕ and t_ϕ . This completes the proof of the lemma. \blacksquare

The next lemma establishes a second implication of the constraints (IC') and (IR).

Lemma 5. *If (IC') and (IR) hold, then for any ϕ with $v_\phi > 0$, we have:*

$$\bar{a}_\phi \in [0, 1].$$

Proof. Consider any $\phi \in \Phi$ and any $v, v' \in \mathbb{R}_+$. (IC') implies $\mathbb{E}_{-\phi}[U(\phi_v, y_{\phi_v})] \geq \mathbb{E}_{-\phi}[U(\phi_v, y_{\phi_{v'}})]$ and $\mathbb{E}_{-\phi}[U(\phi_{v'}, y_{\phi_{v'}})] \geq \mathbb{E}_{-\phi}[U(\phi_{v'}, y_{\phi_v})]$. Adding these two inequalities, and writing them explicitly (using the definition of $U(\cdot)$), yields:

$$\begin{aligned} (v - v') \left(\mathbb{E}_{-\phi}[a_{\phi_v}] - \mathbb{E}_{-\phi}[a_{\phi_{v'}}] \right) &\geq \left(\mathbb{E}_{-\phi}[M(\phi_{v'}, y_{\phi_{v'}})] - \mathbb{E}_{-\phi}[M(\phi_v, y_{\phi_{v'}})] \right) \\ &\quad + \left(\mathbb{E}_{-\phi}[M(\phi_v, y_{\phi_v})] - \mathbb{E}_{-\phi}[M(\phi_{v'}, y_{\phi_v})] \right) \end{aligned}$$

Now the concavity of $M(\cdot)$ in v from Assumption 2 yields

$$\mathbb{E}_{-\phi}[M(\phi_{v'}, y_{\phi_{v'}})] - \mathbb{E}_{-\phi}[M(\phi_v, y_{\phi_{v'}})] \geq \mathbb{E}_{-\phi}[m(\phi_{v'}, y_{\phi_{v'}})](v' - v)$$

and

$$\mathbb{E}_{-\phi}[M(\phi_v, y_{\phi_v})] - \mathbb{E}_{-\phi}[M(\phi_{v'}, y_{\phi_v})] \geq \mathbb{E}_{-\phi}[m(\phi_v, y_{\phi_v})](v - v')$$

which upon substitution in the previous inequality yields:

$$(v - v') \left(\mathbb{E}_{-\phi}[a_{\phi_v} - m(\phi_v, y_{\phi_v})] - \mathbb{E}_{-\phi}[a_{\phi_{v'}} - m(\phi_{v'}, y_{\phi_{v'}})] \right) \geq 0$$

so that we may immediately conclude that $\mathbb{E}_{-\phi}[a_{\phi_v} - m(\phi_v, y_{\phi_v})]$ is non-decreasing in v . But (5) and (IR) imply that for any $v_\phi \geq 0$,

$$\mathbb{E}_{-\phi}[U(\phi, y_\phi)] = \int_{v'=0}^{v_\phi} \mathbb{E}_{-\phi}[a_{\phi_{v'}} - m(\phi_{v'}, y_{\phi_{v'}})] \geq 0.$$

which with the fact that $\mathbb{E}_{-\phi}[a_{\phi_v} - m(\phi_v, y_{\phi_v})]$ is non-decreasing in v immediately lets us conclude that

$$\mathbb{E}_{-\phi}[a_{\phi_v} - m(\phi_v, y_{\phi_v})] \geq 0$$

for all $v > 0$. On the other hand, the fact that $a_\phi \in \{0, 1\}$ and $m(\phi, y_\phi) \geq 0$ implies

$$\mathbb{E}_{-\phi}[a_\phi - m(\phi, y_\phi)] \leq 1$$

for all ϕ . Together, these two inequalities establish the claim. ■

The next lemma establishes a final consequence of the fact that we cannot allocate more inventory than available:

Lemma 6. *For any feasible policy $y^T \in \mathcal{Y}$, we have:*

$$\mathbb{E} \left[\sum_{\phi \in h^T} a_\phi - m(\phi, y_\phi) \right] = \mathbb{E} \left[\sum_{\phi \in h^T} \bar{a}_\phi \right] \leq x_0.$$

Proof. Since for any feasible policy we have $\sum_{\phi \in h^T} a_\phi \leq x_0$, and since $m(\phi, y_\phi) \geq 0$ under Assumption 2 Part 3, the claim is immediate. \blacksquare

We are now ready to revisit our relaxation to the optimal dynamic mechanism design problem (2). Specifically, recall the relaxed problem, (3), that we presented at the outset of this section, whose optimal value we denote by $\bar{J}^*(x_0, T)$:

$$\begin{aligned} \max_{y^T \in \mathcal{Y}} \quad & \mathbb{E} \left[\sum_{\phi \in h^T} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \right) \bar{a}_\phi \right] \\ \text{s.t.} \quad & \mathbb{E} \left[\sum_{\phi \in h^T} \bar{a}_\phi \right] \leq x_0 \\ & \bar{a}_\phi \in [0, 1], \quad \forall \phi \text{ with } v_\phi > 0 \end{aligned}$$

Now, Lemmas 4, 5, and 6, immediately yield the main result for this section:

Proposition 1. *The optimal value of the problem (3) is an upper bound to that of the optimal mechanism design problem, (2):*

$$\bar{J}^*(x_0, T) \geq J^*(x_0, T).$$

As it turns out this relaxed problem will permit an exact analysis that we delve into in the next section. That analysis will in turn permit a comparison with the expected revenues under the static price policy.

5. Static Prices: Asymptotic Optimality and Tight Uniform Guarantee

As discussed in Section 1, Gallego and Van Ryzin [1994], in a seminal piece of research demonstrated the efficacy of static prices. That research has formed the basis of a large body of RM research since. Here we ask whether those conclusions remain valid when customers are forward looking by completing the proofs of our main results, Theorems 1 and 2. Recall the the static price policy is

defined according to:

$$\pi_t^{\text{FP}} = \bar{F}^{-1} \left(\min \left\{ \frac{x_0}{\lambda T}, \bar{F}(v^*) \right\} \right) \triangleq \pi^{\text{FP}}$$

for all t such that $X_{t-} > 0$. Further, recall that since it is dominant strategy for customers to make myopic purchasing decisions when prices stay constant (Lemma 1), the sales process is, in fact Poisson with intensity $\lambda \bar{F}(\pi^{\text{FP}})$, so that the revenue under such a static pricing policy is

$$J_{\pi^{\text{FP}}, \tau^{\text{FP}}}(x_0, T) = \pi^{\text{FP}} \mathbb{E} \left[\min \left(N_{\pi^{\text{FP}}}^{\lambda T}, x_0 \right) \right]$$

Our goal in this Section will be to complete the proofs for our main theoretical guarantees, Theorems 1 and 2. We will proceed by comparing $J_{\pi^{\text{FP}}, \tau^{\text{FP}}}(x_0, T)$ with $\bar{J}^*(x_0, T)$, the upper bound derived in Proposition 1.

5.1. The Asymptotic Optimality of Static Prices: Proof of Theorem 1

In this section we complete the proof of Theorem 1. Recall that in the fluid regime, we consider a sequence of problems parameterized by n . In the n th problem, we have $x_0^{(n)} = nx_0$, and $\lambda^{(n)} = n\lambda$ and denote by a superscript (n) quantities relevant to the n th model in this scaling. All other aspects of the model – namely the customer utility model, and the horizon T , stay unchanged. Our goal is to establish that

$$\frac{J_{\pi^{\text{FP}}, \tau^{\text{FP}}}^{(n)}(nx_0, T)}{J^{*,(n)}(nx_0, T)} = 1 - O\left(\frac{1}{\sqrt{n}}\right).$$

Our key task in this section is to establish that the optimal value of the relaxed mechanism design problem (3), $\bar{J}^*(x_0, T)$ is upper bounded by the optimal value to the fluid program, $\lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}})$. Then using the fact (see, for instance Gallego and Van Ryzin [1994]) that

$$\frac{\pi^{\text{FP}} \mathbb{E} \left[\min \left(N_{\pi^{\text{FP}}}^{\lambda T}, nx_0 \right) \right]}{n \lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}})} = 1 - O\left(\frac{1}{\sqrt{n}}\right)$$

we will have completed a proof of Theorem 1.

Lemma 7. *The optimal value attained in the fluid customer optimization problem (1) is an upper bound to the optimal value of the relaxed mechanism design problem (3):*

$$\lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}}) \geq \bar{J}^*(x_0, T).$$

Proof. Consider the following Lagrangian relaxation of the relaxed mechanism design problem (3):

$$\begin{aligned} \max_{y^T \in \mathcal{Y}} \quad & \mathbb{E} \left[\sum_{\phi \in h^T} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} - \eta \right) \bar{a}_\phi \right] + \eta x_0 \\ \text{s.t.} \quad & \bar{a}_\phi \in [0, 1], \quad \forall \phi \text{ with } v_\phi > 0 \end{aligned} \quad (6)$$

and denote by $\bar{J}^{*,\eta}(x_0, T)$ its optimal value. Now for any feasible mechanism $y^T \in \mathcal{Y}$ in the relaxed mechanism design problem (3), and any $\eta \geq 0$, we have

$$\eta \left(x_0 - \mathbb{E} \left[\sum_{\phi \in h^T} \bar{a}_\phi \right] \right) \geq 0.$$

It follows that $\min_{\eta \geq 0} \bar{J}^{*,\eta}(x_0, T) \geq \bar{J}^*(x_0, T)$; a statement of weak duality. Now, for any feasible y^T to the program above, (6), we require $\bar{a}_\phi \in [0, 1]$ if $v_\phi > 0$. So for all such ϕ , such that $v_\phi > 0$, we have

$$\left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} - \eta \right) \bar{a}_\phi \leq \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} - \eta \right) \mathbf{1} \left\{ v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \geq \eta \right\}$$

Moreover, the set on which $v_\phi = 0$ is of measure zero by assumption, so that for any $\eta \geq 0$, we immediately have:

$$\begin{aligned} \bar{J}^{*,\eta}(x_0, T) &\leq \mathbb{E} \left[\sum_{\phi \in h^T} \left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} - \eta \right) \mathbf{1} \left\{ v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \geq \eta \right\} \right] + \eta x_0 \\ &= \lambda T \mathbb{E} \left[\left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} - \eta \right) \mathbf{1} \left\{ v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \geq \eta \right\} \right] + \eta x_0 \end{aligned}$$

where the first equality is Wald's identity. Now, let $\hat{\eta} = \pi^{\text{FP}} - \frac{\bar{F}(\pi^{\text{FP}})}{f(\pi^{\text{FP}})}$ in the above inequality and observe that

$$\begin{aligned} \bar{J}^{*,\hat{\eta}}(x_0, T) &\leq \lambda T \mathbb{E} \left[\left(v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} - \hat{\eta} \right) \mathbf{1} \left\{ v_\phi - \frac{\bar{F}(v_\phi)}{f(v_\phi)} \geq \hat{\eta} \right\} \right] + \hat{\eta} x_0 \\ &= \lambda T \int_0^\infty \left(v - \frac{\bar{F}(v)}{f(v)} - \hat{\eta} \right)^+ f(v) dv + \hat{\eta} x_0 \\ &= \lambda T \int_{v=\pi^{\text{FP}}}^\infty \left(v - \frac{\bar{F}(v)}{f(v)} - \hat{\eta} \right) f(v) dv + \hat{\eta} x_0 \\ &= \lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}}) + \hat{\eta} (x_0 - \lambda T \bar{F}(\pi^{\text{FP}})) \\ &= \lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}}). \end{aligned}$$

The second equality above uses the fact that $v - \frac{\bar{F}(v)}{f(v)}$ is non-decreasing by Assumption 1. The third equality uses the fact that $\int_{v=p}^{\infty} v f(v) - \bar{F}(v) dv = p \bar{F}(p)$. For the final equality (essentially, a complementary slackness condition), notice that we must have by definition of π^{FP} that $x_0 - \lambda T \bar{F}(\pi^{\text{FP}}) \geq 0$ and that when $x_0 - \lambda T \bar{F}(\pi^{\text{FP}}) > 0$ we have $\pi^{\text{FP}} = v^*$, so that $\hat{\eta} = 0$. We have thus shown that

$$\lambda T \bar{F}(\pi^{\text{FP}}) \geq \bar{J}^{*,\hat{\eta}}(x_0, T) \geq \min_{\eta \geq 0} \bar{J}^{*,\eta}(x_0, T) \geq \bar{J}^*(x_0, T)$$

which is the result. ■

We are now in a position to establish the asymptotic optimality of the static price policy. Specifically, Theorem 3 of Gallego and Van Ryzin [1994] establishes

$$\frac{\pi^{\text{FP}} \mathbb{E} \left[\min \left(N_{\pi^{\text{FP}}}^{\lambda T}, x_0 \right) \right]}{\lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}})} \geq 1 - \frac{1}{\sqrt{\min \left(x_0, \lambda T \bar{F}(v^*) \right)}},$$

and recall that $J_{\pi^{\text{FP}}, \tau \pi^{\text{FP}}}(x_0, T) = \pi^{\text{FP}} \mathbb{E} \left[\min \left(N_{\pi^{\text{FP}}}^{\lambda T}, x_0 \right) \right]$. Now, in the n th problem, we consider an arrival rate of $n\lambda$, and initial inventory of nx_0 , so that

$$\frac{J_{\pi^{\text{FP}}, \tau \pi^{\text{FP}}}^{(n)}(nx_0, T)}{n \lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}})} \geq 1 - \frac{1}{\sqrt{n \min \left(x_0, \lambda T \bar{F}(v^*) \right)}}.$$

We establish in Lemma 7 that $n \lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}}) \geq \bar{J}^{*,(n)}(nx_0, T)$, and Proposition 1 shows that $\bar{J}^{*,(n)}(nx_0, T) \geq J^{*,(n)}(nx_0, T)$, so that we have shown

$$\frac{J_{\pi^{\text{FP}}, \tau \pi^{\text{FP}}}^{(n)}(nx_0, T)}{J^{*,(n)}(nx_0, T)} \geq 1 - \frac{1}{\sqrt{n \min \left(x_0, \lambda T \bar{F}(v^*) \right)}}.$$

This completes the proof of Theorem 1.

5.2. A (Tight) Constant Factor Guarantee for Static Prices

While the previous section established a performance guarantee in the fluid regime, it is also possible to obtain a constant factor guarantee that is valid uniformly in all model parameters. We do this by directly showing that

$$\frac{\pi^{\text{FP}} \mathbb{E} \left[\min \left(N_{\pi^{\text{FP}}}^{\lambda T}, x_0 \right) \right]}{\lambda T \bar{F}(\pi^{\text{FP}})} \geq 1 - \frac{1}{e}.$$

Interestingly, this intermediate result also yields a uniform performance guarantee on the static price policy in the setting of myopic customers, i.e. the setting studied by Gallego and Van Ryzin [1994]. Given the long history of the problem, and the lack of any such constant factor guarantee in antecedent literature, this intermediate result is of independent interest.

Lemma 8.

$$\frac{\pi^{\text{FP}} \mathbb{E} \left[\min \left(N_{\pi^{\text{FP}}}^{\lambda T}, x_0 \right) \right]}{\lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}})} \geq 1 - \frac{1}{e}.$$

Proof. For brevity of notation, denote $\bar{\lambda} \triangleq \lambda T \bar{F}(\pi^{\text{FP}})$ and $\bar{N} \triangleq N_{\pi^{\text{FP}}}^{\lambda T}$ throughout this proof. The definition of π^{FP} implies $\bar{\lambda} \leq x_0$. We have:

$$\begin{aligned} \frac{\mathbb{E} \left[\min \left(N_{\pi^{\text{FP}}}^{\lambda T}, x_0 \right) \right]}{\lambda T \bar{F}(\pi^{\text{FP}})} &\triangleq \frac{1}{\bar{\lambda}} \mathbb{E} \left[\min \left(\bar{N}, x_0 \right) \right] \\ &= \frac{1}{\bar{\lambda}} \left(\mathbb{E} \left[\bar{N} \right] - \mathbb{E} \left[\left(\bar{N} - x_0 \right)^+ \right] \right) \\ &= 1 - \mathbb{E} \left[\left(\bar{N} - x_0 \right)^+ \right] / \bar{\lambda} \\ &= 1 - \sum_{n=1}^{\infty} n \frac{e^{-\bar{\lambda}} \bar{\lambda}^{x_0+n-1}}{(x_0+n)!} \end{aligned} \tag{7}$$

Now $\bar{\lambda} \leq x_0$ by the definition of π^{FP} . Further, $e^{-\bar{\lambda}} \bar{\lambda}^{x_0+n-1}$ is non-decreasing in $\bar{\lambda}$ on $n \geq 1$ since

$$\frac{\partial}{\partial \bar{\lambda}} \ln \left(e^{-\bar{\lambda}} \bar{\lambda}^{x_0+n-1} \right) = -1 + \frac{x_0 + n - 1}{\bar{\lambda}} \geq 0$$

provided $n \geq 1$. Thus, (7) yields

$$\begin{aligned} \frac{\mathbb{E} \left[\min \left(N_{\pi^{\text{FP}}}^{\lambda T}, x_0 \right) \right]}{\lambda T \bar{F}(\pi^{\text{FP}})} &\geq 1 - \sum_{n=1}^{\infty} n \frac{e^{-x_0} x_0^{x_0+n-1}}{(x_0+n)!} \\ &= 1 - \sum_{n=1}^{\infty} ((x_0+n) - x_0) \frac{e^{-x_0} x_0^{x_0+n-1}}{(x_0+n)!} \\ &= 1 - \sum_{n'=0}^{\infty} \frac{e^{-x_0} x_0^{x_0+n'}}{(x_0+n')!} + \sum_{n=1}^{\infty} \frac{e^{-x_0} x_0^{x_0+n}}{(x_0+n)!} \\ &= 1 - \frac{e^{-x_0} x_0^{x_0}}{x_0!} \\ &\geq 1 - \frac{1}{e} \end{aligned}$$

where the last inequality follows from that fact that $\frac{e^{-x_0} x_0^{x_0}}{x_0!}$ is non-increasing in x_0 on $x_0 \geq 1$. This

completes the proof. ■

We established in Lemma 7 that $\lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}}) \geq \bar{J}^*(x_0, T)$, and Proposition 1 showed that $\bar{J}^*(x_0, T) \geq J^*(x_0, T)$, so that with Lemma 8, we have:

$$\frac{J_{\pi^{\text{FP}}, \tau \pi^{\text{FP}}}(x_0, T)}{J^*(x_0, T)} \geq 1 - \frac{1}{e}.$$

This completes the proof of Theorem 2.

Interestingly, this constant factor guarantee also implies a stronger guarantee for the class of ‘Robust Pricing Policies’ proposed by Chen and Farias [2015]. In that paper the authors establish that so-called robust pricing policies provided at least 29% of the revenue under an optimal mechanism. The class of utility models considered in the present paper subsumes the class of utility models studied in Chen and Farias [2015], and the static price policy is trivially a robust pricing policy thereby improving the Chen and Farias [2015] guarantee from 29% to $\sim 63.2\%$.

Let us summarize what we have established in this section. We set out to compare the performance of static price policies against a family of selling mechanisms that subsumed dynamic pricing. Specifically, our benchmark, which is the optimal dynamic mechanism for the problem at hand, includes virtually any selling format one may imagine. Antecedent research suggests that this dynamic mechanism design problem is essentially intractable [Pai and Vohra, 2013]. The principle insight in this paper is the surprising fact that a mechanism as simple as a static posted price is, to a first order, optimal. We made this notion precise by showing that the expected revenue under an optimally set static price is optimal in the regime where inventory and demand grow large. We complemented this result with a uniform performance guarantee that is valid in any regime. To further round out these results, we explore the tightness of our analysis and the necessity of our modeling assumptions in the next section.

6. Tight Problem Instances and Sub-Optimality for General Disutilities

This section seeks to answer two key questions, both of which are related to the tightness of our analysis and the applicability of our results. Specifically, we ask:

1. Is our uniform performance guarantee (Theorem 2) tight? Foreshadowing the answer to this question (we show the guarantee is in fact tight) we will also explore computationally how ‘quickly’ one gets from the level of performance loss in that guarantee to the optimality suggested by Theorem 1.

2. Were our assumptions on the customer disutility function $M(\cdot, \cdot)$ necessary? Here we give an example of a disutility function for which our guarantees do not hold. The example will serve to illustrate what can go wrong if customer disutility grows sufficiently ‘rapidly’ with valuation.

6.1. A Tight Problem Instance

Theorem 2 shows that the expected revenue under the static price policy π^{FP} is at least within a factor of $1 - 1/e$ of that under an optimal dynamic mechanism. This analysis is potentially loose for a number of reasons, the most important one perhaps being that we compared ourselves against an upper bound derived via a relaxation to the optimal dynamic mechanism design problem. Surprisingly, the guarantee is in fact tight, as we now illustrate.

Example 1. (*Tight Problem Instances*) As an example of a tight problem instance, we consider a problem with the following desiderata. First, there is a single unit of inventory; $x_0 = 1$. Second, customer values are uniformly distributed on the unit interval so that $F(v) = v$ for $v \in [0, 1]$. Finally, all customers are ‘fully patient’ so that $M(\phi, y) = 0$ for all (ϕ, y) with $\tau_y \geq t_\phi$. We continue to parameterize by λ the rate of customer arrival.

As we will discuss momentarily, the optimal dynamic mechanism for the problem instance above is simply conducting a Myerson auction. Using this fact, we can establish that the performance guarantee in Theorem 2 is tight for the family of examples above.

Proposition 2. *For the family of problems defined in Example 1, we have:*

$$\limsup_{\lambda \rightarrow \infty} \frac{J_{\pi^{\text{FP}}, \tau^{\pi^{\text{FP}}}}(x_0, T)}{J^*(x_0, T)} \leq 1 - \frac{1}{e}$$

Proof. Let N be a Poisson random variable of rate λT . We first show that:

$$J^*(x_0, T) \geq \left(1 - \frac{4}{\lambda T}\right) \mathbb{P}\left(N > \frac{\lambda T}{2}\right)$$

Observe that for the example at hand, an optimal mechanism is simply for the seller to wait until time T and proceed to conduct a (static) revenue maximizing auction at this time for all arrivals prior to that time; for concreteness let us assume the auction conducted is the Myerson auction (second price with reserve). This is clearly a feasible mechanism. To see why this is optimal, we note that in this setting, any $y^T \in \mathcal{Y}$ can be interpreted as a randomized allocation and payment

rule for a static revenue maximizing auction with N bidders and a single product. We conclude

$$J^*(x_0, T) = \mathbb{E} \left[1 - \frac{2(1 - 2^{-(N+1)})}{N+1} \right].$$

But,

$$\begin{aligned} \mathbb{E} \left[1 - \frac{2(1 - 2^{-(N+1)})}{N+1} \right] &\geq \mathbb{E} \left[1 - \frac{2}{N} \right] \\ &\geq \mathbb{E} \left[1 - \frac{2}{N} \middle| N > \frac{\lambda T}{2} \right] \mathbb{P} \left(N > \frac{\lambda T}{2} \right) \\ &\geq \left(1 - \frac{4}{\lambda T} \right) \mathbb{P} \left(N > \frac{\lambda T}{2} \right). \end{aligned} \tag{8}$$

We next establish an upper bound on the performance of the static price policy. Observe that by definition of the static price policy we have that for $\lambda > 2/T$, $\bar{F}(\pi^{\text{FP}}) = 1/\lambda T$. Consequently, for $\lambda > 2/T$, we have:

$$\begin{aligned} J_{\pi^{\text{FP}}, \tau^{\pi^{\text{FP}}}}(x_0, T) &= \pi^{\text{FP}} \mathbb{E} \left[1 - \left(1 - \bar{F}(\pi^{\text{FP}}) \right)^N \right] \\ &\leq \mathbb{E} \left[1 - \left(1 - \bar{F}(\pi^{\text{FP}}) \right)^N \right] \\ &= \mathbb{E} \left[1 - \left(1 - \frac{1}{\lambda T} \right)^N \right] \\ &\leq 1 - \left(1 - \frac{1}{\lambda T} \right)^{\lambda T}, \end{aligned} \tag{9}$$

where the first inequality follows from the property that $\pi^{\text{FP}} \leq 1$, the second inequality follows from the property that the function a^N is convex for any $a > 0$, Jensen's inequality, and the property that $\mathbb{E}[N] = \lambda T$. The result now follows from (8) and (9), since $\lim_{\lambda \rightarrow \infty} \left(1 - \frac{1}{\lambda T} \right)^{\lambda T} = \frac{1}{e}$ and $\lim_{\lambda \rightarrow \infty} \mathbb{P} \left(N > \frac{\lambda T}{2} \right) = 1$. \blacksquare

The result above shows an example where the gap between the static-price revenue and that under an optimal dynamic mechanism is indeed approximate 37% so that the bound in Theorem 2 is tight. In contrast, Theorem 1 suggests that the static price is optimal in the fluid regime. As such, one is led to wonder whether the performance loss exhibited in the above example quickly mitigates as we change problem parameters, allowing, say inventory to grow large. With that in mind, consider the following numerical experiment: we assume customer valuations are exponentially distributed with unit rate. Further we assume $\lambda = 1$ and $T = 10$. We then numerically compare

the performance of the fixed price policy to an upper bound on the value of an optimal dynamic mechanism, reporting the performance metric:

$$\text{LB}^{\text{FP}}(x_0, T) \triangleq \frac{J_{\pi^{\text{FP}}, \tau^{\pi^{\text{FP}}}}(x_0, T)}{\lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}})}.$$

(Recall that Lemma 7 and Proposition 1 together established that $J^*(x_0, T) \leq \lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}})$).

Table 1: Performance of the Static Price Policy π^{FP} .

x_0	1	2	4	8
$\text{LB}^{\text{FP}}(x_0, T)$	0.63	0.72	0.84	1.00

Notice that for an inventory level of one unit, the performance loss implied by the table above is again $\sim 37\%$. However, this quickly declines with further units of inventory. We see that even in a decidedly non-asymptotic setting, the static price policy already leaves little room for improvement. In fact, this is the *core reason*, the very intuitive results of Gallego and Van Ryzin [1994] have proved so influential in revenue management.

We next turn our attention to the modeling assumptions we made on customer disutility in Section 2.1 with the goal of developing a better understanding of the aspects of our assumptions that appear crucial to our guarantees.

6.2. Sub-Optimality In The Face of General Disutilties

Recall from our discussion on the modeling of customer disutility (i.e., the function $M(\cdot, \cdot)$) in Section 2.1 that we required that a customer’s disutility be concave and non-decreasing in customer valuation (part 4 of Assumption 2). As we discussed there, the key restriction in that assumption (as implied by the condition of concavity) is in requiring that customer disutility not increase ‘too fast’ with valuation. While in Section 2.1 we provided a number of examples of disutility functions in the literature (both theoretical and empirical) that satisfy our assumptions, we seek to go in the opposite direction in this section. We ask what happens if customer disutility did in fact increase super-linearly in valuation. To that end, consider the following class of *deadline-based* disutilities, given by:

$$M(\phi, y) = v_\phi \mathbf{1}_{\{\tau_\phi - t_\phi > d(v_\phi)\}},$$

where $d(\cdot)$ maps valuations to a ‘deadline’. It is easy to see that for a suitable choice of the function $d(\cdot)$, this specification leads to a discontinuity in the dependence of disutility on valuation wherein keeping y and the other components of ϕ fixed, $M(\phi, y)$ jumps from 0 to v , as v is increased beyond

a threshold. Putting aside the relative merits and de-merits of this specification for now, we first focus on showing that static prices are *sub-optimal* for a specification such as the one above⁷. To that end, consider a setting where v is uniformly distributed on the unit interval, and the ‘deadline function’ is,

$$d(v) = T\mathbf{1}_{\{v \leq 1/2\}},$$

so that customer with valuations less than one half are fully patient, while the remaining customers are fully myopic. Figure 1 plots the relationship between $M(\phi_v, y)$ and v for any fixed y with $\tau_y > t_{\phi_v}$. Under this model, the customer disutility function $M(\phi_v, y)$ is not concave in v , i.e.,




Figure 1: Relationship between $M(\phi_v, y)$ and v when $\tau_y > t_{\phi_v}$.

Assumption 2 is violated. Specifically, keeping y and all other components of ϕ fixed, the disutility jumps from 0 (for any value of $v < 1/2$) to $1/2$ at $v = 1/2$, and then increases linearly from there.

To further simplify our analysis of what could go wrong here, let us consider the situation where inventory is unlimited so that $x_0 = \infty$. Now, in this setting the static price policy would set $\pi^{\text{FP}} = 1/2$, which will garner expected revenue

$$\lambda T \pi^{\text{FP}} \bar{F}(\pi^{\text{FP}}) = \frac{1}{4} \lambda T.$$

Now consider the following alternative pricing policy that instantaneously drops prices at the very end of the horizon:

$$\hat{\pi}_t = \begin{cases} 1/2 & \text{if } t < T \\ 1/4 & \text{if } t = T \end{cases}.$$

⁷Of course, we require that the deadline function $d(\cdot)$ has a non-trivial dependence on v . Else, the requirements of Assumption 2 are trivially satisfied

It is simple to verify that under policy $\hat{\pi}$, a candidate equilibrium stopping rule is

$$\tau_{\phi}^{\hat{\pi}} = \begin{cases} t_{\phi} & \text{if } v_{\phi} < 1/4 \text{ or } v_{\phi} \geq 1/2 \\ T & \text{if } v_{\phi} \in [1/4, 1/2) \end{cases}$$

That is, customers with utility greater than $1/2$ will purchase immediately, customers with utility between $1/4$ and $1/2$ will wait until the end of the horizon and then purchase, while customers with utility less than $1/4$ will leave immediately upon arrival without a purchase. This yields expected revenue:

$$\frac{1}{2} \cdot \lambda T \frac{1}{2} + \frac{1}{4} \cdot \lambda T \frac{1}{4} = \frac{5}{16} \lambda T.$$

which improves on the static price revenue by a factor of 25%. Since this relative improvement is independent of the value of λ , and since x_0 was chosen to be unbounded, we see that one cannot hope for asymptotic optimality in this setting. We have thus identified a class of disutility functions that do not satisfy our assumptions, and for which static prices are *not* asymptotically optimal. It makes sense to pursue the design of more sophisticated pricing policies or more general mechanisms in such a setting.

The deadline based disutility functions above have some interesting features – for instance, they allow us to ‘ignore’ customers who arrived prior to a finite time in the past and thereby allow for a succinct representation of state in dynamic programming analyses. On the other hand, such a specification implies that a customer prefers lotteries (as opposed to a deterministic outcome) with respect to the time of assignment τ_{ϕ} which is potentially unrealistic (Azevedo and Gottlieb [2012]). As discussed earlier, the deadline based disutility function has a ‘jump’ when viewed as a function of valuation, keeping other quantities fixed. The crux of Assumption 2 was the requirement that this same function be concave thereby placing a restriction on the rate at which disutility may grow with valuation. Succinctly, provided disutility grows sufficiently slowly (essentially, sub-linearly) with valuation, static prices suffice. If on the other hand, we permitted disutility to grow rapidly with valuation (exemplified by the jump in the deadline-based disutility function), more sophisticated policies are called for.

7. Concluding Remarks

This paper has focused on a canonical problem of revenue management and shown that static prices are, to a first order, optimal for a broad class of customer utility functions. The economic message here is simple and clear and reinforces the message that static pricing policies – or ‘everyday low

prices’ in the vernacular of the dynamic pricing literature – can be surprisingly effective. This message was first delivered by Gallego and Van Ryzin [1994] at a time when search costs were in effect high (e-commerce and the widespread use of the internet did not exist at the time). As such at that time, it was fair to assume that customers were effectively myopic since strategizing on the timing of a purchase was hard. That assumption has become increasingly questionable in the last decade, and with it the key message on the efficacy of static prices. The present paper resolves that conundrum for what we believe is a broad class of utility models that find a broad base of support in multiple streams of literature. There are several remarks that are worth making in concluding:

Non-Stationary Arrival Rates: It is possible to generalize our analysis to the setting where customers arrive according to a Poisson process with a *non*-homogenous rate λ_t . To begin, there is no change required in the proof of Proposition 1. The rest of the changes essentially follow those made by Gallego and Van Ryzin [1994] in their own extension to non-stationary arrival rates. Specifically, replacing λ by $\frac{1}{T} \int_0^T \lambda_t dt$ in the definition of π^{FP} yields a static price policy that is optimal for (1) as shown by Gallego and Van Ryzin [1994]. That same substitution can be made with no further argument in the proof of Lemma 7, and finally, we rely on the extension to non-stationary arrival rates provided by Gallego and Van Ryzin [1994] to their main result (Theorem 3) to complete the proof of our Theorem 1 for non-stationary arrival rates. The proof of Lemma 8 (which immediately yields our Theorem 2) again follows from the observation made by Gallego and Van Ryzin [1994] (via a change of clock argument) that the expected revenue under a fixed price policy, for any price p , is given by $p\mathbb{E} \left[\min(\bar{N}, x_0) \right]$ where \bar{N} is a Poisson random variable with rate $\int_0^T \lambda_t \bar{F}(p) dt$.

The Power to Commit: Any analysis invoking the principles of mechanism design will typically call for an assumption that the seller has the ability to credibly ‘commit’ to a mechanism. Indeed, this is true of even the simplest mechanisms (such as the design of the optimal static auction). As already discussed, antecedent literature in revenue management (and more generally, in mechanism design) has provided a variety of arguments in support of the power to commit so that the assumption is broadly accepted. Nonetheless, it is worth noting that in the absence of the ability to verify (after the fact) that the seller has indeed stuck to her commitment, such an assumption is less palatable. In the case of dynamic mechanisms it is often the case that such a verification is difficult without the seller revealing a great deal of information (including at least the history of all customer allocations). Happily, in the case of a static price, the situation is a lot simpler. In particular, it is trivial for any customer to verify at the end of the selling season that the seller

has deviated from a static price mechanism simply by having observed the price trajectory over the season. Given that myopic behavior is dominant under a static price, the seller may as well employ the optimal static price, so there is no need for the buyer to verify whether the price used was indeed optimal.

Further Research Questions: We believe the present paper sets the stage for an exciting set of further research questions. For instance, how well can one approximate a general disutility function by one in the class we permit? Can we extend our analysis to sub-linear dis-utilities? Another direction is considering more general revenue management problems. Perhaps, most obviously, the Gallego and Van Ryzin [1994] analysis showing the asymptotic optimality of static prices led to follow-on analysis of a similar pricing policy in the ‘network’ revenue management setting. What can one say in such a setting when customers are strategic? Would a similar fixed price policy work well? It is clear that the more general setting creates a few new complications, namely, the ability to substitute across products and the fact that we have multiple resources to name two. Given the vast landscape of RM techniques motivated by the original ideas in Gallego and Van Ryzin [1994], we see this latter direction for future work as particularly exciting.

References

- Philipp Afèche and Haim Mendelson. Pricing and priority auctions in queueing systems with a generalized delay cost structure. *Management Science*, 50(7):869–882, 2004.
- Philipp Afèche and J Michael Pavlin. Optimal price/lead-time menus for queues with customer choice: Segmentation, pooling, and strategic delay. *Management Science*, 2016.
- Simon P Anderson and Regis Renault. Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model. *The RAND Journal of Economics*, pages 719–735, 1999.
- Baris Ata and Tava Lennon Olsen. Near-optimal dynamic lead-time quotation and scheduling under convex-concave customer delay costs. *Operations Research*, 57(3):753–768, 2009.
- Yossi Aviv and Amit Pazgal. Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing & Service Operations Management*, 10(3):339–359, 2008.
- Eduardo M Azevedo and Daniel Gottlieb. Risk-neutral firms can extract unbounded profits from consumers with prospect theory preferences. *Journal of Economic Theory*, 147(3):1291–1299, 2012.

- Omar Besbes and Ilan Lobel. Intertemporal price discrimination: Structure and computation of optimal policies. *Management Science*, 61(1):92–110, 2015.
- Simon Board and Andrzej Skrzypacz. Revenue management with forward looking buyers. *Forthcoming in Journal of Political Economy*, 2013.
- Christian Borgs, Ozan Candogan, Jennifer Chayes, Ilan Lobel, and Hamid Nazerzadeh. Optimal multiperiod pricing with service guarantees. *Management Science*, 60(7):1792–1811, 2014.
- René Caldentey and Gustavo Vulcano. Online auction and list price revenue management. *Management Science*, 53(5):795–813, 2007.
- René Caldentey, Ying Liu, and Ilan Lobel. Intertemporal pricing under minimax regret. *Available at SSRN 2357083*, 2015.
- Yiwei Chen and Vivek F Farias. Robust dynamic pricing with strategic customers. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, pages 777–777, 2015.
- Peter A Diamond. A model of price adjustment. *Journal of economic theory*, 3(2):156–168, 1971.
- Sherwin Doroudi, Mustafa Akan, Mor Harchol-Balter, Jeremy Karp, Christian Borgs, and Jennifer T Chayes. Priority pricing in queues with a continuous distribution of customer valuations. 2013.
- Glenn Ellison and Alexander Wolitzky. A search cost model of obfuscation. *The RAND Journal of Economics*, 43(3):417–441, 2012.
- Guillermo Gallego and Garrett Van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020, 1994.
- Jérémie Gallien. Dynamic mechanism design for online commerce. *Operations Research*, 54(2):291–310, 2006.
- Itai Gurvich, Martin Lariviere, and Can Ozkan. Coverage, coarseness and classification: Determinants of social efficiency in priority queues. 2016.
- Nima Haghpanah and Jason Hartline. Reverse mechanism design. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 757–758. ACM, 2015.
- Zhenyu Hu, Xin Chen, and Peng Hu. Dynamic pricing with gain-seeking reference price effects. *Operations Research*, 64(1):150–157, 2016.

- AK Katta and Jay Sethuraman. Pricing strategies and service differentiation in an M/M/1 queue - A profit maximization perspective. Technical report, Working Paper, 2005.
- Cinar Kilcioglu and Costis Maglaras. Revenue maximization for cloud computing services. *ACM SIGMETRICS Performance Evaluation Review*, 43(3):76–76, 2015.
- Jun Li, Nelson Granados, and Serguei Netessine. Are consumers strategic? Structural estimation from the air-travel industry. *Management Science*, 60(9):2114–2137, 2014.
- Qian Liu and Garrett J Van Ryzin. Strategic capacity rationing to induce early purchases. *Management Science*, 54(6):1115–1131, 2008.
- Yan Liu and William L Cooper. Optimal dynamic pricing with patient customers. *Operations Research*, 63(6):1307–1319, 2015.
- Ilan Lobel. Dynamic pricing with heterogeneous patience levels. Working Paper, New York University, 2017.
- Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- Ken Moon, Kostas Bimpikis, and Haim Mendelson. Randomized markdowns and online monitoring. *Available at SSRN 2589386*, 2015.
- Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- Hamid Nazerzadeh and Ramandeep S Randhawa. Near-optimality of coarse service grades for customer differentiation in queueing systems. *Available at SSRN 2438300*, 2015.
- Özalp Özer and Robert Phillips. *The Oxford handbook of pricing management*. OUP Oxford, 2012.
- Mallesh M Pai and Rakesh Vohra. Optimal dynamic auctions and simple index rules. *Mathematics of Operations Research*, 38(4):682–697, 2013.
- Rafael Rob. Equilibrium price distributions. *The Review of Economic Studies*, 52(3):487–504, 1985.
- Maher Said. Auctions with dynamic populations: Efficiency and revenue maximization. *Journal of Economic Theory*, 147(6):2419–2438, 2012.
- Nancy L Stokey. Intertemporal price discrimination. *The Quarterly Journal of Economics*, pages 355–371, 1979.

Kalyan T Talluri and Garrett J Van Ryzin. *The theory and practice of revenue management*, volume 68. Springer Science & Business Media, 2006.

Gustavo Vulcano, Garrett Van Ryzin, and Costis Maglaras. Optimal dynamic auctions for revenue management. *Management Science*, 48(11):1388–1407, 2002.

Zizhuo Wang. Intertemporal price discrimination via reference price effects. *Operations Research*, 64(2):290–296, 2016.

Rui Yin, Yossi Aviv, Amit Pazgal, and Christopher S Tang. Optimal markdown pricing: Implications of inventory display formats in the presence of strategic customers. *Management Science*, 55(8):1391–1408, 2009.