

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

 Yiwei Chen,^a Vivek F. Farias,^{b,c} Nikolaos Trichakis^{b,c}

^a Carl H. Lindner College of Business, University of Cincinnati, Cincinnati, Ohio 45221; ^b Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; ^c Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

Contact: chen3yw@ucmail.uc.edu,  <https://orcid.org/0000-0001-6243-868X> (YC); vivekf@mit.edu,

 <https://orcid.org/0000-0002-5856-9246> (VFF); ntrichakis@mit.edu,  <https://orcid.org/0000-0002-8324-9148> (NT)

Received: April 11, 2017

Revised: December 18, 2017; April 8, 2018

Accepted: May 24, 2018

Published Online in *Articles in Advance*:
August 1, 2019

<https://doi.org/10.1287/mnsc.2018.3203>

Copyright: © 2019 INFORMS

Abstract. We consider a canonical revenue management (RM) problem wherein a monopolist seller posts prices for multiple products that are for sale over a fixed horizon so as to maximize expected revenues. Products are differentiated and subject to joint capacity constraints. Arriving 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 has established that *static* prices are asymptotically optimal for this problem in the regime where inventory and demand grow large. In stark contrast, for the case where customers are forward looking, available results in mechanism design and dynamic pricing suggest substantially more complicated prescriptions. Notably, these results apply to settings with merely a single product type, and they are also often constrained by restrictive assumptions on customer type. We demonstrate that static prices surprisingly remain asymptotically optimal in the face of strategic customers for a multi-product setting and for a broad class of customer utility models. For the single-product case, we further show that an optimally set static price guarantees the seller revenues that are within at least 63.2% of that under an optimal dynamic pricing policy, irrespective of regime. For utility models outside the class we consider, we show that static prices need not be asymptotically optimal. Nevertheless, the class of customer utility models we consider is parsimonious, enjoys empirical support, and subsumes many of the models considered for this problem in existing mechanism design research. We allow for multidimensional customer types and for a customer's disutility from waiting to be positively correlated with his valuation. Therefore, our findings are likely to be robust and provide for a canonical RM problem a simple prescription that is near-optimal across a broad set of modeling assumptions.

History: Accepted by Yinyu Ye, optimization.

Keywords: dynamic pricing • mechanism design • approximation algorithms • revenue management • strategic customers • asymptotic analysis

1. Introduction

Consider the following canonical revenue management (RM) problem: A monopolist seller faces some fixed time horizon over which she can sell multiple differentiated products, subject to joint capacity constraints, via an anonymous, posted price mechanism. Customers considering to purchase one of the products arrive over time. Should a customer choose to make a purchase, he must pay a price equal to the associated price posted at the time of his purchase. The seller's goal is to maximize expected sales revenue.

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 and “leave the system.” Indeed, given the appropriate assumptions on the customer arrival process, this problem can be solved by using dynamic programming. In fact, the seller can get away with doing something even simpler: charge fixed

prices. In an early foundational paper, Gallego and Van Ryzin (1994) studied a single-product version of the problem and established that if the seller chose to maintain a single price 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. Asymptotic optimality of static policies was subsequently established for the multiproduct version and other closely related variants of the problem (Talluri and Van Ryzin 2006). The insights from this body of work are not merely theoretical; they have been borne out in a host of practical RM applications ranging from the problem discussed above to substantially more complicated problems.

Technological changes have brought into question the basic premise of myopic customers. 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 intertemporal price discrimination on the part of the seller, a feature essentially absent from the model with myopic customers.

Substantially less is known about the setting where customers are forward looking. If one were to take the path of assuming that customers were strategic—as opposed to some boundedly rational model of behavior—existing research on this problem divides broadly into two threads:

- The first 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 multidimensional types. Approximation algorithms for the problem have also recently been proposed, but the mechanisms remain sophisticated with restrictive assumptions on customer utility models.

- The second thread 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.

Notably, all papers in both threads also consider merely the single-product case à la Gallego and Van Ryzin (1994). Against this backdrop, the present paper makes the following contribution:

We demonstrate that for a broad class of customer utility models, static pricing is asymptotically optimal in the regime where inventory and demand grow large. We establish our result for the multiproduct case, where differentiated products are sold subject to joint capacity constraints. For the single-product case, 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.

Our work thus bears a strikingly simple economic message: The asymptotic optimality of static prices established in the aforementioned foundational work in RM 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 and 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 multidimensional 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.

Notably, optimality of static prices does not emerge from the choice of regime, but rather from the customer utility class we consider. In particular, we show via an example in Section 5.1 that if customer utilities lie outside this broad class, static prices could be suboptimal, even when inventory and demand grow arbitrarily large, as in the scaling regime of Gallego and Van Ryzin (1994). Put differently, our work introduces a classification of utility models into those that admit static pricing as an optimal policy and those that potentially call for more sophisticated policies.

It is also important to note that customers in our study do not choose which product to buy. Rather, each customer has her own desired product that she wishes to purchase, but is forward looking—that is, strategizes only over when (and if) to make that purchase. All papers in the literature that study choice behavior do so by assuming customers who are not forward looking. On the flipside, all papers in the literature that study forward-looking customers, including ours, do so by ignoring choice behavior. Considering forward-looking customers who also choose what to buy appears to be very challenging. To start with, considering both these behavioral aspects would further increase the dimension of a customer’s private attributes. For the purposes of the methodology in this paper, this increased dimensionality would make the analysis from its second step (Proposition 1) onward inapplicable, calling for novel lines of attack.

A detailed discussion of our contributions vis-à-vis the extant literature follows in Section 1.1. 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 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, discusses the key features of the parsimonious customer utility model we assume, and 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, as presented in Section 4. In Section 5, we establish that static pricing is asymptotically optimal by demonstrating lower bounds on the static price policy that are directly comparable to the upper bound derived from our relaxation. 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. In Section 6, we focus on the single-product setting and establish our regime-independent constant factor guarantee, alongside a problem instance showing this guarantee to be tight. Section 7 concludes with comments on possible extensions and 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 texts 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 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 single-product setting where available inventory and demand grow large. Follow-up papers extended this result for the multiproduct case (Talluri and Van Ryzin 2006). 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. (2017) and Li et al. (2014), has established that this forward-looking behavior is highly prevalent. Interestingly, the paper by Moon et al. (2017) 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 the aforementioned foundational work to the setting where customers are forward looking, for a broad class of customer utility models.

Antecedent literature on RM in the face of strategic customers that is most relevant to this 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 (2016) consider a discrete time version of the same model and, assuming a *finite* horizon, compute the optimal dynamic mechanism. Board and Skrzypacz (2016) 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 intertemporal 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. However, when deadlines are private information, the authors illustrate that the optimal dynamic mechanism is substantially harder to characterize.

Fiat et al. (2016) study a revenue maximization problem in a context wherein a customer's product valuation and the deadline to receive the product are both her private information. The optimal mechanism is characterized as follows. The seller offers a menu of contracts to the customer. Each contract specifies a delivery day and offers a randomized price drawn from a distribution that is a function of the specified delivery day. The authors adopt a new way of combining and ironing revenue curves. They use the duality approach to compute the optimal allocation rule.

In recent work, Chen and Farias (2018) 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 (2018), and, as already discussed, we allow for a customer's valuation and his disutility from a delay to be positively correlated, which is something that Chen and Farias (2018) 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 to study single-product settings, and ignoring distinctions such as discrete time modeling versus continuous time modeling, and so forth, we consider a general setting. Specifically, we allow for (a) multiple products, (b) a rich class of customer utility models, (c) heterogeneity in the sensitivity to delay, and (d) inventory to be limited. Despite this generality, we show that a simple policy—a fixed price—is asymptotically optimal. For the special case of a single product, we also provide a uniform performance guarantee.

Setting Price Schedules in the Presence of Strategic Customers. This stream of relevant literature foregoes optimal mechanisms to focus exclusively on committing to (potentially time-varying) price schedules. Among the first papers in this vein is Stokey (1979). She considered a class of customer utility functions subsumed by the model we study wherein the functional form prescribing a customer's 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 to be 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 does 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 at which disutility increases with valuation. Loosely speaking, as long as this increase is sublinear (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. (2016) 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 mold 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) and Wang (2016).

Vis-à-vis the work above on setting optimal price schedules, our work sheds light on the conundrum of when to use “promotions” (or nonstatic 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 in charge of selling n differentiated products to customers in continuous time over a fixed selling season $[0, T]$. Products, indexed by $j \in \{1, \dots, n\}$, correspond to bundles of m different resources, indexed by $i \in \{1, \dots, m\}$. In particular, the j th product is a bundle including $A_{ij} \in \{0, 1\}$ units of the i th resource. Both resources and products are indivisible. In the beginning of the selling season, the seller has some given inventory of resources $x_0 \in \mathbb{R}^m$ available, with no replenishment opportunity thereafter. That is, at $t = 0$ the seller possesses x_0^i units of the i th resource.¹

During the selling season, the seller implements an anonymous posted price mechanism by dynamically posting prices for each product. Let $\pi_t \in \mathbb{R}^n$ be the posted prices at time t . In case a customer decides to purchase a unit of product j at time t , the seller generates revenues of π_t^j , and her i th resource inventory is reduced by A_{ij} . We denote the resource inventory the seller possesses at time t by $X_t \in \mathbb{R}^m$. Of course, X_t depends on π_t , but we suppress this dependence to ease notation. We require π_t to be left continuous and to depend only on the history of the pricing and inventory processes—that is, to be adapted to \mathcal{F}_{t-} , where $\mathcal{F}_t = \sigma(\pi_t, X_t)$. In addition, we require $\pi_t^j = \infty$ if $X_{t-}^i = 0$ for some i such that $A_{ij} = 1$, and $\pi_t^j < \infty$ otherwise.

Products are highly differentiated, and each arriving customer desires to purchase a unit of a particular

product type only. Customers desiring to purchase product j arrive over time, according to an exogenous Poisson process of rate λ_t^j at each time t . An arriving customer is also endowed with a valuation for his desired product, $v \in \mathbb{R}_+$, and a collection of K^j attributes, $\theta \in \mathbb{R}_+^{K^j}$. 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

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

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 made a purchase or not. Specifically, $a_\phi = 1$ if customer ϕ decides to purchase and the seller has sufficient resource inventory,² that is, $X_{\tau_\phi-}^i > 0$ for all i such that $A_{ij_\phi} = 1$; otherwise $a_\phi = 0$. The customer pays then the seller the amount $p_\phi = \pi_{\tau_\phi}^{j_\phi} a_\phi$ and garners utility

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

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 will discuss our assumptions on this structure shortly.

We assume that a customer’s type ϕ except his desired product j_ϕ is private information. Recall that a customer’s type is specified by the tuple $(j_\phi, t_\phi, v_\phi, \theta_\phi)$. This is in contrast with the typical model that specifies type based only on time of arrival and valuation—that is, (t_ϕ, v_ϕ) ; see, for instance, Aviv and Pazgal (2008), Board and Skrzypacz (2016), Caldentey and Vulcano (2007), Gallien (2006), and 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.

The arrival times of customers who consider the same product 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 RM literature on strategic customers. Aviv and Pazgal (2008), Besbes and Lobel (2015), Board and Skrzypacz (2016), Gallien (2006), Vulcano et al. (2002), and Yin et al. (2009) all make such an assumption and point out that a primary motivation for dynamic pricing is intertemporal price discrimination, which remains relevant despite the assumption.

Because 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 by

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

Put a different way, any restrictions to the nature of the correlation of v_ϕ and θ_ϕ 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.³ After we formalize the disutility model below, we provide further discussion and examples. We make no assumptions on the correlation between t_ϕ and θ_ϕ .

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

Assumption 1. *The virtual value function of the valuation distribution for each product j , $v - \frac{\bar{F}^j(v)}{f^j(v)}$, is nondecreasing in v and has a nonnegative root v_j^* .*

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 the following:

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 nondecreasing 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 payment 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 is 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. The paper by Stokey (1979) on intertemporal price discrimination provides a foundation on which to make such an assumption. Modern papers in RM and service operations more generally, also make such an assumption; see, for instance, Katta and Sethuraman (2005), Gallien (2006), Aviv and Pazgal (2008), Doroudi et al. (2013), Kilcioglu and Maglaras (2015), Afèche and Pavlin (2016), Board and Skrzypacz (2016), Moon et al. (2017), Gurvich et al. (2018), and Nazerzadeh and Randhawa (2018). 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 sublinear. 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. 2017). A natural model for the 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, Rob (1985), Anderson and Renault (1999), 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,

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

If $h^j(\cdot)$ were a nonnegative, nondecreasing, 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 customers' disutility from a delay in service systems, including Doroudi et al. (2013), Afèche and Pavlin (2016), and Gurvich et al. (2018), assume such a model, taking $h^j(\cdot)$ to be a linear function. Nazerzadeh and Randhawa (2018) and Katta and Sethuraman (2005) assume that $h^j(\cdot)$ is sublinear; 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

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

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

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

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

which in turn is a canonical model both in the economics-oriented literature on dynamic pricing for perishable products such as Board and Skrzypacz (2016) and also the revenue-management literature—for example, Gallien (2006) and Aviv and Pazgal (2008). 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, although 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 who 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 (2016) 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 at least one of the required bundled resources for the customer's desired product j_ϕ is out of stock, then the price for product j_ϕ posted by the seller is formally infinite, so that the customer will choose to leave without making a purchase (so that $a_\phi = 0$). Otherwise, if no other customers present themselves at time t , then customer ϕ 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 for either the same or different products (an unlikely event but one that cannot be ruled out), the seller makes the allocation in a random order. If at least one resource does not have enough inventory to satisfy all customers' demand, then clearly some customers will not be able to make a purchase; denote by B_t^ϕ the random indicator that the seller allocates product j_ϕ to customer ϕ if he presents himself to make a purchase at time t . Of course, $B_t^\phi = 1$ if $X_{t-}^i > 0$ for all

resource i with $A_{ij_\phi} > 0$ and ϕ is the only customer to request a product at time t and $B_t^\phi = 0$ if $X_{t-}^i = 0$ for some resource i with $A_{ij_\phi} = 0$. 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

$$U^*(\phi, t) \triangleq \begin{cases} v_\phi - \pi_t^{j_\phi} - M(\phi, (t, 1, \pi_t^{j_\phi})) & \text{if } B_t^\phi = 1 \\ -M(\phi, (t, 0, 0)) & \text{otherwise.} \end{cases}$$

Customers strategize about the time of their purchase and use stopping rules contingent on their type that constitute a symmetric Bayes Perfect Nash equilibrium. Such an equilibrium can be formally defined by a map τ^π from types to stopping rules.⁴ In particular, at each point of time t , each customer can observe his desired product's historical prices up to time t (historical sales are only observed by the seller). For customer type ϕ , $\tau^\pi(\phi) \triangleq \tau_\phi$ is a stopping rule with respect to the filtration generated by product j_ϕ 's historical price process $\mathcal{P}_t^{j_\phi} = \sigma(\{\pi_s^{j_\phi} : s \in [0, t]\})$. The stopping rule is derived as a solution to the optimal stopping problem

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

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

Now consider that the seller uses the pricing policy π , and let τ^π be an equilibrium stopping rule for such a policy. Similarly, let p^π be the induced equilibrium payment process. The seller's expected revenue is then given by

$$J_{\pi, \tau^\pi}(x_0, T) = \mathbb{E} \left[\sum_{\phi \in h^T} p_\phi^\pi \right],$$

where $h^t \triangleq \{\phi : t_\phi \leq t\}$ is the set of customer types that arrive up to time t , for all $t \in [0, T]$. 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 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) induces 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. Because 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. Efficacy of Static Prices

As discussed in the Introduction, Gallego and Van Ryzin (1994) proposed the use of a simple *static price policy* for revenue management problems of the type we have just discussed. They showed that in a single-product setting ($n = m = 1$) where inventory and the customer arrival rate grow large, such a policy is asymptotically optimal. Follow-up papers generalized this result for multiple products. This body of work 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 intertemporal 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 intertemporal price discrimination is limited for the broad class of utility models we consider. Specifically, static pricing remains to be, 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 what follows in this section, we first consider a static price policy that is optimal in a "fluid" setting. Under such a static 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}_+^n , and consider the following fluid optimization problem:

$$\begin{aligned} & \underset{\hat{\pi}}{\text{maximize}} && \sum_{j=1}^n \int_0^T \lambda_t^j \hat{\pi}_t^j \bar{F}^j(\hat{\pi}_t^j) dt && (1) \\ & \text{subject to} && \sum_{j=1}^n A_{ij} \int_0^T \lambda_t^j \bar{F}^j(\hat{\pi}_t^j) dt \leq x_0^i, \quad \forall i. \end{aligned}$$

This problem treats customers as myopic and infinitesimal (hence, fluid). It is easy to show that it admits an optimal solution that is static. In particular, an optimal solution to the following optimization problem is also optimal for (1):

$$\begin{aligned} & \underset{\{\hat{\pi}^j \in \mathbb{R}_+, \forall j\}}{\text{maximize}} && \sum_{j=1}^n \left(\int_0^T \lambda_t^j dt \right) \hat{\pi}^j \bar{F}^j(\hat{\pi}^j) \\ & \text{subject to} && \sum_{j=1}^n A_{ij} \left(\int_0^T \lambda_t^j dt \right) \bar{F}^j(\hat{\pi}^j) \leq x_0^i, \quad \forall i. \end{aligned}$$

Let $\pi^{\text{FP}}(x_0, \int_0^T \lambda_t dt)$ be an optimal solution and $J^f(x_0, \int_0^T \lambda_t dt)$ be the optimal value of this optimization problem. To ease notation, unless necessary, we henceforth drop the arguments in π^{FP} and J^f . Note that in the single-product case—that is, when the seller sells a single product produced using a single resource ($m = n = 1$)—Gallego and Van Ryzin (1994) showed that π^{FP} takes the form

$$\pi^{\text{FP}}\left(x_0, \int_0^T \lambda_t dt\right) = \bar{F}^{-1}\left(\min\left\{\frac{x_0}{\int_0^T \lambda_t dt}, \bar{F}(v^*)\right\}\right).$$

We consider the following static pricing policy $\pi^{\text{FP}} = \{\pi_t^{\text{FP},j} : j = 1, \dots, n, t \in [0, T]\}$:^{5,6}

$$\pi_t^{\text{FP},j} = \begin{cases} \pi^{\text{FP},j} & \text{if } X_{t-}^i > 0, \quad \forall i \text{ with } A_{ij} > 0 \\ \infty & \text{otherwise.} \end{cases}$$

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 nonincreasing on $t \geq t_\phi$ on every sample path, because of the first two requirements of Assumption 2. In fact, if the disutility function $M(\cdot, \cdot)$ were strictly positive for positive delays, myopic behavior would be a strongly dominant strategy.

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 nontrivial 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. We will produce this upper bound by allowing the seller to use 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 inasmuch as it yields a useful upper 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 notation. 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, $a_\phi \in \{0, 1\}$ is an indicator for whether a unit of the ϕ -customer's desired product is allocated, and $p_\phi \geq 0$ is the price paid by the customer. Note that, unlike the dynamic pricing setting, the customer *explicitly reports his type* to the seller in this setup, although he may potentially lie. In particular, if a customer of type ϕ is truth telling, he reports his type as ϕ , otherwise as some type $\hat{\phi} \neq \phi$. Encoded in this report is the customer's time of arrival. Of course, if the customer lies, his reported time of arrival, $t_{\hat{\phi}}$, cannot be earlier than his true time of arrival, t_ϕ . When the seller receives a report of type ϕ , she then determines whether that customer is allocated his desired 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 products by overconsuming any resource: $\sum_{\phi \in h^T} A_{ij_\phi} a_\phi \leq x_0^i$, almost surely (a.s.) for all i .

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:

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

Denote by $J^*(x_0, T)$ the optimal value obtained in the problem above. Chen and Farias (2018) establish that in the single-product, single-resource setting, for any dynamic pricing policy there exists a direct dynamic mechanism that satisfies the constraints of (2) and has the objective value equal to the seller's revenue under the dynamic pricing policy. By extending their arguments in a multiproduct setting, we have the following result:

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

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

A formal proof is omitted for brevity, but can be readily derived based on the work of Chen and Farias (2018). 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 the seller 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. Principal Results

Our principal result establishes 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 consider the fluid regime studied originally by Gallego and Van Ryzin (1994), whereby inventory and the scale of demand grow large simultaneously. Now our setting is similar, 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 intertemporal price discrimination and mechanism design more generally—for example, Besbes and Lobel (2015) and Liu and Cooper (2015).

In particular, 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 for the broad class of customer utility models, we consider, the static price policy is asymptotically optimal in the “fluid regime.” Specifically, we will show in Section 5 that, provided Assumptions 1 and 2 are satisfied,

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

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. Our result provides a significant generalization to the conclusions drawn by Gallego and Van Ryzin (1994). Whereas their conclusions rest heavily on the assumption that customers were myopic and short lived, our analysis shows that those conclusions are robust to potentially long-lived customers who strategize on the timing of their purchase.

Furthermore, for the single-product case, we are able to derive another performance guarantee for static prices that is uniform and nonasymptotic—that is, it is relevant over all parameter regimes. In particular, in Section 6.1, we show that, provided $n = m = 1$ and Assumptions 1 and 2 are satisfied, we have

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

In addition, we show that this guarantee is also tight in the sense that a specific problem instance achieves the bound implicit in the guarantee. 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 for the single-product setting. 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.

For utility models that lie outside the class we consider—that is, if Assumption 2 does not hold—we show via an example that static prices could be suboptimal even when inventory and demand grow large. This finding reinforces that the fluid regime considered here preserves the potential relevance of intertemporal price discrimination and mechanism design more generally, as noted above.

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

$$\begin{aligned} & \text{maximize} && \mathbb{E} \left[\sum_{\phi \in h^T} \left(v_\phi - \frac{\bar{F}^{j_\phi}(v_\phi)}{f^{j_\phi}(v_\phi)} \right) \bar{a}_\phi \right] && (3) \\ & \text{subject to} && \mathbb{E} \left[\sum_{\phi \in h^T} A_{ij_\phi} \bar{a}_\phi \right] \leq x_0^i, \quad \forall i \\ & && \bar{a}_\phi \in [0, 1], \quad \forall \phi \text{ with } v_\phi > 0. \end{aligned}$$

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

$$\phi_{v'} \triangleq (j_\phi, t_\phi, v', \theta_\phi), \quad \forall v' \geq 0,$$

and consider the following weaker incentive compatibility constraint:

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

(IC') is a relaxation of (IC) because 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),

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

But because (IR) requires $\mathbb{E}_{-\phi} [U(\phi_0, y_{\phi_0})] \geq 0$, we must have (4). Therefore, we have $\mathbb{E}_{-\phi} [p_{\phi_0}] = 0$ and $\mathbb{E}_{-\phi} [M(\phi_0, y_{\phi_0})] = 0$.

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

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

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 third equality follows from (4). Consequently,

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

The first equality follows from the definition of $U(\cdot)$. The second equality follows from our application of the envelope theorem above. The third equality follows from the definition of $\bar{a}_{\phi_{v'}}$, and the fourth equality from

the property that $\mathbb{E}_{-\phi} [M(\phi_0, y_{\phi_0})] = 0$. Finally, by the assumed concavity of $M(\cdot)$ in v_ϕ , we have the inequality that $M(\phi, y_\phi) \geq v_\phi m(\phi, y_\phi) + M(\phi_0, y_{\phi_0})$. ■

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 ,*

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

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:

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

Here, the expectation $\mathbb{E}_{v_\phi}^{j_\phi}[\cdot]$ in the first and the fifth equalities is with respect to v_ϕ that has the probability density function (p.d.f.) $f^{j_\phi}(\cdot)$, 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 uses the fact that v_ϕ is independent of θ_ϕ and t_ϕ . This completes the proof of the lemma. □

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

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

and

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

which upon substitution in the previous inequality yields:

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

Therefore, we conclude that $\mathbb{E}_{-\phi}[a_{\phi_v} - m(\phi_v, y_{\phi_v})]$ is nondecreasing 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, coupled with the fact that $\mathbb{E}_{-\phi}[a_{\phi_v} - m(\phi_v, y_{\phi_v})]$ is nondecreasing in v , lets us conclude readily that

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

for all $v > 0$. Finally, 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 complete the proof. ■

The next lemma establishes that we cannot allocate more inventory than available:

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

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

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

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 $\bar{J}^*(x_0, T)$. Lemmas 4–6 yield the following result:

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 enable a comparison with the expected revenues under the static price policy.

5. Asymptotic Optimality of Static Prices

We are now ready to derive our principal result—namely, that static prices are asymptotically optimal for our fairly general revenue management problem in the face of strategic consumers, for a broad class of utility models. In particular, our analysis proof proceeds 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, and yields the following result.

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

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

Before presenting the proof, let us summarize what Theorem 1 enables us to establish. 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 essentially an intractable optimal dynamic mechanism problem (Pai and Vohra 2013), includes virtually any selling format one may imagine. The principal insight in this paper is the surprising fact that a mechanism as simple as a static posted price is, to a first order, optimal. Theorem 1 makes 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.

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 $\lambda_t^{(n)} = n\lambda_t$ for all $t \in [0, T]$ 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.

We first 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, $J^f(x_0, \int_0^T \lambda_t dt)$.

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

$$J^f\left(x_0, \int_0^T \lambda_t dt\right) \geq \bar{J}^*(x_0, T).$$

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

$$\begin{aligned} & \underset{y^T \in \mathcal{Y}}{\text{maximize}} \quad \mathbb{E}\left[\sum_{\phi \in \mathcal{H}^T} \left(v_\phi - \frac{\bar{F}^\phi(v_\phi)}{f^\phi(v_\phi)} - \sum_{i=1}^m A_{ij_\phi} \eta^i\right) \bar{a}_\phi\right] + \sum_{i=1}^m \eta^i x_0^i \\ & \text{subject to} \quad \bar{a}_\phi \in [0, 1], \forall \phi \text{ with } v_\phi > 0, \end{aligned} \quad (6)$$

where $\eta \in \mathbb{R}_+^m$. Let $\bar{J}^{*,\eta}(x_0, T)$ be its optimal value. Now, for any feasible mechanism $y^T \in \mathcal{Y}$ in the relaxed mechanism design problem (3), and any $\eta \in \mathbb{R}_+^m$, we have

$$\eta^i \left(x_0^i - \mathbb{E}\left[\sum_{\phi \in \mathcal{H}^T} A_{ij_\phi} \bar{a}_\phi\right]\right) \geq 0, \forall i.$$

It follows that $\min_{\eta \in \mathbb{R}_+^m} \bar{J}^{*,\eta}(x_0, T) \geq \bar{J}^*(x_0, T)$, 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

$$\begin{aligned} & \left(v_\phi - \frac{\bar{F}^\phi(v_\phi)}{f^\phi(v_\phi)} - \sum_{i=1}^m A_{ij_\phi} \eta^i\right) \bar{a}_\phi \\ & \leq \left(v_\phi - \frac{\bar{F}^\phi(v_\phi)}{f^\phi(v_\phi)} - \sum_{i=1}^m A_{ij_\phi} \eta^i\right) \mathbf{1}\left\{v_\phi - \frac{\bar{F}^\phi(v_\phi)}{f^\phi(v_\phi)} \geq \sum_{i=1}^m A_{ij_\phi} \eta^i\right\}. \end{aligned}$$

Moreover, the set on which $v_\phi = 0$ is of measure zero by assumption, so that for any $\eta \in \mathbb{R}_+^m$, we immediately have

$$\begin{aligned} \bar{J}^{*,\eta}(x_0, T) & \leq \mathbb{E}\left[\sum_{\phi \in \mathcal{H}^T} \left(v_\phi - \frac{\bar{F}^\phi(v_\phi)}{f^\phi(v_\phi)} - \sum_{i=1}^m A_{ij_\phi} \eta^i\right) \right. \\ & \quad \cdot \mathbf{1}\left\{v_\phi - \frac{\bar{F}^\phi(v_\phi)}{f^\phi(v_\phi)} \geq \sum_{i=1}^m A_{ij_\phi} \eta^i\right\} \left. + \sum_{i=1}^m \eta^i x_0^i\right] \\ & = \sum_{j=1}^n \left(\int_0^T \lambda_t^j dt\right) \mathbb{E}_v^j \left[\left(v - \frac{\bar{F}^j(v)}{f^j(v)} - \sum_{i=1}^m A_{ij} \eta^i\right) \right. \\ & \quad \cdot \mathbf{1}\left\{v - \frac{\bar{F}^j(v)}{f^j(v)} \geq \sum_{i=1}^m A_{ij} \eta^i\right\} \left. + \sum_{i=1}^m \eta^i x_0^i\right], \end{aligned}$$

where the expectation $\mathbb{E}_v^j[\cdot]$ is with respect to v with the p.d.f. $f^j(\cdot)$, and the equality follows from Wald's identity.

Now, for each i , let $\hat{\eta}^i$ be the optimal dual variable associated with the inventory constraint for the i th resource in optimization problem (1). We have

$$\begin{aligned} \bar{J}^{*,\hat{\eta}}(x_0, T) & \leq \sum_{j=1}^n \left(\int_0^T \lambda_t^j dt\right) \mathbb{E}_v^j \left[\left(v - \frac{\bar{F}^j(v)}{f^j(v)} - \sum_{i=1}^m A_{ij} \hat{\eta}^i\right) \right. \\ & \quad \cdot \mathbf{1}\left\{v - \frac{\bar{F}^j(v)}{f^j(v)} \geq \sum_{i=1}^m A_{ij} \hat{\eta}^i\right\} \left. + \sum_{i=1}^m \hat{\eta}^i x_0^i\right] \\ & = \sum_{j=1}^n \left(\int_0^T \lambda_t^j dt\right) \int_0^\infty \left(v - \frac{\bar{F}^j(v)}{f^j(v)} - \sum_{i=1}^m A_{ij} \hat{\eta}^i\right)^+ \\ & \quad \cdot f^j(v) dv + \sum_{i=1}^m \hat{\eta}^i x_0^i \\ & = \sum_{j=1}^n \left(\int_0^T \lambda_t^j dt\right) \int_{v=\pi^{\text{FP},j}}^\infty \left(v - \frac{\bar{F}^j(v)}{f^j(v)} - \sum_{i=1}^m A_{ij} \hat{\eta}^i\right) \\ & \quad \cdot f^j(v) dv + \sum_{i=1}^m \hat{\eta}^i x_0^i \\ & = \sum_{j=1}^n \left(\int_0^T \lambda_t^j dt\right) \pi^{\text{FP},j} \bar{F}^j(\pi^{\text{FP},j}) \\ & \quad + \sum_{i=1}^m \hat{\eta}^i \left(x_0^i - \sum_{j=1}^n \left(\int_0^T \lambda_t^j dt\right) A_{ij} \bar{F}^j(\pi^{\text{FP},j})\right) \\ & = \sum_{j=1}^n \left(\int_0^T \lambda_t^j dt\right) \pi^{\text{FP},j} \bar{F}^j(\pi^{\text{FP},j}). \end{aligned}$$

The second equality follows from the first-order condition in optimization problem (1) that

$$\begin{aligned} & \frac{\partial}{\partial v} \left(v \bar{F}^j(v) - \sum_{i=1}^m A_{ij} \hat{\eta}^i\right) \Big|_{v=\pi^{\text{FP},j}} \\ & = -f^j(\pi^{\text{FP},j}) \left(\pi^{\text{FP},j} - \frac{\bar{F}^j(\pi^{\text{FP},j})}{f^j(\pi^{\text{FP},j})} - \sum_{i=1}^m A_{ij} \hat{\eta}^i\right) = 0, \forall j, \end{aligned}$$

and Assumption 1 that $v - \frac{\bar{F}^j(v)}{f^j(v)}$ is nondecreasing. The third equality uses the fact that $\int_{v=p}^\infty v f^j(v) - \bar{F}^j(v) dv = p \bar{F}^j(p)$. The fourth equality follows from the complementary slackness condition in optimization problem (1) that

$$\sum_{i=1}^m \hat{\eta}^i \left(x_0^i - \sum_{j=1}^n \left(\int_0^T \lambda_t^j dt\right) A_{ij} \bar{F}^j(\pi^{\text{FP},j})\right) = 0, \forall i.$$

We have thus shown that

$$J^f \geq \bar{J}^{*,\hat{\eta}}(x_0, T) \geq \min_{\eta \in \mathbb{R}_+^m} \bar{J}^{*,\eta}(x_0, T) \geq \bar{J}^*(x_0, T),$$

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 (1997) establishes

$$\frac{J_{\pi^{\text{FP}}, \tau^{\text{FP}}}(\mathbf{x}_0, T)}{J^f(\mathbf{x}_0, \int_0^T \lambda dt)} \geq 1 - \frac{\sum_{i=1}^m \left(\max_{j: A_{ij} > 0} \pi^{\text{FP}, j} \right) \sqrt{\sum_{j=1}^n A_{ij} \left(\int_0^T \lambda_t^j dt \right) \bar{F}_i(\pi^{\text{FP}, j})}}{2 \sum_{j=1}^n \pi^{\text{FP}, j} \left(\int_0^T \lambda_t^j dt \right) \bar{F}_i(\pi^{\text{FP}, j})}.$$

Now, in the n th problem, we consider an arrival rate process of $\{n\lambda_t : t \in [0, T]\}$, and initial inventory of $n\mathbf{x}_0$, so that

$$\frac{J_{\pi^{\text{FP}}, \tau^{\text{FP}}}^{(n)}(n\mathbf{x}_0, T)}{J^f(n\mathbf{x}_0, \int_0^T n\lambda dt)} \geq 1 - \frac{\sum_{i=1}^m \left(\max_{j: A_{ij} > 0} \pi^{\text{FP}, j} \right) \sqrt{\sum_{j=1}^n A_{ij} \left(\int_0^T \lambda_t^j dt \right) \bar{F}_i(\pi^{\text{FP}, j})}}{2 \sum_{j=1}^n \pi^{\text{FP}, j} \left(\int_0^T \lambda_t^j dt \right) \bar{F}_i(\pi^{\text{FP}, j})} \frac{1}{\sqrt{n}}.$$

We establish in Lemma 7 that $J^f(n\mathbf{x}_0, T) \geq \bar{J}^{*(n)}(n\mathbf{x}_0, T)$, and Proposition 1 shows that $\bar{J}^{*(n)}(n\mathbf{x}_0, T) \geq J^{*(n)}(n\mathbf{x}_0, T)$, so that we have shown

$$\frac{J_{\pi^{\text{FP}}, \tau^{\text{FP}}}^{(n)}(n\mathbf{x}_0, T)}{J^{*(n)}(n\mathbf{x}_0, T)} \geq 1 - \frac{\sum_{i=1}^m \left(\max_{j: A_{ij} > 0} \pi^{\text{FP}, j} \right) \sqrt{\sum_{j=1}^n A_{ij} \left(\int_0^T \lambda_t^j dt \right) \bar{F}_i(\pi^{\text{FP}, j})}}{2 \sum_{j=1}^n \pi^{\text{FP}, j} \left(\int_0^T \lambda_t^j dt \right) \bar{F}_i(\pi^{\text{FP}, j})} \frac{1}{\sqrt{n}}.$$

This completes the proof of Theorem 1. ■

5.1. Suboptimality for General Disutilities

We demonstrate that our assumptions on the customer disutility function $M(\cdot, \cdot)$ were necessary for the asymptotic optimality of static prices. In particular, we give an example of a disutility function for which our guarantee does not hold. The example will serve to illustrate what can go wrong if customer disutility grows sufficiently “rapidly” with valuation.

In particular, recall from our discussion in Section 2.1 on the modeling of customer disutility—that is, the function $M(\cdot, \cdot)$ —that we required that a customer’s disutility be concave and nondecreasing 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. Although 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 superlinearly in valuation. To that end, consider the class of *deadline-based* disutilities given by

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

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 demerits of this specification for now, we focus on showing that static prices are *suboptimal* for a specification such as the one above.⁷ 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 customers with valuations less than one half are fully patient, whereas 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 —that is, 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 a single-product ($n = m = 1$) setting where inventory is also unlimited, $x_0 = \infty$. Now, in this setting, the static price policy would set $\pi^{\text{FP}} = 1/2$, which will garner expected revenue

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

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

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

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

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

Customers with valuation greater than $1/2$ will purchase immediately, and customers with valuation between $1/4$ and $1/2$ will wait until the end of the horizon and then purchase, whereas customers with valuation less than $1/4$ will leave immediately upon arrival without a purchase. This yields the following 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%. Because this relative improvement is independent

of the value of λ , and because 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 in such a setting.


The deadline-based disutility functions above have some interesting features. For instance, for deadlines shorter than the horizon, they enable 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. But, 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, sublinearly) with valuation, static prices suffice. If disutility can grow rapidly with valuation (exemplified by the jump in the deadline-based disutility function), more sophisticated policies are called for.

We further explore the necessity of the concavity assumption in Section A.1 by studying the same setting as here, but with a convex disutility of the form

$$M(\phi, y) = v_\phi^\alpha \mathbf{1}_{\{\tau_y > t_\phi\}},$$

where $\alpha \geq 1$. We focus on values of α that are in the neighborhood of 1 and find that dynamic policies continue to dominate static ones, as long as $\alpha > 1$. In other words, for unlimited inventory, static pricing appears to

Figure 1. (Color online) Relationship Between $M(\phi_v, y)$ and v when $\tau_y > t_\phi$



be (even asymptotically) suboptimal for convex disutilities that get arbitrarily close to linear. Furthermore, we study the role of inventory by reverting back to the setting with deadline-based disutilities, but with finite inventory (see Section A.2). Interestingly, we find that for “sufficiently” scarce inventory, static pricing remains asymptotically optimal. Loosely speaking, these findings suggest that, for unlimited inventory, the boundary we identify in our work within which static pricing suffices for good performance appears to be tight, whereas for scarce inventory it appears to be not.

6. Efficacy of Static Prices in the Single-Product Setting

In this section, we leverage the special structure of the single-product setting ($n = m = 1$) to extend our analysis. Specifically, for this important setting, we show first that the static price policy π^{FP} enjoys a constant factor performance guarantee that applies uniformly across all parameter regimes. Second, we show this guarantee to be tight.

We begin by establishing some useful facts for the single-product setting. Recall that

$$\pi^{\text{FP}} = \bar{F}^{-1} \left(\min \left\{ \frac{x_0}{\int_0^T \lambda_t dt}, \bar{F}(v^*) \right\} \right).$$

If we let $\bar{\lambda} \triangleq \left(\int_0^T \lambda_t dt \right) \bar{F}(\pi^{\text{FP}})$, the optimal value of the fluid optimization problem (1) is precisely $J^f = \bar{\lambda} \pi^{\text{FP}}$. Furthermore, note that when customers behave 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 $\lambda_t \bar{F}(\pi^{\text{FP}})$ so that

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

where \bar{N} is a Poisson random variable with parameter $\bar{\lambda}$.

6.1. A Constant Factor Guarantee for Static Prices

Although the previous section established a performance guarantee in the fluid regime for the general setting, in the single-product setting we study here, we can obtain the following constant factor guarantee that is valid uniformly in all model parameters.

Theorem 2. *Provided $n = m = 1$ and Assumptions 1 and 2 are satisfied, we have*

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

Proof. The first inequality follows from Lemma 7, $J^f \geq \bar{J}^*(x_0, T)$, and Proposition 1, $\bar{J}^*(x_0, T) \geq J^*(x_0, T)$.

Next, we prove the second inequality. To this end, we have

$$\begin{aligned} \frac{J_{\pi^{\text{FP}}, \tau^{\text{FP}}}(x_0, T)}{J^f} &= \frac{1}{\bar{\lambda}} \mathbb{E}[\min(\bar{N}, x_0)] \\ &= \frac{1}{\bar{\lambda}} \left(\mathbb{E}[\bar{N}] - \mathbb{E}[(\bar{N} - x_0)^+] \right) \\ &= 1 - \mathbb{E}[(\bar{N} - x_0)^+] / \bar{\lambda} \\ &= 1 - \sum_{n=1}^{\infty} n \frac{e^{-\bar{\lambda}} \bar{\lambda}^{x_0+n-1}}{(x_0+n)!}. \end{aligned} \quad (7)$$

Note that $\bar{\lambda} \leq x_0$ by the definition of π^{FP} . Furthermore, $e^{-\bar{\lambda}} \bar{\lambda}^{x_0+n-1}$ is nondecreasing in $\bar{\lambda}$ on $n \geq 1$ because

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

Thus, (7) yields

$$\begin{aligned} \frac{J_{\pi^{\text{FP}}, \tau^{\text{FP}}}(x_0, T)}{J^f} &\geq 1 - \sum_{n=1}^{\infty} n \frac{e^{-x_0} x_0^{x_0+n-1}}{(x_0+n)!} \\ &= 1 - \sum_{n=1}^{\infty} ((x_0+n) - x_0) \frac{e^{-x_0} x_0^{x_0+n-1}}{(x_0+n)!} \\ &= 1 - \sum_{n'=0}^{\infty} \frac{e^{-x_0} x_0^{x_0+n'}}{(x_0+n')!} + \sum_{n=1}^{\infty} \frac{e^{-x_0} x_0^{x_0+n}}{(x_0+n)!} \\ &= 1 - \frac{e^{-x_0} x_0^{x_0}}{x_0!} \\ &\geq 1 - \frac{1}{e}, \end{aligned}$$

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

Interestingly, the second inequality in this theorem also yields a uniform performance guarantee on the static price policy in the setting of myopic customers—that is, 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. In addition, the constant factor guarantee also implies a stronger guarantee for the class of “Robust Pricing Policies” proposed by Chen and Farias (2018). In that paper, the authors establish that so-called robust pricing policies provide 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 (2018), and the static price policy is trivially a robust pricing policy, thereby improving the Chen and Farias (2018) guarantee from 29% to $\sim 63.2\%$.

6.2. 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, the customer arrival rate is a constant, λ . Third, 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 $\tau_y \geq t_\phi$.

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*

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

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, it is optimal for the seller to wait for all customers to arrive and proceed to conduct a (static) revenue-maximizing auction at time T . In particular, it is optimal to conduct a second-price Myerson auction with reserve price, which 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.

In this setting, it can be readily seen that the optimal reserve price is $1/2$. Therefore, the seller collects revenues as follows. If the highest submitted bid is less than $1/2$, the seller collects no revenues. Suppose now that the highest submitted bid is higher than $1/2$. If the second highest bid is less than $1/2$, the seller collects precisely $1/2$; otherwise, she collects revenues equal to the second highest bid. Therefore, if there are N bids submitted, the seller’s expected revenues are equal to

$$\begin{aligned} &\int_{1/2}^1 \int_0^{1/2} \frac{1}{2} f_{N,N-1}(u_{[N]}, u_{[N-1]}) du_{[N-1]} du_{[N]} \\ &+ \int_{1/2}^1 \int_{1/2}^{u_{[N]}} u_{[N-1]} f_{N,N-1}(u_{[N]}, u_{[N-1]}) du_{[N-1]} du_{[N]}, \end{aligned}$$

where $u_{[N]}$, $u_{[N-1]}$ are the largest and second-largest submitted bids and

$$f_{N,N-1}(u_{[N]}, u_{[N-1]}) = N(N-1)u_{[N-1]}^{N-2}.$$

is their joint probability density. By evaluating the integrals and taking expectation over number of arrivals N , we obtain that the seller's expected revenues equal

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

But,

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

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

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

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), because $\lim_{\lambda \rightarrow \infty} \left(1 - \frac{1}{\lambda T} \right)^{\lambda T} = \frac{1}{e}$ and $\lim_{\lambda \rightarrow \infty} \mathbb{P}(N > \frac{\lambda T}{2}) = 1$. ■

The result above shows an example where the gap between the static-price revenue and that under an optimal dynamic mechanism is indeed approximately 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. Furthermore, 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:

$$\text{LB}^{\text{FP}}(x_0, T) \triangleq \frac{J_{\pi^{\text{FP}}, \tau^{\pi^{\text{FP}}}}(x_0, T)}{J^f}.$$

[Recall that Lemma 7 and Proposition 1 together established that $J^*(x_0, T) \leq J^f$.]

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 nonasymptotic setting, the static price policy already leaves little room for improvement. In fact, this is the *core reason* that the very intuitive results of Gallego and Van Ryzin (1994) have proved so influential in revenue management.

7. Concluding Remarks

This paper has focused on a canonical revenue management problem 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 because 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.

In concluding, it is worth remarking on the seller's power to commit. In particular, 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

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

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 use the optimal static price, so there is no need for the buyer to verify whether the price used was indeed optimal.

Finally, 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 sublinear disutilities? Another direction is considering more general revenue management problems. For example, problems wherein the seller offers products that are not highly differentiated and may be substitutable. In such a setting, customers could strategize not just over when to buy, but over what to buy, too. This added dimension in the customers' strategy space would pose unique technical challenges that have not been addressed before. Attempting to derive performance guarantees for static prices in that setting would be a particularly exciting direction for future work.

Appendix. Performance for Nonconcave Disutility

We extend our analysis in Section 5.1 and study the performance of static pricing under nonconcave disutilities. First, we consider (strictly) convex disutilities and find dynamic pricing to dominate static pricing, even as the disutility is arbitrarily close to being linear. Second, we consider deadline-based disutilities, but with limited inventory, and find static pricing to remain asymptotically optimal as long as inventory is sufficiently "scarce." Loosely speaking, these findings suggest that, for unlimited inventory, the boundary we identify in our work within which static pricing suffices for good performance, appears to be tight, whereas for scarce inventory it appears to be not.

A.1. Convex Disutilities

Consider the same single-product setting where inventory is also unlimited as in Section 5.1, but suppose that the customers' disutility takes the form

$$M(\phi, y) = v_\phi^\alpha \mathbf{1}_{\{\tau_y > t_\phi\}},$$

where $\alpha \geq 1$. Note that, unless $\alpha = 1$, this disutility function is convex in v_ϕ and violates the concavity requirement in Assumption 2. We will focus on a neighborhood of α that contains values that are strictly greater than 1 (and get arbitrarily

close to 1). For each such value, we will construct a dynamic pricing policy that improves upon the static pricing in terms of expected revenues by a factor that is independent of the value of λ .

In particular, consider a pricing policy, similar to the one we studied in Section 5.1, that instantaneously drops prices at the very end of the horizon:

$$\hat{\pi}_t = \begin{cases} p_H & \text{if } t < T \\ p_L & \text{if } t = T, \end{cases}$$

for some prices p_H, p_L . We will select these prices jointly with threshold valuations $0 \leq v_L \leq v_H \leq 1$, such that the following becomes an equilibrium stopping rule under policy $\hat{\pi}$:

$$\tau_{\hat{\pi}}^\phi = \begin{cases} t_\phi & \text{if } v_\phi \in [v_H, 1] \\ T & \text{if } v_\phi \in [v_L, v_H]. \end{cases}$$

Customers with valuation greater than v_H will purchase immediately, and customers with valuation between v_L and v_H will wait until the end of the horizon and then purchase, whereas customers with valuation less than v_L will leave immediately upon arrival without a purchase.

Specifically, let $v_H = \frac{\alpha}{\alpha+1}$ and v_L be such that $v_L \leq \frac{1}{2}$ and

$$v_H - v_H^\alpha = v_L - v_L^\alpha. \quad (\text{A.1})$$

Uniqueness of such v_L can be readily verified for any $\alpha \in (1, 2]$. We then select prices $p_L = v_L - v_L^\alpha$ and $p_H = v_H^\alpha + p_L$. Using straightforward algebra, it can be verified that our selected prices indeed induce the equilibrium stopping rule above, for our selected threshold valuations.

The alternative pricing policy $\hat{\pi}$ we constructed yields expected revenue

$$R \triangleq \lambda T p_H (1 - v_H) + \lambda T p_L (v_H - v_L).$$

By substituting, we get that the revenues are

$$R = \lambda T (f + (1 - v_L)g),$$

where $f \triangleq v_H^\alpha (1 - v_H)$ and $g \triangleq v_H - v_L^\alpha$. Treating R, f, g , and v_L as functions of α , we get that $f(1) = \frac{1}{4}$ and $g(1) = 0$, yielding

$$R(1) = \frac{1}{4} \lambda T,$$

which precisely matches the static pricing policy. To prove our claim, it suffices to show that $R'(1) > 0$.

To calculate $R'(1)$, first note that

$$R' = \lambda T (f' + (1 - v_L)g' - g v_L').$$

For f' and g' , we have that

$$f'(1) = -\frac{\log 2}{4}, \quad g'(1) = \frac{\log 2}{2}.$$

By (A.1), we can upper-bound v_L by the stationary point of $v^\alpha - v$ in v —that is $v_L < \left(\frac{1}{\alpha}\right)^{\frac{1}{\alpha-1}}$. Therefore, we conclude that $v_L(1) < \frac{1}{e}$. Furthermore, by equating the derivative of $v_H - v_H^\alpha - v_L + v_L^\alpha$ to 0 and using the Implicit Function Theorem, one can obtain an expression for v_L' . Using it, we get

$$g v_L' = \frac{(v_L^\alpha \log v_L + g')g}{1 - \alpha v_L^{\alpha-1}}.$$

Taking the limit, we get

$$g(1)v_L'(1) = \lim_{\alpha \rightarrow 1} g v_L' = -\frac{(\log 2)(\log 2 + 2v_L(1) \log v_L(1))}{8(1 + \log v_L(1))}.$$

For $v_L(1) < 1/e$, it can be readily checked that $g(1)v'_L(1) \leq 0.02$. Combining all these facts, we then get that

$$\begin{aligned} R'(1) &= \lambda T(f'(1) + (1 - v_L(1))g'(1) - v'_L(1)g(1)) \\ &= \lambda T\left(-\frac{\log 2}{4} + (1 - v_L(1))\frac{\log 2}{2} - v'_L(1)g(1)\right) \\ &> \lambda T\left(-\frac{\log 2}{4} + (1 - \frac{1}{e})\frac{\log 2}{2} - 0.02\right) \\ &> 0.025 \lambda T \\ &> 0. \end{aligned}$$

A.2. Deadline-Based Disutilities with Scarce Inventory

We consider the deadline disutility model studied in Section 5.1, except that we assume the seller's initial inventory is limited:

$$x_0 \leq \frac{1}{4} \lambda T.$$

Therefore, the seller can admit at most one-quarter of the mean of the total number of customers arriving over the entire season.

First, we compute an upper bound of the seller's optimal expected revenue, defined as the seller's optimal expected revenue in an auxiliary setting, wherein the seller has superior information that she is able to tell whether a customer's valuation is strictly greater than $1/2$ (myopic customers) or no greater than $1/2$ (fully patient customers). Because all myopic customer valuations are strictly greater than all fully patient customer valuations, the seller's optimal dynamic mechanism is as follows:

(1) For customers whose valuations are strictly greater than $1/2$ (myopic customers), the seller implements an anonymous posted dynamic pricing policy $\{\pi_t^H \in [\frac{1}{2}, 1] : t \in [0, T]\}$ over the entire season.

(2) For customers whose valuations are no greater than $1/2$ (fully patient customers), the seller conducts a static revenue maximizing auction at time T .

We denote by N_t^H the total number of myopic customers who purchase up to time t . Therefore, N_t^H is a Poisson random variable with parameter $\lambda \int_{s=0}^t (1 - \pi_s^H) ds$. We denote by $\Phi_L \triangleq \{\phi : v_\phi \leq 1/2\}$ the collection of all arriving fully patient customers. In the fully patient customer group Φ_L , a customer ϕ 's virtual value function is $v_\phi - 2(1 - 2v_\phi)$.

Therefore, the seller's optimal expected revenue in this superior information setting is the optimal value of the following optimization problem:

$$\begin{aligned} \max_{\substack{\{\pi_t^H \in [1/2, 1] : t \in [0, T]\} \\ \{a_\phi \in \{0, 1\} : \phi \in \Phi_L\}}} & \mathbb{E} \left[\int_{t=0}^T \pi_t^H dN_t^H + \sum_{\phi \in \Phi_L} (v_\phi - 2(1 - 2v_\phi)) a_\phi \right] \\ \text{subject to} & \int_{t=0}^T dN_t^H + \sum_{\phi \in \Phi_L} a_\phi \leq x_0 \quad \text{a.s.} \end{aligned} \quad (\text{A.2})$$

Now, we establish an upper bound of this optimization problem. We consider another optimization problem with the same objective function and the following relaxed inventory constraint:

$$\mathbb{E} \left[\int_{t=0}^T dN_t^H + \sum_{\phi \in \Phi_L} a_\phi \right] \leq x_0.$$

We denote by η the dual variable associated with this constraint. We select $\eta = 1 - 2\frac{x_0}{\lambda T}$. Because $x_0 \leq \frac{1}{4} \lambda T$, we have $\eta \geq \frac{1}{2}$. Therefore, following from the weak duality theorem, the optimization problem (A.2) is upper-bounded by the following function:

$$\begin{aligned} & \max_{\substack{\{\pi_t^H \in [1/2, 1] : t \in [0, T]\} \\ \{a_\phi \in \{0, 1\} : \phi \in \Phi_L\}}} \mathbb{E} \left[\int_{t=0}^T \pi_t^H dN_t^H + \sum_{\phi \in \Phi_L} (v_\phi - 2(1 - 2v_\phi)) a_\phi \right] \\ & + \eta \left(x_0 - \mathbb{E} \left[\int_{t=0}^T dN_t^H + \sum_{\phi \in \Phi_L} a_\phi \right] \right) \\ & = \max_{\substack{\{\pi_t^H \in [1/2, 1] : t \in [0, T]\} \\ \{a_\phi \in \{0, 1\} : \phi \in \Phi_L\}}} \mathbb{E} \left[\int_{t=0}^T (\pi_t^H - \eta) dN_t^H \right. \\ & \left. + \sum_{\phi \in \Phi_L} (v_\phi - 2(1 - 2v_\phi) - \eta) a_\phi \right] + \eta x_0 \\ & = \max_{\substack{\{\pi_t^H \in [1/2, 1] : t \in [0, T]\} \\ \{a_\phi \in \{0, 1\} : \phi \in \Phi_L\}}} \mathbb{E} \left[\int_{t=0}^T (\pi_t^H - \eta) \lambda (1 - \pi_t^H) dt \right. \\ & \left. + \sum_{\phi \in \Phi_L} (v_\phi - 2(1 - 2v_\phi) - \eta) a_\phi \right] + \eta x_0 \\ & = \max_{\pi_t^H \in [1/2, 1]} \lambda T (\pi_t^H - \eta) (1 - \pi_t^H) \\ & + \mathbb{E} \left[\sum_{\phi \in \Phi_L} (v_\phi - 2(1 - 2v_\phi) - \eta)^+ \right] + \eta x_0 \\ & = \max_{\pi_t^H \in [1/2, 1]} \lambda T (\pi_t^H - \eta) (1 - \pi_t^H) + \eta x_0 \\ & = x_0 \left(1 - \frac{x_0}{\lambda T} \right). \end{aligned}$$

The second equality follows from theorem II in Brémaud (1981). The fourth equality follows from the property that for $v_\phi \leq \frac{1}{2}$,

$$\begin{aligned} v_\phi - 2(1 - 2v_\phi) - \eta &\leq \frac{1}{2} - 2 \left(1 - 2 \cdot \frac{1}{2} \right) - \eta = \frac{1}{2} - \eta \\ &= \frac{1}{2} - \left(1 - 2 \frac{x_0}{\lambda T} \right) \leq \frac{1}{2} - \left(1 - 2 \cdot \frac{1}{4} \right) = 0. \end{aligned}$$

Next, we show that this upper-bound revenue can be asymptotically achieved under a fixed price policy $\pi^{\text{FP}} = 1 - \frac{x_0}{\lambda T}$. Under this policy, the total number of customers who are willing to purchase is a Poisson random variable with parameter x_0 . Therefore, following from the proof in theorem 3 of Gallego and Van Ryzin (1994), in a sequence of problems parameterized by n with $\lambda^{(n)} = n\lambda$ and $x_0^{(n)} = nx_0$, we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} J_{\pi^{\text{FP}}, \tau, \pi^{\text{FP}}}^{(n)}(nx_0, T) = x_0 \left(1 - \frac{x_0}{\lambda T} \right).$$

Therefore, the example above indicates that if the seller's inventory is limited, then a static price policy can be asymptotically optimal, even in the presence of a delay disutility function that is not concave in valuation.

Endnotes

¹To ease notation, we use bold font to denote vectors that suppress dummy product or resource superscript indices.

²Multiple customers revealing themselves to the seller at the same time are allocated inventory in random order.

³Modeling potential correlation of v_ϕ and θ_ϕ the way we do here does not imply that our analysis can handle any correlation type, but rather only the ones consistent with the structural properties we require from disutility M .

⁴In the sequel we will at times, with an abuse of notion, use this map and the corresponding stopping rules interchangeably.

⁵To ease notation, we suppress dummy product and resource superscript indices when we study the single-product setting throughout this paper.

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

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

References

- Afèche P, Mendelson H (2004) Pricing and priority auctions in queueing systems with a generalized delay cost structure. *Management Sci.* 50(7):869–882.
- Afèche P, Pavlin JM (2016) Optimal price/lead-time menus for queues with customer choice: Segmentation, pooling, and strategic delay. *Management Sci.* 62(8):2412–2436.
- Anderson SP, Renault R (1999) Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model. *RAND J. Econom.* 30(4):719–735.
- Ata B, Olsen TL (2009) Near-optimal dynamic lead-time quotation and scheduling under convex-concave customer delay costs. *Oper. Res.* 57(3):753–768.
- Aviv Y, Pazgal A (2008) Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing Service Oper. Management* 10(3):339–359.
- Azevedo EM, Gottlieb D (2012) Risk-neutral firms can extract unbounded profits from consumers with prospect theory preferences. *J. Econom. Theory* 147(3):1291–1299.
- Besbes O, Lobel I (2015) Intertemporal price discrimination: Structure and computation of optimal policies. *Management Sci.* 61(1):92–110.
- Board S, Skrzypacz A (2016) Revenue management with forward looking buyers. *J. Political Econom.* 124(4):1046–1087.
- Borgs C, Candogan O, Chayes J, Lobel I, Nazerzadeh H (2014) Optimal multiperiod pricing with service guarantees. *Management Sci.* 60(7):1792–1811.
- Brémaud P (1981) *Point Processes and Queues: Martingale Dynamics* (Springer, New York).
- Caldentey R, Vulcano G (2007) Online auction and list price revenue management. *Management Sci.* 53(5):795–813.
- Caldentey R, Liu Y, Lobel I (2016) Intertemporal pricing under minimax regret. *Operations Res.* 65(1):104–129.
- Chen Y, Farias VF (2018) Robust dynamic pricing with strategic customers. *Math. Oper. Res.* 43(4):1119–1142.
- Diamond PA (1971) A model of price adjustment. *J. Econom. Theory* 3(2):156–168.
- Doroudi S, Akan M, Harchol-Balter M, Karp J, Borgs C, Chayes JT (2013) Priority pricing in queues with a continuous distribution of customer valuations. Technical Report CM-CS-13-109, Computer Science Department, Carnegie Mellon University, Pittsburgh.
- Ellison G, Wolitzky A (2012) A search cost model of obfuscation. *RAND J. Econom.* 43(3):417–441.
- Fiat A, Goldner K, Karlin AR, Koutsoupias E (2016) The Fedex problem. *Proc. 2016 ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 21–22.
- Gallego G, Van Ryzin G (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* 40(8):999–1020.
- Gallego G, Van Ryzin G (1997) A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* 45(1):24–41.
- Gallien J (2006) Dynamic mechanism design for online commerce. *Oper. Res.* 54(2):291–310.
- Gurvich I, Lariviere MA, Ozkan C (2018) Coverage, coarseness and classification: Determinants of social efficiency in priority queues. *Management Sci.* 65(3):1061–1075.
- Haghpanah N, Hartline J (2015) Reverse mechanism design. *Proc. 16th ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 757–758.
- Hu Z, Chen X, Hu P (2016) Dynamic pricing with gain-seeking reference price effects. *Oper. Res.* 64(1):150–157.
- Katta AK, Sethuraman J (2005) Pricing strategies and service differentiation in an M/M/1 queue—A profit maximization perspective. Working paper, Columbia University, New York.
- Kilcioglu C, Maglaras C (2015) Revenue maximization for cloud computing services. *ACM SIGMETRICS Performance Evaluation Rev.* 43(3):76.
- Li J, Granados N, Netessine S (2014) Are consumers strategic? Structural estimation from the air-travel industry. *Management Sci.* 60(9):2114–2137.
- Liu Y, Cooper WL (2015) Optimal dynamic pricing with patient customers. *Oper. Res.* 63(6):1307–1319.
- Liu Q, Van Ryzin GJ (2008) Strategic capacity rationing to induce early purchases. *Management Sci.* 54(6):1115–1131.
- Lobel I (2017) Dynamic pricing with heterogeneous patience levels. Working paper, New York University, New York.
- Milgrom P, Segal I (2002) Envelope theorems for arbitrary choice sets. *Econometrica* 70(2):583–601.
- Moon K, Bimpikis K, Mendelson H (2017) Randomized markdowns and online monitoring. *Management Sci.* 64(3):1271–1290.
- Nazerzadeh H, Randhawa RS (2018) Near-optimality of coarse service grades for customer differentiation in queueing systems. *Production Oper. Management* 7(3):578–595
- Özer Ö, Phillips R (2012) *The Oxford Handbook of Pricing Management* (Oxford University Press, Oxford, UK).
- Pai MM, Vohra R (2013) Optimal dynamic auctions and simple index rules. *Math. Oper. Res.* 38(4):682–697.
- Rob R (1985) Equilibrium price distributions. *Rev. Econom. Stud.* 52(3):487–504.
- Said M (2012) Auctions with dynamic populations: Efficiency and revenue maximization. *J. Econom. Theory* 147(6):2419–2438.
- Stokey NL (1979) Intertemporal price discrimination. *Quart. J. Econom.* 93(3):355–371.
- Talluri KT, Van Ryzin GJ (2006) *The Theory and Practice of Revenue Management*, International Series in Operations Research and Management Science, vol. 68 (Springer Science & Business Media, New York).
- Vulcano G, Van Ryzin G, Maglaras C (2002) Optimal dynamic auctions for revenue management. *Management Sci.* 48(11):1388–1407.
- Wang Z (2016) Intertemporal price discrimination via reference price effects. *Oper. Res.* 64(2):290–296.
- Yin R, Aviv Y, Pazgal A, Tang CS (2009) Optimal markdown pricing: Implications of inventory display formats in the presence of strategic customers. *Management Sci.* 55(8):1391–1408.