MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Articles in Advance, pp. 1–19
https://pubsonline.informs.org/journal/msom
ISSN 1523-4614 (print), ISSN 1526-5498 (online)

Intertemporal Price Discrimination via Randomized Promotions

Hongqiao Chen,[a,*] Ming Hu,[b] Jiahua Wu[c]

[a] School of Management & Engineering, Nanjing University, Nanjing, Jiangsu 210093, China; [b] Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada; [c] Imperial College Business School, Imperial College London, London SW7 2AZ, United Kingdom

*Corresponding author

Contact: chq@nju.edu.cn, https://orcid.org/0000-0002-6776-6954 (HC); ming.hu@rotman.utoronto.ca, https://orcid.org/0000-0003-0900-7631 (MH); j.wu@imperial.ac.uk, https://orcid.org/0000-0002-2254-1858 (JW)

Received: November 19, 2021
Revised: June 9, 2022; September 19, 2022
Accepted: November 15, 2022
Published Online in Articles in Advance: March 6, 2023

https://doi.org/10.1287/msom.2023.1194

Abstract. ***Problem definition*:** The undesirable but inevitable consequence of running promotions is that consumers can be trained to time their purchases strategically. In this paper, we study randomized promotions, where the firm randomly offers discounts over time, as an alternative strategy of intertemporal price discrimination. ***Methodology/results*:** We consider a base model where a monopolist sells a single product to a market with a constant stream of two market segments. The segments are heterogeneous in both their product valuations and patience levels. The firm precommits to a price distribution, and in each period, a price is randomly drawn from the committed distribution. We characterize the optimal price distribution as a randomized promotion policy and show that it serves as an intertemporal price discrimination mechanism such that high-valuation customers would make a purchase immediately at a regular price upon arrival, and low-valuation customers would wait for a random promotion. Compared against the optimal cyclic pricing policy, which is optimal within the strategy space of all deterministic pricing policies, the optimal randomized pricing policy beats it if low-valuation customers are sufficiently patient and the absolute discrepancy between high and low customer valuations is large enough. We extend the model in three directions. First, we consider the case where a portion of customers are myopic and would never wait. We show that the existence of myopic customers is detrimental to the firm's profitability, and the expected profit from an optimal randomized pricing policy decreases as the proportion of myopic customers in the population increases. Second, we consider Markovian pricing policies where prices are allowed to be intertemporally correlated in a Markovian fashion. This additional maneuver allows the firm to reap an even higher profit when low-valuation customers are sufficiently patient by avoiding consecutive promotions but, on average, running the promotion more frequently with a smaller discount size. Lastly, we consider a model with multiple customer segments and show that a two-point price distribution remains optimal, and our conclusion from the two-segment base model still holds under certain conditions that are adopted in the literature. ***Managerial implications*:** Our results imply that the firm may want to deliberately randomize promotions in the presence of forward-looking customers.

Funding: This work was supported by the National Natural Science Foundation of China [Grant 72201124], the Natural Sciences and Engineering Research Council of Canada [Grants RGPIN-2015-06757 and RGPIN-2021-04295], and the Youth Project of the Humanities and Social Science Foundation of the Ministry of Education of China [Grant 22YJC630006].
Supplemental Material: The online appendix is available at https://doi.org/10.1287/msom.2023.1194.

Keywords: pricing and revenue management • OM-marketing interface • dynamic programming • randomized pricing • strategic consumer behavior

1. Introduction

Dynamic or contingent pricing is now widely adopted by firms in a range of sectors, such as the airline and e-commerce industry, and is proven to be a strategy that can improve a firm's bottom line substantially. It is especially prevalent in online retail markets, even for durable goods, because retailers are no longer inhibited by physical price tags that require time and effort to change manually. Many firms manipulate prices frequently and, from the customers' point of view, in a somewhat random fashion. For example, Figure 1 displays the price trajectory of Apple AirPods Pro on Amazon from its launch in November 2019 to April 2020. The price varied between its list price of $249.00 and a discounted price of $234.99. If a price-sensitive consumer wants to wait for the discount of $14, it is not clear when the promotion might happen outside the Christmas shopping season and how long it will last.

Figure 1. (Color online) Pricing Trajectory of Apple AirPods Pro on Amazon from Its Launch in November 2019 to April 2020

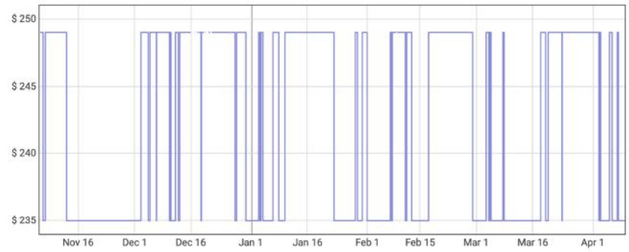

Source: https://keepa.com/#!product/1-B07ZPC9QD4.

When customers are forward-looking, a seller may benefit from deliberately varying prices over time in a *random* fashion. Consider a price-insensitive or delay-sensitive consumer who wants to buy an AirPods Pro arrives at Amazon. Although the consumer is fully aware of the possibility of a promotion in the near future by checking the historical price trajectory through websites such as CamelCamelCamel or Keepa, the random occurrence of such a price discount may deter the consumer from waiting and result in their purchase at the list price upon arrival. Moon et al. (2017, p. 1276) empirically show the success of an established North American specialty retailer in implementing a committed state-independent randomized markdown policy through avoiding "'conditioning' or 'training' its customers to wait for markdowns."[1]

In comparison with running promotions on predictable time instances, there are clearly pros and cons associated with randomized pricing. On the one hand, under a randomized price policy, customers cannot reliably predict future prices, making it hard to time their purchases. Hence randomized pricing would be able to alleviate strategic customer behavior to some extent compared with a deterministic markdown policy, especially when it was close to a predetermined promotion time. However, the flip side of the coin is that, although customers cannot predict future prices accurately, they would expect some price variations at any point in time and still have a tendency to delay their purchase regardless of when their consumption need arises. Thus, it is unclear under what conditions randomized pricing would dominate a deterministic pricing policy, such as a static pricing policy or a cyclic pricing policy. The former completely eliminates strategic customer behavior, and the latter can intertemporally price-discriminate customers (Conlisk et al. 1984). Therefore, a full-blown theoretical investigation is needed to complement the empirical study by Moon et al. (2017).

To that end, we address a set of questions that require formal modeling and a complete scan of the parameter and strategy space. For example, how may randomized promotions and pricing work to benefit the firm? Is it possible that randomized pricing is even worse than a static pricing policy?[2] Under what conditions does randomized pricing outperform the optimal deterministic pricing such as cyclic pricing[3] that has been well studied in the literature and vice versa? Is it beneficial to further expand the search for the optimal pricing strategy from randomized prices independently drawn from a price distribution to more sophisticated randomized pricing policies, such as Markovian pricing, where the price process evolves in a Markovian fashion, for example, whether there will be a promotion in the next period depends on whether there is one in the current period?

We answer these questions with a theoretical model in discrete time. In particular, we consider a base model where a monopolist sells a single product to two segments of customers. Customers arrive sequentially at a constant rate, and they are heterogeneous in both product valuations and patience levels. Upon arrival, a customer can choose to purchase the product immediately, delay purchases for future periods with the hope of a lower price, or leave immediately. The decision of randomized pricing for the firm is to decide on and commit to a price distribution before the sales horizon starts, from which a price is randomly drawn in each period of the sales horizon.

First, we characterize customers' best responses to randomized pricing and structural properties that optimal randomized pricing policies must obey. It turns out, under any optimal randomized pricing policies, high-valuation customers would either purchase or leave immediately upon arrival, whereas low-valuation customers would either buy immediately if the price happens to be sufficiently low or delay the purchase. We then construct a simple two-point price distribution with the higher price to be the valuation of high-type customers and prove its optimality by showing that such a distribution can attain an upper bound of the firm's expected profit. Under this two-point price distribution (when it is nondegenerate to a single mass point), high-valuation customers would always purchase immediately upon arrival, whereas low valuations would always purchase with the lower discounted price either immediately or in a future period. That is, the optimal randomized pricing policy serves as an intertemporal price discrimination mechanism such that high-valuation customers would make a purchase immediately at a regular price upon arrival, and low-valuation customers would wait for a promotion.

Next, we investigate the firm's profitability, consumer surplus, social welfare, and surplus allocation across the two market segments under the optimal randomized pricing policy. We show that, compared with the optimal static pricing policy, the optimal randomized pricing policy yields a higher profit when low-valuation customers are sufficiently patient. Hence, randomized pricing either dominates or reduces to static pricing. If under the optimal static pricing policy, the low-valuation customers are priced out of the market, then the optimal randomized pricing policy leads to (weakly) higher surplus for both market segments and hence higher social welfare (which is the total surplus of the firm and all customers). Otherwise, the optimal randomized pricing policy results in a lower surplus difference between the two market segments. Thus, shifting from static pricing to randomized pricing leads to a (weakly) higher profit for the firm and, moreover, either (weakly) higher social welfare or a (weakly) fairer surplus allocation among customers.

We further evaluate the performance of randomized pricing policies by benchmarking against cyclic pricing policies. We first show that, in our setup, cyclic pricing policies arise to be optimal within the strategy space of all deterministic policies. Under a cyclic pricing policy, the same pricing schedule within a cycle is repeated over time. We show that an optimal cyclic pricing policy follows a (weakly) markdown pattern. Interestingly, neither randomized pricing nor cyclic pricing always dominates the other over the full parameter space. However, when low-valuation customers are sufficiently patient, and the absolute discrepancy in customers' valuations is sufficiently large, an optimal randomized pricing policy is more profitable than an optimal cycle deterministic pricing policy. The underlying rationale is as follows. When the absolute discrepancy in customers' valuations is sufficiently large, it is important for the firm to capture most of the surplus from high-valuation customers. This can be achieved under the optimal randomized pricing policies by a low frequency of promotions, and thus the majority of high-valuation customers would make a purchase at a price equal to their valuation. Meanwhile, the low-valuation customers wait for the promotion as they are sufficiently patient. Conversely, the optimal cyclic pricing policy follows a predetermined price trajectory known to all customers. Thus, the firm needs to compensate high-valuation customers to prevent them from waiting to purchase at a discounted price, especially when it is closer to the end of a cycle. This compensation cuts into the firm's profitability, leading to the underperformance of cyclic pricing policies compared with randomized pricing policies, where high-valuation customers pay at their valuation most of the time.

Last, we extend the base model in three directions. First, we consider the case where a portion of customers are myopic and would never wait. We show that, in the presence of myopic customers, the optimal randomized pricing policy still follows a two-point distribution, which is composed of a regular price and a discounted price that clears the market of strategic low-valuation customers from time to time. As the proportion of myopic customers in the population increases, the firm may either run promotions less frequently with a lower discounted price or run promotions more frequently with a higher discounted price. The optimal strategy depends on the ratio of valuations between the two customer segments. Overall, the existence of myopic customers is detrimental to the firm's profitability, and the expected profit from an optimal randomized pricing policy decreases as the proportion of myopic customers in the population increases. Second, we consider a Markovian pricing policy, where prices are allowed to be intertemporally correlated in a Markovian fashion. We show that Markovian pricing may not always dominate randomized pricing. However, when it generates a higher profit, it does so by avoiding consecutive promotions and extracting weakly more profit from low-valuation customers and strictly more profit from high-valuation customers than an optimal randomized pricing policy. When low-valuation customers are sufficiently patient, the extra maneuver with Markovian pricing policies leads to a higher expected profit through, on average, running the promotion more frequently with a smaller discount size. Third, we consider a model with multiple customer segments. We show that, under certain conditions that are also used in the literature, the optimal randomized pricing policy still follows a two-point price distribution, under which there exists a cutoff on customers' valuations such that any customer with a valuation less than the cutoff would choose to wait, whereas any customer with a valuation greater than or equal to that cutoff would either purchase or leave immediately upon arrival. This randomized pricing policy is guaranteed to outperform an optimal static pricing policy when customers of the lowest valuation are sufficiently patient. Under the same set of conditions that guarantee the optimality of a two-point price distribution for randomized pricing, we show that when the absolute discrepancies in customers' valuations are large enough, an optimal randomized pricing policy is always more profitable than an optimal cyclic pricing policy.

Our results imply that firms may want to deliberately randomize promotions in the presence of forward-looking customers who sequentially arrive because doing so can lead to higher profits. As a default assumption in the literature of dynamic mechanism design, commitment to a randomized pricing policy is practically easier to enforce nowadays than ever, thanks to a provable track record displayed by price tracker websites such as CamelCamelCamel or Keepa. E-commerce websites can develop a reputation of running random promotions to train consumers to make an immediate purchase as they see an acceptable price upon arrival rather than to wait for a lower price.

2. Literature Review

Our paper is related to the economic literature on the usage of price dispersion for the purpose of price discrimination. The origin of this stream of literature can be traced back to the 1970s. Back then, economists have come to recognize that the "law of one price" is no law at all (Varian 1980). A significant degree of price dispersion is observed across a wide range of industries. The first few papers, including but not limited to Salop and Stiglitz (1977), Salop (1977), and Reinganum (1979), seek to provide a rationale regarding why multiple retailers offer an identical item at different prices contemporaneously. This type of price dispersion is also referred to as *spatial* price dispersion. Our work is more closely related to a different kind of price dispersion, namely *temporal* price dispersion, where prices vary over time, and at any moment, a cross-section of the market would exhibit price dispersion. Temporal price dispersion is first studied in the seminal paper by Varian (1980), who shows that the existence of temporal price dispersion is sustained by information heterogeneity in prices. Varian (1980) assumes that there are two segments of consumers, that is, consumers who are informed of all current prices and consumers who are uninformed. Informed consumers would only purchase from the lowest-priced seller, whereas uninformed consumers would visit sellers randomly. The author shows that when there are both informed and uninformed consumers, there is a symmetric mixed equilibrium in which all sellers randomize according to the same price distribution. In that setting, the existence of temporal price dispersion is sustained by information heterogeneity in prices. Since then, temporal price dispersion has been studied with various twists from the information perspective (see Baye and Morgan 2001 and references therein). The randomized pricing policy considered in our paper is also one type of temporal price dispersion. By contrast, we focus on one single firm rather than multiple firms. We unravel the effectiveness of randomized pricing policies by a single firm in serving as an intertemporal price discrimination mechanism.

Along the line of intertemporal price discrimination, researchers have tried to tackle the intertemporal pricing problems with sequentially arriving strategic customers who are forward-looking and able to time their purchases. With strategic customers, the optimal class of policies within the strategy space of all deterministic policies has been shown to be *cyclic* pricing policies. For example, Conlisk et al. (1984) and Besbes and Lobel (2015) study a setting very close to ours with an infinite sales horizon and a steady stream of strategic customers of heterogeneous valuation or patience levels. A subtle difference is that Conlisk et al. (1984) adopt a discounting factor in capturing the customers' cost of waiting, whereas Besbes and Lobel (2015) assume an exogenously given shopping window over which a customer compares prices. By contrast to the shopping window assumption, we assume customers have waiting cost per period as an alternative assumption to a discounting factor[4] and fully endogenize strategic customer behavior as a solution to a dynamic program. Similar to Conlisk et al. (1984) and Besbes and Lobel (2015), we also show that the optimal deterministic pricing policy is a cyclic pricing policy among deterministic policies in our setting. Then, we identify conditions under which randomized pricing outperforms cyclic pricing.

Sobel (1984) extends Conlisk et al. (1984) to account for competition and shows that, in the symmetric equilibrium, competing symmetric sellers use randomized (mixed) strategies. The equilibrium strategy has the following form: There exists a finite time length during which all sellers charge the regular price to sell to high-valuation customers and after which any seller may run a promotion. When a promotion or multiple promotions occur, all of the accumulated low-valuation customers buy from the lowest-price seller, and the process repeats. Sinitsyn (2017) builds on Sobel (1984) but assumes that the sellers must schedule their price promotion in advance over a finite horizon; that is, each seller commits to a future period in which to run a promotion to sell to low-valuation customers. The author shows that sellers often use mixed strategies in equilibrium, choosing the future promotion period according to a probability distribution function. In contrast to these papers that study competition where randomized (mixed) strategies may naturally arise in equilibrium, we study whether a monopolist seller can do better by expanding its strategy space from deterministic pricing policies to randomized pricing policies.

In a parallel setup to the line of Conlisk et al. (1984), Besanko and Winston (1990) assume a finite number of customers lining up in the beginning of the sales horizon and show the subgame perfect Nash equilibrium pricing policy involves intertemporal price discrimination through price skimming. Liu and Zhang (2013) study a competitive version of this setup for two firms offering vertically differentiated products. In contrast, we follow the line of Conlisk et al. (1984) by assuming the entry of new customers over time, which captures the sequential customer arrivals in many practical settings, such as purchasing an AirPods Pro online as customers' needs arise over time.

Aviv and Pazgal (2008) show that strategic customer behavior suppresses the benefits of price segmentation. Chen et al. (2019) point out that static prices eliminate any strategic customer behavior and show its competitive ratio is $1 - 1/e$ for a broad class of customer utility models. In our model, we characterize sufficient conditions under which randomized pricing reduces to static pricing or strictly outperforms static pricing. In addition, various operational-level maneuvers, such as rationing by deliberately understocking products (Liu and van Ryzin 2008), implementing a dry period with zero inventory (Chen and Chu 2020), displaying one unit of inventory at a time (Yin et al. 2009), and offering price guarantees under Markovian pricing (Wu et al. 2020),[5] have been proposed to reduce the propensity of consumers to strategically wait for discounts. Contributing to this literature, we study randomized pricing as a mechanism to alleviate strategic customer behavior. We identify under what conditions such randomized pricing policies can indeed benefit the firm over static pricing or cyclic pricing.

There are two papers that are closest to this paper: Moon et al. (2017) and Wu et al. (2014). Moon et al. (2017) introduce randomized pricing as a mechanism of intertemporal price discrimination and emphasize the customers' active monitoring as a critical factor in determining the success of randomized pricing. Other than that Moon et al. (2017) is a structural estimation model and ours is an analytical model, there are a couple of distinct differences. First, the models have many differences. Moon et al. (2017) study a continuous-time model with an infinite sales horizon, a finite initial inventory level, and a finite market size, whereas we consider a discrete-time model with an infinite horizon, no capacity constraint, and a constant stream of new customers (essentially infinitely many customers). Ours is a common setting to study strategic consumer behavior (Conlisk et al. 1984, Besbes and Lobel 2015). Moon et al. (2017) assume that the firm has the information on the remaining numbers of customers in the market, whereas we do not. Second, the main focuses are different. Moon et al. (2017) compare randomized pricing with contingent pricing dependent on the inventory level and the remaining numbers of strategic customers.[6] As our model has neither inventory constraint nor the information of outstanding strategic customers, state-dependent contingent pricing reduces to a static pricing policy. In addition to an analytical comparison between randomized and static pricing, we also theoretically compare randomized pricing with cyclic pricing, which is proven to be optimal among deterministic policies.[7]

Like ours, Wu et al. (2014) also build a model to study the impact of randomized pricing policies on a firm's profitability. However, our paper differs from Wu et al. (2014) in the following essential aspects. First, Wu et al. (2014) adopt an exogenously given shopping window to capture customers' strategic waiting behavior. In their paper, high-valuation customers are assumed to wait for up to one period, and low-valuation customers would wait for up to a given finite number of periods. In contrast, we assume that each segment of customers have their own cost of waiting per time period and fully endogenize strategic waiting behavior as a solution to an optimal stopping time problem for each individual customer. Second, the price distribution considered in Wu et al. (2014) is beforehand restricted to a two-point distribution. By contrast, we start with general price distributions and characterize the sufficient conditions under which a two-point distribution is optimal. Last, with our general model, we provide a much more comprehensive set of comparisons among randomized, static, and cyclic pricing.

In sum, the main contributions of this paper are threefold. First, we complement the empirical work of Moon et al. (2017) to theorize randomized pricing. Second, we show after expanding the strategy space of the firm from deterministic policies to random policies, the firm can do strictly better and identify conditions under which the firm is better off with randomized pricing, compared with cyclic pricing. Third, we show that the firm can do even better with Markovian pricing, compared with time-independent randomized pricing, and characterize the conditions under which a two-point price distribution remains optimal with multiple customer segments.

3. Model of Randomized Pricing

We consider a monopolist who sells a single product over an infinite horizon. The seller uses a randomized pricing policy. That is, the firm chooses and announces a distribution of prices ex ante, and then a random price P is drawn from the chosen distribution for each period. In our base model, prices are assumed to be independent and identically distributed across different periods, which will be relaxed in Section 4.2 where the price process follows a Markovian process. We denote the price mass or density function by $f(p)$ and the cumulative distribution function by $F(p)$. Because no customer would make a purchase at an infinitely high price, we just consider the support of P to be finite. As general mechanism design without commitment presents well-known difficulties (Bester and Strausz 2001), commitment to a (potentially randomized) policy is the most common approach in the literature nowadays.[8]

Customers are infinitesimal and arrive sequentially over time. We consider a discrete-time model, and the customer arrival rate is normalized to one for each time period. Each customer purchases up to one unit of the product. Upon arrival, a customer decides among three options: purchase immediately in the current period, wait in the system with the hope of a more favorable price in future periods, or leave immediately. If a customer chooses to delay the purchase, a waiting cost is incurred for each period.[9] Customers are heterogeneous along two dimensions, namely, valuation and patience level. We assume there are two customer segments.[10] A fraction, α, of low-valuation customers (bargain hunters) value the product at v_L, and the remaining, $1-\alpha$, fraction of customers value the product at v_H, where $v_H > v_L$. The one-period waiting costs for low- and high-valuation customers are denoted by c_L and c_H, respectively. We do not require that low-valuation customers be more patient than high-valuation customers, that is, c_L can be larger than c_H.

We assume that customers benefit from waiting, that is, a customer's waiting cost per period is less than their valuation. Otherwise, it becomes trivial that customers would either purchase the product or leave immediately upon arrival.

Assumption (S) (Strategic Customers). *Suppose* $0 < c_L < v_L$ *and* $0 < c_H < v_H$.

3.1. Optimal Customer Response

We first characterize the optimal response of customers under *any* given price distribution. It suffices to consider one segment of customers with a valuation of v and a one-period waiting cost of c. Suppose a customer arrives in period t and makes a purchase in period t'. Their utility is thus given by $v - p_{t'} - (t' - t)c, \ \forall t' \geq t$. Upon arrival, a customer makes a choice among three options: buy immediately (buy), delay purchase (wait), or leave immediately (quit). The customer surplus of quitting and adopting an outside option is assumed to be zero. At period t, their utility from purchasing in period $t'(> t)$ would be uncertain under a randomized pricing policy. Customers are assumed to be rational and make decisions to maximize their utilities. Denote by $V_t(p_t)$ a customer's maximum utility given price p_t in the current period t. V_t would be independent of time because we consider a stationary system with an infinite horizon. Thus, without abuse of notation, we omit the subscript and use V throughout the rest of the paper. Therefore, $V(p_t)$ is given by

$$V(p_t) = \max\{\underbrace{v - p_t}_{\text{(buy)}}, \underbrace{-c + \mathsf{E}[V(P)]}_{\text{(wait)}}, \underbrace{0}_{\text{(quit)}}\}.$$

We show in the following lemma that customers' optimal response is governed by a threshold policy. They would purchase the product immediately if the price in the current period is sufficiently low; otherwise, they would either delay the purchase or leave immediately.

Lemma 1. *There exists a threshold* $\underline{p} \in (0, v]$ *such that*

$$V(p_t) = \begin{cases} v - p_t & \text{if } p_t \leq \underline{p}, \\ [-c + \mathsf{E}[V(P)]]^+ & \text{otherwise.} \end{cases} \tag{1}$$

Whether a customer would wait or leave depends on their valuation v. If a customer's valuation v is sufficiently high, then the customer is better off waiting for a favorable price in future periods. Otherwise, they would either buy now or leave immediately. In case of ties between acting immediately and later, an agent is assumed to act immediately. We apply this tie-breaking rule throughout this paper. In particular, a customer, who is otherwise indifferent between quitting or buying immediately and waiting, is assumed to quit or buy immediately. We formalize this result as follows.

Proposition 1 (Threshold Rule of Customer). *A customer with valuation* v *would be willing to wait and eventually purchase a unit of the product (either immediately or in a future period) if* $v > \underline{v}^* = \max\{v' \mid \mathsf{E}[(v' - P)^+] \leq c\}$, *and their purchase threshold* $\underline{p} = \underline{v}^*$. *Otherwise, the customer would purchase the product if* $v \geq p_t$, *or leave immediately if* $v < p_t$.

Proposition 1 shows that a customer's decision, in particular, whether they will wait or leave, depends critically on the firm's choice of the price distribution. This is not surprising in the sense that if a relatively low price can be drawn from the price distribution with a high probability, then a customer is better off waiting because their expected utility from waiting outweighs their waiting cost. Conversely, if a low price is not likely to be drawn from the price distribution, that is, a customer is expected to wait for a long time before seeing a favorable price, then the best option for the customer is not to wait. Proposition 1 implies that the firm's choice of the price distribution shapes customers' decisions in a critical way.

More precisely, valuation $\underline{v}^* = \max\{v' \mid \mathsf{E}[(v' - P)^+] \leq c\}$ is the break-even point. A customer with a valuation exactly at $\underline{v}^*$ would be indifferent between waiting and not waiting. At this break-even point $v = \underline{v}^*$, the expected benefit of waiting for one more period, $\mathsf{E}[(v - P)^+]$, obtaining a surplus $v - P$ when $P \leq v$ is exactly equal to the one-period waiting cost c. A customer whose valuation is strictly higher than $\underline{v}^*$ would be willing to wait: If the customer observes a favorable low price $p_t \leq \underline{v}^*$, they would buy one unit in period t; otherwise, they would wait, because their expected benefit of waiting for one more period is always strictly higher than the waiting cost c. In period $t + 1$, the exactly *same* threshold rule applies because the last period cost of waiting, c, has been *sunk*. As a result, a customer with $v > \underline{v}^*$ would eventually purchase a unit of the product, either immediately upon arrival or in a future period. For those customers whose valuation is strictly lower than $\underline{v}^*$, they would not wait and only make a purchase upon arrival when they see a price lower than their valuation. For those whose valuation is at $\underline{v}^*$, our tie-breaking rule assumes that they conform to the latter.

Other than the price distribution, it is apparent that customers' own traits of valuation and waiting cost also affect their optimal responses. For a given waiting cost c, perhaps somewhat counterintuitively, a customer is more willing to wait if their valuation is higher because the customer's valuation is more likely to exceed the threshold $\underline{v}^*$. On the other hand, more intuitively, with a given valuation v, a customer is less willing to wait with a higher waiting cost.[11] These observations are important to understand the firm's optimal choice of the price distribution in the presence of multiple market segments with heterogeneous customer valuations and waiting costs.

Las, we formally summarize the firm's expected profit from one customer segment with valuation v and per-period waiting cost c in the following lemma.

Lemma 2 (Firm's Expected Profit from a Single Segment). *Consider a single market segment with customer valuation*

v and waiting cost c. Under an arbitrary randomized price distribution P:

(i) *If $v > \underline{v}^*$, the monopolist's expected profit is* $\mathsf{E}[P \mid P \le \underline{v}^*]$.

(ii) *Otherwise, the monopolist's expected profit is* $F(v)\mathsf{E}[P \mid P \le v]$.

For the case when customers are willing to wait (i.e., $v > \underline{v}^*$), they would wait, potentially for many periods, until they buy at a price lower than $\underline{v}^*$. Hence, the monopolist's expected profit is given by $\mathsf{E}[P \mid P \le \underline{v}^*]$. Of course, one underlying assumption of this expression is that the firm does not discount future payoffs, which is consistent with the long-run average objective adopted in a similar setting of Besbes and Lobel (2015). For the case when customers are not willing to wait (i.e., $v \le \underline{v}^*$), they would only buy upon arrival seeing a price no higher than their valuation v and hence, the monopolist's expected profit is $F(v)\mathsf{E}[P \mid P \le v]$. For both cases, the expected profit is no more than v. As an immediate ramification, the optimal randomized pricing policy is trivial with only one segment of customers: it reduces to an optimal static pricing policy with a price at v.

3.2. Optimal Randomized Pricing Policy

Next, we derive the optimal randomized pricing policy with two segments of customers who differ in both valuations and patience levels. The flow of this section is as follows. We first discuss structural properties that an optimal randomized pricing policy must satisfy. Then, based on those properties, we construct an optimal price distribution and prove its optimality by showing that an upper bound of the expected revenue can be attained by such a price distribution.

A benchmark that we will use throughout the rest of this section is the optimal static pricing policy, under which customers behave myopically and the expected profit is given by $\max\{v_L, (1-\alpha)v_H\}$. That is, the monopolist either sets a static price at v_L to sell to both segments or sets a static price at v_H to sell only to the high-valuation customers. When an optimal static pricing policy outperforms randomized pricing policies, the price distribution in an optimal randomized pricing policy reduces to a constant.

To facilitate the following discussion, we determine critical thresholds that govern customers' decisions as follows, based on the single-segment analysis in Section 3.1. That is, a low (respectively, high) valuation customer would either purchase or wait if and only if $\underline{p}^L < v_L$ (respectively, $\underline{p}^H < v_H$), where

$$\underline{p}^i = \max\{v' \mid \mathsf{E}[(v' - P)^+] \le c_i\}, \qquad i = L, H.$$

3.2.1. Structural Properties.

We start by proving some key features that an optimal randomized pricing policy must satisfy. The properties shed light on market characteristics that the firm can potentially benefit from using randomized pricing policies. They also narrow down the candidate policies and set the stage for us to derive the optimal randomized pricing policy.

Lemma 3 (Absolute Patience Level). *An optimal randomized pricing policy reduces to a static pricing policy if $c_L \ge c_H$.*

Waiting cost parameters c_L and c_H indicate the absolute patience levels for low- and high-valuation customers, respectively. Lemma 3 shows that, under any randomized pricing policy, the monopolist cannot obtain a higher profit than that from an optimal static pricing policy when high-valuation customers are more patient. This result is not surprising in the sense that when high-valuation customers are more patient, they can wait for at least as long as low-valuation customers and thus would make a purchase at a price no higher than what low-valuation customers would pay. The highest price that a low-valuation customer would pay is v_L. Because the firm cannot effectively separate the two segments of customers when $c_L \ge c_H$, the optimal expected profit per period can be no higher than v_L, which is the expected profit under the static price v_L.

Lemma 4 (Relative Patience Level). *An optimal randomized pricing policy reduces to a static pricing policy if $v_L/c_L \le v_H/c_H$.*

In Lemma 4, we show a weaker condition, $v_L/c_L \le v_H/c_H$ (i.e., $c_L/c_H \ge v_L/v_H$), under which the firm would not be better off from a randomized pricing policy. Essentially, the ratio of a customer's valuation and their per-period waiting cost, that is, v_i/c_i, $i = L, H$, indicates their relative patience level. It manifests the maximum number of time periods that a customer could potentially wait before their utility becomes negative. Lemma 4 shows that the firm can only effectively price discriminate two segments when low-valuation customers are relatively more patient. Moon et al. (2017) show, in an online appendix, a special case of Lemma 4 when the high-valuation customers have zero waiting cost, that is, $c_H = 0$.

Furthermore, as shown in Lemma OS.1 in the online appendix, $v_L/c_L > v_H/c_H$ is a necessary condition for low-valuation customers to wait, and for high-valuation customers to either purchase or leave immediately under any randomized pricing policy. Intuitively, this is a situation that benefits the firm the most because the firm can charge a higher price for high-valuation customers while at the same time setting a lower price randomly once in a while and clearing the market of hanging around low-valuation customers. This intuition is formally confirmed in the following proposition.[12]

Proposition 2 (Customer Behavior Under an Optimal Randomized Pricing Policy). *Under any optimal randomized policy that outperforms the optimal static pricing policy, low-valuation customers would wait, and high-valuation customers would either purchase or leave immediately upon arrival.*

3.2.2. Optimal Price Distribution. Next, we derive the optimal randomized pricing policy. Based on Proposition 2, we only need to focus on those pricing policies under which low-valuation customers would wait, whereas high-valuation customers would not, that is, $v_L > \underline{p}^L$ and $v_H \le \underline{p}^H$. Recall that $\underline{p}^i = \max\{v' \mid \mathsf{E}[(v' - P)^+] \le c_i\}$, $i = L, H$. Because $\mathsf{E}[(v' - P)^+]$ increases in v', $v_L > \underline{p}^L$ and $v_H \le \underline{p}^H$ are equivalent to $\mathsf{E}[(v_L - P)^+] > c_L$ and $\mathsf{E}[(v_H - P)^+] \le c_H$, respectively. An optimal randomized pricing policy that satisfies the preceding conditions can be derived by solving the following optimization problem, with the objective function derived from our preceding discussion and the constraints based on Proposition 1:

$$\begin{aligned} \max_{F(p)} \quad & \alpha \mathsf{E}[P \mid P \le \underline{p}^L] + (1-\alpha)F(v_H)\mathsf{E}[P \mid P \le v_H] \\ \text{s.t.} \quad & \underline{p}^L = \max\{v' \mid \mathsf{E}[(v' - P)^+] \le c_L\}, \\ & \mathsf{E}[(v_L - P)^+] > c_L, \\ & \mathsf{E}[(v_H - P)^+] \le c_H. \end{aligned} \tag{2}$$

This is a nontrivial problem because we need to optimize a distribution function with no obvious structure to exploit. The solution approach we take here is that we first find an upper bound of the expected profit from Problem (2) and then construct a price distribution achieving the upper bound.

Because $\mathsf{E}[P \mid P \le \underline{p}^L] \le \mathsf{E}[P \mid P \le v_L]$, we can remove $\underline{p}^L$ and obtain a relaxation of the previous optimization problem as follows:

$$\begin{aligned} \max_{F(p)} \quad & \alpha \mathsf{E}[P \mid P \le v_L] + (1-\alpha)F(v_H)\mathsf{E}[P \mid P \le v_H] \\ \text{s.t.} \quad & \mathsf{E}[(v_L - P)^+] > c_L, \\ & \mathsf{E}[(v_H - P)^+] \le c_H. \end{aligned} \tag{3}$$

For any feasible solution of this relaxation, we have $F(v_L) < (c_H - c_L)/(v_H - v_L)$. This is because by the two constraints of Problem (3),

$$\begin{aligned} c_H - c_L &> \mathsf{E}[(v_H - P)^+] - \mathsf{E}[(v_L - P)^+] \\ &= F(v_H)\mathsf{E}[v_H - P \mid P \le v_H] - F(v_L)\mathsf{E}[v_L - P \mid P \le v_L] \\ &\ge F(v_L)\mathsf{E}[v_H - P \mid P \le v_L] - F(v_L)\mathsf{E}[v_L - P \mid P \le v_L] \\ &= (v_H - v_L)F(v_L). \end{aligned}$$

As such, we restrict our discussion to those distributions where this condition is satisfied in the following analysis.

Lemma 5. *Let R_P be the support of any optimal price distribution of Problem (3). If there exists $p \in R_P$ such that $p > v_L$, then $p = v_H$.*

Lemma 5 shows an important structural property of an optimal price distribution for Problem (3). The underlying rationale is as follows. First, it does not make sense to charge a price greater than v_H because neither high-valuation nor low-valuation customers would make a purchase at such a price. Second, for any feasible solution to Problem (3), high-valuation customers would either purchase or leave immediately upon arrival. As such, whenever a price greater than v_L but less than or equal to v_H is drawn from the price distribution, high-valuation customers would always purchase immediately, whereas low-valuation customers would wait because the price is greater than their valuation. Thus, the best strategy for the firm is to charge the highest possible price, that is, v_H, such that all surplus from high-valuation customers can be extracted.

Lemma 6 (Upper Bounds). *Let*

$$\Delta_R = \sqrt{\frac{\alpha c_L}{(1-\alpha)(v_H - v_L)}}.$$

(i) *When $\Delta_R \ge 1$, the optimal expected profit from (3) is no more than that from an optimal static pricing policy;*

(ii) *When $\Delta_R < 1$,*

(a) *If $\Delta_R < (c_H - c_L)/(v_H - v_L)$, the optimal expected profit from (3) is no more than $U_1 \equiv \alpha v_L + (1-\alpha) v_H - (1-\alpha)c_L - 2(1-\alpha)(v_H - v_L)\Delta_R$;*

(b) *Otherwise, the optimal expected profit from (3) is no more than*

$$U_2 \equiv \alpha v_L + (1-\alpha)v_H - (1-\alpha)c_H - \alpha c_L \frac{v_H - v_L}{c_H - c_L}.$$

Based on Lemma 5, we can derive the upper bounds of the expected profit from Problem (3), which is summarized in Lemma 6. It turns out that when $\Delta_R \ge 1$, any randomized pricing policy is dominated by the optimal static pricing policy. It is easy to verify that Δ_R increases in c_L, which implies that randomized pricing policies are more likely to be dominated when low-valuation customers are less patient. The underlying rationale is that the effectiveness of randomized pricing policies relies on intertemporal price discrimination of the two segments of customers by charging high and low prices for high- and low-valuation customers, respectively. However, when low-valuation customers are less patient, they become less likely to delay their purchases, and thus the firm would either lose more low-valuation customers or reduce prices to accommodate both low- and high-valuation customers. In the case of $c_L \ge c_H$, that is, high-valuation customers are more patient than low-valuation customers, randomized pricing policies become completely ineffective, as shown in Lemma 3. Similarly, we notice that Δ_R decreases in $v_H - v_L$. That is, randomized pricing policies are more likely to be dominated when the difference in valuations between two segments becomes smaller. When customers' valuations become closer, high-valuation customers can mimic the behavior of low-valuation customers more easily, making it more costly for the firm to price discriminate between the two segments.

When $\Delta_R < 1$, we show the upper bounds of the optimal expected profit from Problem (3) in Lemma 6(ii).

These upper bounds can be achieved by a two-point price distribution, as we will show shortly. In the first best solution, when the firm knows the private information of each individual customer's valuation, the optimal expected profit is given by $\alpha v_L + (1-\alpha)v_H$, which certainly cannot be attained by anonymous posted price mechanisms due to information asymmetry and customers' self-selection behavior. Lemma 6 shows that, under an optimal randomized pricing policy, the profit loss, compared with the first best solution, depends on the market composition α, per-period waiting costs c_L and c_H for both low- and high-valuation customers, and the difference in valuations across two market segments $v_H - v_L$.

Next, we are ready to discuss optimal randomized pricing policies. Similar to Su (2007), we will focus on ϵ-optimal policies as a consequence of our restriction to anonymous posted price mechanisms. An ϵ-optimal policy is defined as follows.

Definition 1 (ϵ-Optimal Policy). For any arbitrarily small $\epsilon > 0$, a policy is ϵ-optimal if its revenue is greater than $1-\epsilon$ fraction of the optimal revenue.[13]

For brevity, we refer to ϵ-optimal policies as "optimal" policies in the rest of the analysis.

Proposition 3 (Optimal Randomized Pricing Policy). *Suppose*

$$\Delta_R = \sqrt{\frac{\alpha c_L}{(1-\alpha)(v_H - v_L)}} < 1.$$

The following two-point price distribution is optimal:

$$f(p) = \begin{cases} \beta^* & \text{if } p = \underline{p}^*, \\ 1-\beta^* & \text{if } p = v_H, \end{cases}$$

where $\beta^* = \min\{\Delta_R, (c_H - c_L)/(v_H - v_L) - \delta\}$, $\underline{p}^* = v_L - c_L/\beta^* - \eta(\delta)$ *and* $\delta > 0$ *is sufficiently small. Moreover,* $\lim_{\delta \searrow 0} \eta(\delta) = 0$.

Proposition 3 shows that the firm can attain the optimal profit with a simple two-point distribution. Under this two-point price distribution, high-valuation customers would always purchase immediately upon arrival, whereas low valuations would always purchase with the lower price of $\underline{p}^*$. The lower price from the price distribution is less than v_L by an amount of c_L/β^*. This is the "price" the firm has to pay to incentivize low-valuation customers to wait (instead of leaving immediately). The compensation increases in c_L but decreases in β^*. It is intuitive that the firm needs to compensate more with larger c_L due to low-valuation customers being less patient. The reason why the compensation decreases in β^* is explained as follows: β^* is the probability of drawing the lower price from the optimal price distribution. When the lower price is drawn, the market will be cleared, and all customers who have arrived thus far will make a purchase immediately and then leave the market. Therefore, if the lower price is drawn less frequently, then low-valuation customers are expected to wait for a longer time on average, rendering more compensation required to incentivize them to wait.

Proposition 4. *There exists a threshold on c_L, below which the optimal randomized pricing policy outperforms the optimal static pricing policy.*[14]

The optimal static pricing policy is either to set the price at v_H to only serve the high-valuation customers, often referred to as the margin strategy (if $v_L < (1-\alpha)v_H$), or to set the price at v_L to serve both segments of customers, often referred to as the volume strategy (if $v_L \geq (1-\alpha)v_H$). The optimal two-point price distribution would allow the firm to achieve an upper bound as specified in Lemma 6. Whether the optimal expected profit from randomized pricing is greater than that from static pricing depends on system parameters. When low-valuation customers are sufficiently patient, the optimal randomized pricing policy is always better than the optimal static pricing policy, as shown in Proposition 4. In particular, when c_L converges to zero, the upper bound given in Lemma 6 converges to the expected profit of the first best solution, that is, $\alpha v_L + (1-\alpha)v_H$. That is, when low-valuation customers are sufficiently patient, an optimal randomized pricing policy is able to almost *perfectly* price discriminate the two segments through occasional random promotions to clear low-valuation customers, allowing the firm to extract almost *all* surplus from customers.

Corollary 1 (Consumer Surplus and Social Welfare). *Compared with the optimal static pricing policy, the optimal randomized pricing policy yields:*

(i) *If $v_L < (1-\alpha)v_H$, higher social welfare, and higher consumer surplus for both low-valuation and high-valuation customers;*

(ii) *If $v_L \geq (1-\alpha)v_H$, lower social welfare, and higher consumer surplus for low-valuation customers, but lower consumer surplus for high-valuation customers when $c_H < v_H - v_L$.*

(*All these statements hold weakly if the optimal randomized pricing policy reduces to a static price.*)

Corollary 1 illustrates the impact of the optimal randomized pricing policy on social welfare and consumer surplus. Let us first discuss social welfare, which is the total surplus of the firm and all customers. Under the optimal randomized pricing policy, each customer is guaranteed to purchase a product either immediately or in future periods. As high-valuation customers always make a purchase under both the optimal randomized and static pricing policies, the loss of efficiency is solely due to the cost of waiting from low-valuation customers. If $v_L < (1-\alpha)v_H$, low-valuation customers are induced to wait and make a purchase under the optimal randomized pricing policy but are excluded from the market under the optimal static pricing policy (because the optimal static price is v_H), leading to higher social welfare under the optimal randomized pricing policy. Conversely, when $v_L \geq (1-\alpha)v_H$, the social welfare is maximized under the optimal static pricing policy (because the optimal static price is v_L) as there is no efficiency loss. Consequently, the optimal randomized pricing policy would only hurt social welfare in this case, with the efficiency loss from waiting by low-valuation customers.

On another note, low-valuation customers would always weakly benefit from randomized pricing. The underlying rationale is that consumer surplus from low-valuation customers is always zero under the optimal static pricing policy: They would either purchase at a price of v_L or leave immediately. However, under the optimal randomized pricing policy, the lower price from the price distribution, which is less than v_L, incentivizes low-valuation customers to wait rather than leave immediately and compensates for their expected wait c_L/β^*, where $1/\beta^*$ is the expected number of periods to see the low price. In addition, the optimal policy also leaves a tiny amount of positive expected surplus, $\eta(\delta)$, for low-valuation customers. If $v_L < (1-\alpha)v_H$, consumer surplus from high-valuation customers would also be higher under the optimal randomized pricing policy. This is because their surplus is always equal to zero under the optimal static pricing policy with the price of v_H; however, under the randomized pricing policy, some high-valuation customers can purchase at the lower price of $\underline{p}^*$, leading to a positive surplus for them. Consequently, using the optimal randomized pricing policy by the firm could result in higher social welfare as long as $v_L < (1-\alpha)v_H$, because both consumer surplus and the firm's profit can be higher. However, if $v_L \geq (1-\alpha)v_H$, consumer surplus from high-valuation customers could be higher under either the optimal randomized or static pricing policy, depending on the system parameters.

Last, we compare the difference in consumer surplus between low- and high-valuation customers under the two alternative pricing policies. Many researchers have shown that customers would compare with their peers as a reference to evaluate their own payoffs (Ho and Su 2009). As a result, the difference in the surplus from purchasing the product between the two segments of customers can be regarded as a measure of fairness. We show in the following corollary that the optimal randomized pricing policy would make the surplus allocation fairer if $v_L \geq (1-\alpha)v_H$.

Corollary 2 (Fairness). *Compared with the optimal static pricing policy, the discrepancy in consumer surplus between low- and high-valuation customers is smaller under the optimal randomized pricing policy if* $v_L \geq (1-\alpha)v_H$.

The reason that randomized pricing can be fairer is exactly due to the same reason why higher social welfare can be generated. If $v_L \geq (1-\alpha)v_H$, under the optimal static pricing policy of charging v_L, high-valuation customers enjoy a positive surplus by hiding their identity, whereas low-valuation customers have zero surpluses. However, under the optimal randomized pricing policy, better price discrimination reduces the surplus of high-valuation customers while still leaving a tiny amount of surpluses on the table for the low-valuation customers. As a result, the difference between the surpluses of the two segments is reduced. Combining the two corollaries, we see that implementing the optimal randomized pricing, as opposed to the optimal static pricing, either achieves higher social welfare benefiting all or a fairer situation among all customers.

3.3. Optimal Deterministic Pricing: Cyclic Pricing

In this section, we seek to gauge the performance of randomized pricing by comparing it against deterministic pricing, which has been studied extensively in the literature. To that end, we first need to derive the optimal deterministic pricing policy in our setting, which has not been studied before when customers are forward-looking and heterogeneous in both valuations and per-period waiting costs. Under a deterministic pricing policy, the firm needs to determine a sequence of prices $\mathbf{p} = \{p_t\}_{t\in\mathbb{N}}$. The sequence of prices is public information, announced to all customers. For a customer of valuation v and per-period waiting cost c, their utility from buying in period t' is given by $v - p_{t'} - (t'-t)c$ if they arrive in period t. They would thus compare net utilities from buying in all periods $t' \geq t$ and choose to buy in the period where their utility is maximized, conditional on it being nonnegative.

We first show that one does not need to consider beyond cyclic pricing for the optimal deterministic pricing policy. Under a cyclic pricing policy, there exists an integer T such that $p_t = p_{t+T}$ for all t, where T is the cycle length. The firm's long-run average expected profit with a cyclic pricing policy $\mathbf{p}$ is given by

$$\Pi(\mathbf{p}) = \lim_{T'\to\infty} \frac{1}{T'} \sum_{t=1}^{T'} \pi_t(\mathbf{p}),$$

where $\pi_t(\mathbf{p})$ is the profit from customers arriving in period t. We have $\pi_t(\mathbf{p}) = \pi_{t+T}(\mathbf{p})$ due to prices being cyclic, and thus an alternative formulation of the firm's long-run average profit is given by $\Pi(\mathbf{p}) = [\sum_{t=1}^{T} \pi_t(\mathbf{p})]/T$.

Proposition 5 (Optimality of Cyclic Pricing). *Cyclic deterministic pricing is optimal within the strategy space of deterministic policies.*[15]

Proposition 5 significantly narrows down candidates for optimal deterministic policies. Our goal is to find a cyclic pricing policy $\mathbf{p}$ such that $\Pi(\mathbf{p})$ is maximized. The flow of this section is as follows. We first discuss the structural properties of optimal cyclic policies and then construct an optimal cyclic policy based on those structural results.

3.3.1. Structural Properties. First, it is easy to verify that the firm would not charge a price higher than v_H, the highest valuation from customers, under any optimal cyclic pricing policy. That is, if $\mathbf{p} = \{p_t\}_{t\in\mathbb{N}}$ is an optimal cyclic pricing policy, we have $p_t \le v_H$ for any t.

Lemma 7. *An optimal cyclic pricing policy reduces to a static pricing policy if* $c_L \ge c_H$.

As the counterpart of Lemma 3, Lemma 7 illustrates the impact of customers' patience levels on the structure of the optimal cyclic pricing policy. When high-valuation customers are more patient than low-valuation customers, cyclic pricing policies do not perform better than the optimal static pricing policy. This is because high-valuation customers would wait for at least as long as low-valuation customers and pay no more than low-valuation customers, rendering any intertemporal price discrimination mechanism, including cyclic pricing policies, ineffective at all.

Next, we characterize the behavior of high-valuation customers under an optimal cyclic pricing policy in Proposition 6. Both cyclic deterministic and randomized pricing policies are variations of intertemporal price discrimination mechanisms, and thus customers' behavior under the optimal policies is similar. Proposition 6 shows that the optimal cyclic pricing policy will be designed in such a way that high-valuation customers would have no incentive to wait.

Proposition 6. *Under any optimal cyclic pricing policy, high-valuation customers would always purchase immediately upon arrival.*

3.3.2. Optimal Cyclic Deterministic Pricing Policy. Proposition 6 sets the stage for us to derive optimal cyclic pricing policies. For this purpose, we assume $c_L < c_H$ in view of Lemma 7. If this assumption fails, the optimal cyclic pricing policy reduces to a static price. We construct a cyclic pricing policy and prove its optimality in Proposition 7. With such a policy, prices weakly decrease over time within each cycle.

Proposition 7 (Optimal Cyclic Deterministic Pricing Policy). *An optimal cyclic pricing policy is given by*

$$p_{t'} = v_H - [v_H - v_L + (T-1)c_L + \delta - (T-t')c_H]^+,\quad t' = 1, 2 \ldots, T,$$

where T is the cyclic length, and $\delta > 0$ is arbitrarily small. Under this policy, all low-valuation customers buy in period T, and high-valuation customers buy immediately without delay.

Consistent with Conlisk et al. (1984), Proposition 7 shows that the optimal cyclic pricing policy follows a markdown pattern. Within a cycle, the prices between two consecutive periods either stay constant at v_H or drop by a size no more than c_H, which eliminates the incentive for high-valuation customers to wait. More precisely, in general, within a cycle, the prices stay constant at v_H for some time (which can be zero), drop by a size no more than c_H, and then finally drop by a size of exactly c_H to the end-of-cycle price $v_L - (T-1)c_L - \delta$ (see Figure 2(a) for an illustration). Conversely, low-valuation customers would wait patiently in the market upon arrival because their waiting is more than fully compensated. Eventually, the market of low-valuation customers will be cleared at the end of each cycle. There is still a loose part in Proposition 7, which is the determination of the cycle length T. This is a common issue for characterizing the optimal cyclic pricing policy (Conlisk et al. 1984, Besbes and Lobel 2015). We cannot characterize it analytically in general, but we show that it can be computed for a special case (see Proposition 8(i)).

The optimal cyclic pricing policy may reduce to a static pricing policy. The following corollary identifies two sufficient conditions for such a degeneracy.

Corollary 3. *If* (i) $c_H/c_L \le 2/(1-\alpha)$ *or* (ii) $v_H - v_L \le c_L/(1-\alpha)$, *an optimal cyclic pricing policy reduces to a static pricing policy.*

By Lemma 7, if $c_H/c_L \le 1$, the optimal cyclic pricing policy reduces to a static price. Corollary 3(i) further shows a weaker condition for such a degeneracy. That is, when the patience level of high-valuation customers relative to low-valuation customers is below a threshold of $2/(1-\alpha)$, cyclic pricing reduces to static pricing, and that situation is more likely to occur if there is a larger fraction of low-valuation customers. In addition, Corollary 3(ii) shows that if $v_H - v_L \le c_L/(1-\alpha)$, that is, $(1-\alpha)(v_H - v_L) \le c_L$, the optimal cyclic pricing policy reduces to a static price, in which case the benefit of intertemporal price discrimination through having low-valuation customers wait for one period is no more than the waiting cost. Corollary 3 implies that we cannot say cyclic pricing always outperforms static pricing.

3.4. Performance Comparison

Neither can we have the same ranking of the performances of randomized pricing and cyclic pricing for all scenarios. We identify the following sufficient conditions under which one does better than the other.

Proposition 8 (Randomized Pricing vs. Cyclic Deterministic Pricing). *Suppose* $v_H - v_L > c_L/(1-\alpha)$.

(i) *If* $v_H - v_L$ *is below the threshold of* $c_H/\left(1+\left(\sqrt{(1-\alpha)/\alpha}\right)/2\right)$, *the optimal cyclic pricing policy is in*

Figure 2. (Color online) Optimal Cyclic Prices

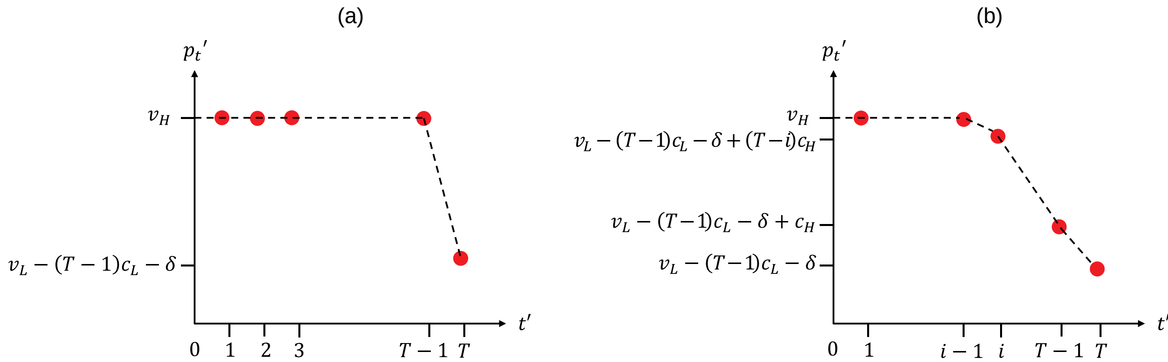

the form of the first $T-1$ periods priced at v_H and the last period priced at $v_L-(T-1)c_L-\delta$ with the optimal cycle length $T=\lfloor T_1\rfloor$ or $\lfloor T_1\rfloor+1$, and yields a higher profit than an optimal randomized pricing policy, where $T_1=\sqrt{(1-\alpha)(v_H-v_L-c_L)/(\alpha c_L)}$ and $\lfloor x\rfloor$ represents the greatest integer that is no more than x.

(ii) *If v_H-v_L is sufficiently large, the optimal cyclic pricing policy is in the general form of a series of markdowns and yields a lower profit than an optimal randomized pricing policy.*

For the rest of this section, we assume $v_H-v_L>c_L/(1-\alpha)$. Otherwise, by Corollary 3(ii), the optimal cyclic pricing policy reduces to a static pricing policy, then we can resort to the comparisons between the optimal randomized and static pricing policies studied in Section 3.2.

Proposition 8(i) provides a sufficient condition under which cyclic pricing outperforms randomized pricing. When the valuation difference between the two segments v_H-v_L is below a threshold, the optimal cyclic pricing policy is in the form of the first $T-1$ periods priced at v_H and the last period priced at $v_L-(T-1)c_L-\delta$ (see Figure 2(a) for an illustration). That is, under the optimal cyclic pricing, the frequency of running a promotion is $1/T$. Recall that under the optimal randomized pricing policy, the frequency of running a promotion is β^*. If we equalize these two frequencies, the revenue earned from the high-valuation customers would be the same. However, the waiting cost compensation under the cyclic deterministic policy is lower than that under the randomized policy because under the cyclic deterministic policy, the wait time for a promotion is strictly less than T periods, but under the randomized policy, that is expected to be $1/\beta^*$ periods. As a result, it takes less for the cyclic deterministic policy to compensate the low-valuation customers and hence retain more revenue than the randomized policy. Therefore, the optimal cyclic pricing policy has a lower promotion frequency, charges a higher promotion price, and earns a higher revenue than the optimal randomized pricing policy. To illustrate this, we consider an example with parameters $v_H=10$, $v_L=5$, $c_H=8$, $c_L=0.3125$, and $\alpha=0.5$. The optimal randomized pricing policy is to run a promotion with probability $\beta=0.25$ and a discount price $\underline{p}^*=3.75$, which achieves a profit of 6.09. On the other hand, the optimal cyclic pricing policy is $\{10,10,10,4.06-\delta\}$ with a cycle of four periods and achieves a profit of almost 6.29.

Proposition 8(ii) shows that randomized pricing performs better than cyclic pricing when the valuation difference between the two segments v_H-v_L is sufficiently large. In this case, the optimal cyclic deterministic policy is in the general form of a series of markdowns as illustrated in Figure 2(b). The underlying rationale can be understood by considering a high-valuation customer who arrives at the second-to-last period under a cyclic deterministic policy. If they waited for a period, they would, for sure, obtain the discounted price that is targeted to the low type. However, this is not the case under randomized pricing. In any period, if a high-valuation customer decided to wait, there would be no guarantee at all they would get a discounted price in the next period. If the discrepancy between v_H and v_L is large enough, it is important for the firm to capture most of the surplus from high-valuation customers. This can be achieved under optimal randomized pricing policies by maintaining a low frequency of offering the discounted price $\underline{p}^*$, and thus the majority of high-valuation customers would make a purchase at the high price v_H upon their arrivals. Given that the low-valuation customers are sufficiently patient, the firm can still capture those customers at the occasional promotions. On the other hand, the optimal cyclic pricing policy follows a predetermined price trajectory known to all customers. Thus, the firm either gives up the low-valuation customers completely by not running promotions at all or adopts a sequence of markdowns by capturing the low-valuation at the end of a cycle. However, neither is ideal. The former loses the low segment compared with the optimal randomized pricing, and the latter needs to compensate high-valuation customers to prevent them from waiting through a series of markdowns. This compensation cuts into the firm's profitability, again leading to a lower profit compared with the optimal randomized pricing. To illustrate this, we consider an example with parameters $v_H = 10$, $v_L = 1$, $c_H = 5$, $c_L = 0.1$, and $\alpha = 0.5$. The optimal randomized pricing policy is to run a promotion with probability $\beta = 0.1054$ and a discount price $\underline{p}^* = 0.05$, which achieves a profit of 4.50. On the other hand, the optimal cyclic pricing policy is $\{10, 10, 10, 10, 10, 10, 10, 10, 5.1 - \delta, 0.1 - \delta\}$ with a cycle of 10 periods and achieves a profit of almost 4.31.

4. Extensions

In this section, we extend the base model in three directions. We first consider the case where a portion of customers are myopic and study the impact of myopic customers on the optimal randomized pricing policy, as well as the firm's expected profit. Then we consider a Markovian pricing policy, where prices are allowed to be intertemporally correlated in a Markovian fashion, and last, consider a general model with multiple customer segments.

4.1. Myopic Customers

In practice, it may be the case that some customers are myopic in the sense that they are not patient or sophisticated enough to wait for promotions in the future. In this section, we extend our base model to account for myopic customers and study how the optimal randomized pricing policy varies with the proportion of myopic customers in the population.

Suppose that $\gamma \in [0, 1]$ fraction of customers are strategic, and the remaining customers are myopic, who will never wait. We assume that whether a customer is myopic is independent of their valuation. The rest of the settings are the same as the base model in Section 3. Thus, $\alpha\gamma$ (respectively, $(1-\alpha)\gamma$) fraction of low-valuation (respectively, high-valuation) customers have a per-period waiting cost c_L (respectively, c_H). On the other hand, $\alpha(1-\gamma)$ (respectively, $(1-\alpha)(1-\gamma)$) fraction of low-valuation (respectively, high-valuation) customers would never wait, that is, their per-period waiting cost is equal to ∞.

It is easy to show that Proposition 2 still holds for strategic customers by following a similar analysis. That is, under any optimal randomized policy that outperforms the optimal static pricing policy, strategic low-valuation customers would wait, and strategic high-valuation customers would either purchase or leave immediately upon arrival. Consequently, we can derive the optimal randomized pricing policy by solving the following optimization problem:

$$\begin{aligned}\max_{F(p)} \quad & \alpha\gamma \mathsf{E}[P \mid P \le \underline{p}^L] + \alpha(1-\gamma)F(v_L)\mathsf{E}[P \mid P \le v_L] \\ & + (1-\alpha)F(v_H)\mathsf{E}[P \mid P \le v_H] \\ \text{s.t.} \quad & \underline{p}^L = \max\{v' \mid \mathsf{E}[(v' - P)^+] \le c_L\}, \\ & \mathsf{E}[(v_L - P)^+] > c_L, \\ & \mathsf{E}[(v_H - P)^+] \le c_H.\end{aligned} \tag{4}$$

The objective function is derived based on Lemma 2, and the constraints are given by Proposition 1. Problem (4) reduces to Problem (2) when $\gamma = 1$. Denote

$$\tilde{\Delta}_R = \sqrt{\frac{\alpha\gamma c_L}{(1-\alpha)(v_H - v_L) - \alpha(1-\gamma)v_L}}.$$

We show in the proof of Proposition 9 that an optimal randomized pricing policy reduces to a static pricing policy when either $\tilde{\Delta}_R \ge 1$, or $\tilde{\Delta}_R$ does not exist, that is, $(1-\alpha)(v_H - v_L) \le \alpha(1-\gamma)v_L$.

Proposition 9 (Optimal Randomized Pricing Policy with Myopic Customers). *If $\tilde{\Delta}_R < 1$, the following two-point price distribution is optimal:*

$$f(p) = \begin{cases} \tilde{\beta}^* & \text{if } p = \underline{\tilde{p}}^*, \\ 1 - \tilde{\beta}^* & \text{if } p = v_H, \end{cases}$$

where $\tilde{\beta}^* = \min\{\tilde{\Delta}_R, (c_H - c_L)/(v_H - v_L) - \delta\}$, $\underline{\tilde{p}}^* = v_L - c_L/\tilde{\beta}^* - \eta(\delta)$ and $\delta > 0$ is sufficiently small. Moreover, $\lim_{\delta \searrow 0} \eta(\delta) = 0$.

In the presence of myopic customers, the optimal randomized pricing policy still follows a two-point distribution that is composed of a regular price equal to v_H and a discounted price that clears the market of strategic low-valuation customers from time to time. As shown in Corollary 4(i), if $v_L \le (1-\alpha)v_H$, it is optimal to run promotions less frequently with a lower discounted price as the proportion of myopic customers in the population increases; otherwise, the firm is better off running promotions more frequently with a higher discounted price. This is because, on the one hand, if $v_L \le (1-\alpha)v_H$; that is, v_H is high relative to v_L, it is imperative to extract surplus from high-valuation customers. The expected price that a high-valuation customer would pay upon arrival is given by $\tilde{\beta}^*\underline{\tilde{p}}^* + (1-\tilde{\beta}^*)v_H$, which decreases in $\tilde{\beta}^*$. As the proportion of myopic customers increases, high-valuation customers become more important, and thus the firm is better off running promotions less frequently, that is, choosing a smaller $\tilde{\beta}^*$. The distribution converges to a constant of v_H as $\gamma \to 0$. On the other hand, if $v_L > (1-\alpha)v_H$, it is important for the firm to capture surplus from low-valuation customers. As the proportion of myopic customers increases, the firm should run promotions more frequently so that it will be able to sell to more myopic low-valuation customers. Overall, the existence of myopic customers is detrimental to the firm's profitability, and the expected profit from an optimal randomized pricing policy decreases as the proportion of myopic customers in the population increases, as shown in Corollary 4(ii).

Corollary 4. *The following statements hold.*

(i) *If $v_L \le (1-\alpha)v_H$, both $\tilde{\beta}^*$ and $\underline{\tilde{p}}^*$ increase in γ; otherwise, they both decrease in γ;*

(ii) *The expected profit of the optimal randomized pricing policy increases in γ.*

4.2. Markovian Pricing

Rather than drawing prices independently from a price distribution in the base randomized pricing policy, we consider a Markovian pricing policy, where random prices over time form a Markov process (Wu et al. 2014, 2020). A Markovian pricing policy gives more control to the firm as the price trajectory is state dependent, and we want to investigate whether this extra layer of control would benefit the firm.

Motivated by the optimal randomized pricing policy characterized in Proposition 3, we restrain our discussion to a Markovian pricing policy with only two values, that is, a regular price of p and a discounted price of $p_d (\leq v_L)$. We denote the probability transition matrix by

$$\mathbf{M} = \begin{pmatrix} q_L & 1 - q_L \\ q_H & 1 - q_H \end{pmatrix},$$

where q_L is the transition probability from p_d to p_d, and q_H is the transition probability from p to p_d. Similar to our discussion in Section 3, under any optimal Markovian pricing policy, we can show that high-valuation customers would always purchase immediately upon arrival, whereas low-valuation customers would make a purchase immediately if the price in the current period is p_d and wait otherwise. Consequently, the optimal regular price is given by $p = v_H$, which is assumed to be the case in the rest of our analysis.

For a high-valuation customer arriving in period t, their utility of buying immediately is given by $v_H - p_t$. Denote by $V(p_t)$ the maximum utility for a low-valuation customer arriving in period t, conditional on the price in the previous period being v_H. Then, $V(p_t)$ is given by

$$V(p_t) = \begin{cases} v_L - p_d & \text{if } p_t = p_d, \\ -c_L + \mathsf{E}[V(p_{t+1})] & \text{if } p_t = v_H. \end{cases} \tag{5}$$

We first study the structural property for an optimal Markovian pricing policy. Lemma 8 characterizes the necessary and sufficient conditions such that high-valuation customers would always purchase immediately upon arrival, whereas low-valuation customers would either purchase immediately or wait. They are also necessary conditions for a Markovian pricing policy to be optimal.

Lemma 8. *High-valuation customers would always purchase immediately upon arrival, whereas low-valuation customers would either purchase immediately or wait if and only if $q_H(v_L - p_d) > c_L$ and $q_H(v_H - p_d) \leq c_H$.*

Based on Lemma 8, we can derive an optimal Markovian pricing policy and compare its performance with other pricing policies.

Proposition 10 (Optimal Markovian Pricing Policy). *Let*

$$\Delta_M = \sqrt{\frac{\alpha c_L}{(1-\alpha)(v_H - v_L - c_L)}}.$$

(i) *If $(1-\alpha)(v_H - v_L) \leq c_L$ (under which either Δ_M is not well defined as a real number or $\Delta_M \geq 1$), an optimal Markovian pricing policy (and an optimal randomized pricing policy as well) reduces to a static pricing policy.*

(ii) *If $1/2 \leq \Delta_M < 1$ and $(c_H - c_L)/(v_H - v_L) > 1$, an optimal Markovian pricing policy is given by $q_L^* = 0$, $q_H^* = 1$, $p_d^* = v_L - c_L - \eta(\delta)$ for an arbitrarily small $\delta > 0$ and $\lim_{\delta \searrow 0} \eta(\delta) = 0$. Moreover, it is a cyclic pricing policy with a cyclic length of two. The expected profit from this optimal Markovian pricing policy is always greater than that from an optimal randomized pricing policy.*

(iii) *Otherwise, an optimal Markovian pricing policy is given by $q_L^* = 0$, $q_H^* = \min\{\Delta_M/(1-\Delta_M), (c_H - c_L)/(v_H - v_L) - \delta\}$, $p_d^* = v_L - c_L/q_H^* - \eta(\delta)$ for an arbitrarily small $\delta > 0$ and $\lim_{\delta \searrow 0} \eta(\delta) = 0$. The expected profit from this optimal Markovian pricing policy is always greater than that from an optimal randomized pricing policy.*

There are a couple of interesting observations from Proposition 10. First, we show in Proposition 10(i) that, when $(1-\alpha)(v_H - v_L) \leq c_L$, both an optimal Markovian pricing policy and an optimal randomized pricing policy as a special case of the Markovian policy reduce to a static pricing policy. However, when $0 \leq \Delta_M < 1$, an optimal Markovian pricing policy is guaranteed to be more profitable than an optimal randomized pricing policy, although it may reduce to a cyclic pricing policy (see Proposition 10(ii)). Second, q_L^* is equal to zero in the optimal Markovian pricing policy. That is, after running a promotion in one period by charging the lower price p_d, the firm would not want to do it again immediately in the following period. This is sensible because p_d essentially serves as a market clearing price, and all customers who have arrived so far would make a purchase and leave immediately with the lower price p_d. Then in the next period, the system restarts, and the firm would be better off charging the higher price to extract surplus from high-valuation customers. Such an extra tune-up of the Markovian pricing policy potentially avoids running promotions consecutively and leaving surpluses on the table for high-valuation customers. It indeed allows the firm to gain higher profit when $\Delta_M < 1$, which requires the low-valuation customers to be sufficiently patient.

Now we compare the key decision variables in an optimal Markovian pricing policy with those in an optimal randomized pricing policy when Markovian pricing dominates randomized pricing. The results are summarized in Corollary 5.

Corollary 5. *If c_L is sufficiently small, $p_d^* \geq p^*$ and $\pi_L^* \geq \beta^*$, where π_L^* is the steady-state probability of p_d^*. Moreover,*

$L_M^* = 1/\pi_L^* \le L_F^* = 1/\beta^*$, where L_M^* and L_F^* are the expected length of time periods between two consecutive discounts under optimal Markovian pricing and randomized pricing, respectively.

By Proposition 10, if c_L is sufficiently small, the optimal Markovian pricing policy dominates the optimal randomized pricing policy, which further dominates the optimal static pricing policy. Interestingly, the optimal Markovian pricing policy would run the promotion more often than the optimal randomized pricing policy, but with a smaller size of the discount. This is because, as explained, the optimal Markovian pricing policy avoids running promotions consecutively, and as a result, the low-valuation customers need to be compensated more for waiting. To do so, the firm can either increase the frequency of promotions but, at the same time, increase the discount price or, alternatively, decrease the frequency of promotions but offer a more appealing discount price. The former is more profitable for the firm, as it reduces the surplus left to the high-valuation customers arriving in the promotional period. Overall, such a Markovian tune-up ensures incentive compatibility for low-valuation customers to wait while reducing the surplus enjoyed by high-valuation customers if they happen to arrive at a promotion period.

Finally, we confirm our intuition that Markovian pricing fares better than randomized pricing because it extracts weakly more surplus from both market segments.

Corollary 6. *If $v_H - v_L > c_L$ and $\Delta_M < 1$, an optimal Markovian pricing policy can extract weakly more profit from low-valuation customers and strictly more profit from high-valuation customers than an optimal randomized pricing policy.*

Under the condition that $v_H - v_L > c_L$ and $\Delta_M < 1$, by Proposition 10, Markovian pricing dominates randomized pricing. It does so by smartly running promotions and avoiding leaving surpluses on the table to the high-valuation customers, as in randomized pricing, while still keeping the low-valuation customers incentive-compatible to wait for a promotion.

4.3. Multiple Customer Segments

Next, we consider an extension of the base model with n customer segments. A fraction α_i of customers are of type i, who value the product at v_i. Without loss of generality, we assume $v_i < v_{i+1}$, for $i = 1, 2, \ldots, n-1$. The one-period waiting cost for type i customers is denoted by c_i. We make the following assumptions in this section.

Assumption (W) (Customers May Wait). *Suppose $0 < c_i < v_i$, $i = 1, 2, \ldots, n$.*

Assumption (O) (Ordered Waiting Costs). *Suppose $c_1 < c_2 < \cdots < c_n$.*

Assumption (W) is an extension of Assumption (S) for multiple customer segments. The model becomes trivial if it is violated because customers would either purchase the product or leave immediately upon arrival. Assumption (O) implies that the waiting cost per period is higher for customers with a higher valuation. This assumption is innocuous in the sense that, as we will show in Lemma 9 and Lemma A.5 in the online appendix, if $c_1 \ge c_2 \ge \cdots \ge c_n$, both an optimal randomized pricing policy and an optimal deterministic pricing policy reduce to a static pricing policy. In other words, if there exists a pair of i and $j(> i)$ such that $c_i \ge c_j$, neither randomized pricing policies nor deterministic pricing policies are able to effectively price discriminate the two segments. Thus, this scenario is as if they belong to the same segment, and our analysis is still applicable with the number of customer segments reduced by one.

4.3.1. Optimal Randomized Pricing Policy.

Denote the power set of the set $\{1, 2, \ldots, n\}$ by $Q(n)$, and, under any randomized pricing policy, denote the set of customer segments that will wait by $\mathcal{W} \in Q(n)$. The rest of the customers will either purchase or leave immediately upon arrival. The firm's problem can be formulated as the following optimization problem:

$$\begin{aligned}
\max_{F(p),\, \mathcal{W}\in Q(n)} \quad & \sum_{i\in\mathcal{W}} \alpha_i \mathsf{E}[P \mid P \le \underline{p}^i] + \sum_{i\notin\mathcal{W}} \alpha_i F(v_i) \mathsf{E}[P \mid P \le v_i] \\
\text{s.t.} \quad & \underline{p}^i = \max\{v' \mid \mathsf{E}[(v' - P)^+] \le c_i\}, \quad \forall i \in \mathcal{W}, \\
& \mathsf{E}[(v_i - P)^+] > c_i, \quad \forall i \in \mathcal{W}, \\
& \mathsf{E}[(v_i - P)^+] \le c_i, \quad \forall i \notin \mathcal{W}. \qquad (6)
\end{aligned}$$

The objective function is derived from our discussion in Section 3.1, and the constraints are due to Proposition 1. Potentially, we may solve Problem (6) by conditioning on $\mathcal{W}$ and evaluate the corresponding profits from optimal price distributions across all possible $\mathcal{W}$. However, this is a nontrivial problem due to (i) conditioning on $\mathcal{W}$, we still need to optimize the price distribution, where there is no obvious structure that we can exploit; and (ii) more critically, the cardinality of the power set $Q(n)$ increases exponentially in n. That is, the number of optimization problems that we need to evaluate increases exponentially in the number of customer segments. Thus, it demands that we start with exploring the structural properties of the optimal randomized pricing policy.

A static pricing policy is a special case of randomized pricing policies, where the price distribution reduces to a constant. The expected profit from an optimal static pricing policy, under which customers behave myopically, is given by $\max_{1\le j\le n}\{\sum_{i=j}^n \alpha_i v_j\}$. We first show a sufficient condition under which an optimal randomized pricing policy degenerates into a static pricing policy.

Lemma 9 (Impact of Patience Level). *If $v_1/c_1 \le v_2/c_2 \le \cdots \le v_n/c_n$, there exists $k \in \{1, 2, \ldots, n\}$ for any randomized pricing policy such that*

(i) *Any customer of a valuation greater than v_k will wait;*

(ii) *Any customer of a valuation less than or equal to v_k will either purchase or leave immediately upon arrival.*

Under the previous condition, an optimal randomized pricing policy reduces to a static pricing policy.

As the counterpart of Lemma 4, Lemma 9 shows that, under the condition $v_i/c_i \le v_j/c_j$ (i.e., $c_i/c_j \ge v_i/v_j$), $\forall i < j$, the firm would not be better off with any price randomization. We make some observations. First, Lemma 9 provides justification for Assumption (O) because it implies that, under $c_1 \ge c_2 \ge \cdots \ge c_n$ that completely flips the order in Assumption (O), an optimal randomized pricing policy reduces to a static pricing policy. That is, under any randomized pricing policy, the monopolist cannot obtain a higher profit than that from an optimal static pricing policy when customers of higher valuations have lower waiting costs per period. Second, as shown in the proof of Lemma 9, $v_i/c_i > v_j/c_j$, $\forall i < j$, is a necessary condition for customers of a lower valuation to wait and customers of a higher valuation to either purchase or leave immediately under any randomized pricing policy. Intuitively, this is a situation that benefits the firm the most because the firm can charge higher prices for customers of higher valuations while at the same time setting lower prices randomly once in a while and clearing the market of hanging-around customers of lower valuations. However, $v_i/c_i > v_j/c_j$, $\forall i < j$, is not a sufficient condition for this to happen. We are able to identify one sufficient condition.

Assumption (R) (Ranked Valuation Decay). *Suppose $v_1 - c_1 \ge v_2 - c_2 \ge \cdots \ge v_n - c_n$.*

Assumption (R) says that though upon arrival, customer valuations are ordered such that type 1 has the lowest valuation, the order is reversed if all wait. That is, type 1 customers have the highest residual valuation after waiting (for one period), so they have more incentive to wait. An analogous assumption is made in Golrezaei et al. (2020) for two customer types (see Figure 1 therein).[16]

Proposition 11 (Optimality of Two-Price Randomized Pricing Policy). *Under the additional* Assumption (R), *an optimal randomized pricing policy follows a two-point price distribution. In particular, under the optimal two-point price distribution:*

(i) *There exists $k_r \in \{1, \ldots, n\}$ such that any customer with a valuation less than v_{k_r} would choose to wait, while any customer with a valuation greater than or equal to v_{k_r} would either purchase or leave immediately upon arrival;*

(ii) *The lower price point is less than v_1;*

(iii) *The higher price point is equal to v_{h_r}, where $h_r = \arg\max_{k_r \le j \le n} \left(\sum_{i=j}^{n} \alpha_i v_j\right)$.*

Proposition 11 is handy when it comes to identifying the optimal randomized pricing policy. First, it shows that under the additional Assumption (R), customer behavior is sorted in the sense that there exists a cutoff k_r in type, where any customer with valuation less than v_{k_r} would choose to wait, whereas any customer with a valuation greater than or equal to v_{k_r} would either purchase or leave immediately upon arrival. Consequently, instead of solving 2^n optimization problems as suggested in Problem (6), we only need to evaluate n optimization problems to identify the optimal price distribution, which represents a significant improvement computationally. Second, for a given k_r, n segments of customers have three types of behavior under the optimal randomized pricing policy that turns out to follow a two-point distribution: Those low-valuation customers from type 1 to type $k_r - 1$ would wait for the discount price (lower than v_1, so that type 1 customers feel worthwhile waiting), those high-valuation customers from type h_r and above would always buy upon arrival, and those customers with types in between only buy at the discount price or leave immediately upon arrival. The latter intermediate case would not exist when there were two customer segments.

4.3.2. Optimal Cyclic Pricing Policy.

As shown in the proof of Proposition 5, the optimal customers' response under any cyclic pricing policy can be summarized as follows.

Corollary 7 (Optimal Customer Response Under Cyclic Pricing). *The optimal customer response can be characterized as follows: There exist $k_c \in \{1, 2, \ldots, n\}$ and $k_0 \in \{1, 2, \ldots, k_c - 1\}$ such that,*

(i) *Any customer of valuation greater than or equal to v_{k_c} will either buy or leave immediately upon arrival;*

(ii) *Any customer of valuation less than v_{k_0} will leave immediately upon arrival;*

(iii) *Furthermore, there exists τ_i for any type i customer, $k_0 \le i < k_c$, such that they leave immediately upon arrival if they arrive before the τ_i^{th} period within a cycle; otherwise, they will wait to buy at the end of a cycle.*

Following a similar approach as the case with two customer segments, we investigate structural properties and then construct an optimal cyclic policy based on them. The optimal cyclic pricing policy is summarized in the following proposition, and results on structural properties are relegated to Online Appendix B.

Proposition 12 (Optimal Cyclic Pricing). *Under the additional* Assumption (R), *an optimal cyclic pricing policy has the following characterizations:*

(i) *For any $i < k_c - 1$, we have $\tau_i \le \tau_{i+1}$. In particular, $\tau_1 = 1$;*

(ii) *Moreover, if* $v_1 \geq (\alpha_2 + \cdots + \alpha_n)v_2 \geq \cdots \geq \alpha_n v_n$, *the optimal pricing schedule is given by*

$$p_T = \min_{1 \leq i < k_c}\{v_i - (T - \tau_i)c_i\} - \delta$$

and

$$p_t = \min\{v_{k_c}, p_T + (T - t)c_{k_c}\},\ t = 1, 2, \ldots, T - 1,$$

where $\delta > 0$ *is arbitrarily small.*

Similar to that of Proposition 11, to derive more insights and structural properties of the optimal cyclic pricing policy, we restrict ourselves with the additional Assumption (R). We make some observations from Proposition 12. First, under the stipulated condition, customers of the lowest valuation would definitely buy at the end of a cycle, whereas customers of a higher valuation tend to wait as it comes closer to the end of a cycle. Intuitively, this is a very favorable situation for the firm due to the following. (1) Customers of the lowest valuation have the lowest per-period waiting cost. Thus, the incentive provided for their waiting is not too costly for the firm. (2) The firm is able to extract surplus from a large fraction of customers of lower valuations at the end of a cycle.

Moreover, we are able to characterize the optimal pricing schedule if we further assume that $v_1 \geq (\alpha_2 + \cdots + \alpha_n)v_2 \geq \cdots \geq \alpha_n v_n$. In this case, the optimal static price is given by v_1, which is often referred to as volume strategy. Under the optimal cyclic pricing policy, the price level across two consecutive periods either stays constant at v_{k_c} or drops by a size no more than c_{k_c}, which eliminates the incentive for customers of a valuation greater than or equal to v_{k_c} to wait. More precisely, in general, within a cycle, the prices stay constant at v_{k_c} for some time (which can be zero), drop by a size no more than c_{k_c} and then finally drop by a size of exactly c_{k_c} to the end-of-cycle price $\min_{1 \leq i < k_c}\{v_i - (T - \tau_i)c_i\} - \delta$. Conversely, customers of a valuation less than v_{k_c} would wait patiently in the market upon arrival (conditional on being sufficiently close to the end of a cycle) because their waiting is more than fully compensated. Eventually, the market of customers of lower valuations will be cleared at the end of a cycle.

4.3.3. Randomized Pricing vs. Cyclic Pricing.

Performance comparison between optimal randomized pricing policies and optimal cyclic pricing policies with n customer segments is summarized as follows.

Proposition 13 (Randomized Pricing vs. Cyclic Pricing). *Under the additional* Assumption (R) *and* $v_1 \geq (\alpha_2 + \cdots + \alpha_n)v_2 \geq \cdots \geq \alpha_n v_n$, *we have*

(i) *The problem with n customer segments is equivalent to the problem with two customer segments in which the* $\sum_{j=1}^{k_r-1}\alpha_j$ *fraction of customers have valuation* v_{k_r-1} *and per-period waiting cost* c_{k_r-1}, *and the rest fraction of customers have valuation* v_{k_r} *and per-period waiting cost* c_{k_r}.

(ii) *If* $v_{k_r} - v_{k_r-1} > c_{k_r-1}$ *and* $v_{k_r} - v_{k_r-1} < c_{k_r}/\left(1 + \frac{1}{2}\sqrt{\sum_{j=k_r}^{n}\alpha_j / \sum_{j=1}^{k_r-1}\alpha_j}\right)$, *an optimal cyclic pricing policy yields a higher profit than an optimal randomized pricing policy.*

(iii) *There exists a threshold on* $v_{k_c} - v_{k_c-1}$, *above which an optimal cyclic pricing policy yields a lower profit than an optimal randomized pricing policy.*

As the counterpart of Proposition 8, Proposition 13 shows that our conclusion from the two-segment case remains in the general case with n customer segments under certain conditions. In particular, Proposition 13 is derived under the extra Assumption (R) and the condition $v_1 \geq (\alpha_2 + \cdots + \alpha_n)v_2 \geq \cdots \geq \alpha_n v_n$. The former condition ensures that an optimal randomized pricing policy follows a two-point price distribution, with the lower price being less than v_1 and the higher price equal to $v_{h_r} = v_{k_r}$, with the identity ensured by the latter condition. We supplement this theoretical result with comprehensive numerical studies in Online Appendix C for additional insights.

5. Conclusion

In this paper, we consider randomized pricing as an alternative dynamic pricing strategy. We show that, compared with the optimal static pricing policy, the optimal randomized pricing policy always benefits the firm in terms of profitability when customers of the lowest valuation are sufficiently patient. We further gauge the performance of an optimal randomized pricing policy by comparing it against cyclic pricing policies, which arise to be optimal among deterministic policies. In this case, neither policy always dominates the other. However, when the valuation differences among customers are sufficiently large, an optimal randomized pricing policy yields higher profit than cyclic pricing policies. We expect that these results will continue to hold when customers discount utilities for future periods. Furthermore, customers' patience levels can differ significantly among different types of products, such as hedonic versus utilitarian products. Our model provides useful guidance for firms regarding whether randomized pricing can be beneficial, evaluated at the estimated parameters, no matter which type of product it is.

We model randomized pricing as a price distribution that a firm commits itself to, and a price is randomly drawn from this distribution in each time period. This is, arguably, a stylized model for randomized pricing, which allows us to analyze the problem analytically. However, a randomized pricing policy may manifest itself in different ways in practice. The key essence of randomized pricing, compared with deterministic pricing, is that prices shall appear somewhat randomly from the perspective of customers, which hinders their ability to predict prices accurately. For instance, firms might well adopt deterministic pricing algorithms that depend on factors, such as the marginal cost of procuring/selling a product, which are not directly observable to customers. Because of a random disruption at the supply source or a fluctuation in the inventory level, prices generated by the pricing algorithms may vary accordingly, which would seem random in the eye of consumers, as they have no visibility of these price-governing states and thus are not able to predict the future promotions. Our results show that the firm may benefit from this "randomization" without deliberation, especially when the variation in customer valuations is significant.

Acknowledgments

The authors thank Huseyin Topaloglu (department editor), the associate editor, and two anonymous reviewers for valuable comments and suggestions on the manuscript throughout the whole review process.

Endnotes

1 In an online appendix, Moon et al. (2017) also numerically show when randomized pricing outperforms a deterministic markdown schedule.

2 In an online appendix, Moon et al. (2017) analytically show that deterministic markdowns are preferred over the optimal randomized price policy in the extreme case when the waiting cost of high valuation customers is zero.

3 A cyclic pricing policy is commonly observed in practice, even for durable goods. Take games for Nintendo Switch as an example. Other than the usual seasonal promotions (such as discounts on Mario games on March 10th to celebrate "Mario Day" and promotions of horror games around Halloween), digital games would go on sale at Nintendo's digital store each Thursday, so patient gamers will be able to purchase games at a discounted price (source: https://bucketlist.games/do-nintendo-switch-games-ever-go-on-sale-how-to-find-lowest-prices/).

4 We allow customers to have heterogeneous waiting costs per period, whereas Conlisk et al. (1984) assume the same discounting factor for all customers.

5 There are several major differences between our paper and Wu et al. (2020). On the demand side, Wu et al. (2020) assume that customers have an exponentially distributed lifetime with price monitoring costs, whereas in our model, there is neither any price monitoring cost nor a shopping window, but instead, customers incur waiting costs. On the supply side, Wu et al. (2020) assume a menu cost for each price change, whereas we do not have such a cost. Wu et al. (2020) focus on price guarantees under Markovian pricing, whereas we compare randomized versus cyclic pricing.

6 The state in Moon et al. (2017) is time independent.

7 Moon et al. (2017, appendix F) compare randomized pricing with a deterministic price schedule that may be analogous to cyclic pricing in the case of no inventory constraint. For a special case, it is shown that the deterministic price schedule dominates randomized pricing. We identify general conditions under which randomized pricing dominates cyclic pricing.

8 Under the assumption of perfect commitment, a mechanism design problem reduces to an optimization problem subject to the well-known incentive compatibility constraints because of the revelation principle. However, the revelation principle generally fails when the mechanism designer cannot credibly commit to an outcome induced by the mechanism.

9 An alternative way of modeling utility loss because of customer waiting is to discount future payoffs. Charging waiting cost proportional to the wait time is more common in the operations literature. Our main results shall hold qualitatively for the discounting scheme under some conditions.

10 An extension with multiple segments of customers is studied in Section 4.3.

11 Customers' optimal response under randomized pricing policies closely resembles the searching behavior as commonly observed in the search literature (Lippman and Mccall 1976, Weitzman 1979). In a typical search model, an agent pays a cost to sample a new item with a random reward, whereas analogously, in our model, a customer is able to sample a new random price by waiting for one more time period (and pay a per-period waiting cost c). As a result, customers' optimal response, as characterized in Proposition 1, can be viewed as a variation of the well-known reservation price policy. That is, $\underline{v}^*$ is the reservation price for a customer with valuation v and per-period waiting cost c. The customer will stop waiting and purchase immediately if the price in the time period is no greater than v^* (should the customer decide to wait in the first place).

12 We present the result here to facilitate the flow; however, the proof of this result relies on many of the following analyses, and we thus suggest readers to postpone the reading of the proof until the end of this section.

13 More specifically, for any arbitrarily small $\epsilon > 0$, there exists a sufficient small $\delta > 0$ such that the revenue of a policy that depends on δ is greater than $1 - \epsilon$ fraction of the optimal revenue. Because we assume that in case of ties between acting immediately and later, an agent is assumed to act immediately, by offering a sufficiently small δ off a price, we could induce one segment of customers to wait rather than to buy immediately so to achieve intertemporal price discrimination. This is why we achieve the ϵ-optimality because we leave a small amount of surpluses to some customers. But with δ made infinitesimally small, we almost achieve the full optimality. The ϵ-optimality does not affect our comparison between two ϵ-optimal policies because the ϵ values for both policies can be made sufficiently small, so the comparison always holds for sure.

14 We obtain a more general set of sufficient conditions in the proof of Proposition 4.

15 We prove the result under the general case with n customer segments.

16 In particular, if customers of type i have a discount factor $\gamma_i \in (0, 1]$, Assumption (R) would be equivalent to $\gamma_1 v_1 \geq \gamma_2 v_2 \geq \cdots \geq \gamma_n v_n$.

References

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.

Baye MR, Morgan J (2001) Information gatekeepers on the Internet and the competitiveness of homogeneous product markets. *Amer. Econom. Rev.* 91(3):454–474.

Besanko D, Winston WL (1990) Optimal price skimming by a monopolist facing rational consumers. *Management Sci.* 36(5):555–567.

Besbes O, Lobel I (2015) Intertemporal price discrimination: Structure and computation of optimal policies. *Management Sci.* 61(1):92–110.

Bester H, Strausz R (2001) Contracting with imperfect commitment and the revelation principle: The single agent case. *Econometrica* 69(4):1077–1098.

Chen Y, Farias VF, Trichakis NK (2019) On the efficacy of static prices for revenue management in the face of strategic customers. *Management Sci.* 65(12):5535–5555.

Chen Y-J, Chu LY (2020) Synchronizing pricing and replenishment to serve forward-looking customers. *Naval Res. Logist.* 67(5):321–333.

Conlisk J, Gerstner E, Sobel J (1984) Cyclic pricing by a durable goods monopolist. *Quart. J. Econom.* 99(3):489–505.

Golrezaei N, Nazerzadeh H, Randhawa R (2020) Dynamic pricing for heterogeneous time-sensitive customers. *Manufacturing Service Oper. Management* 22(3):562–581.

Ho TH, Su X (2009) Peer-induced fairness in games. *Amer. Econom. Rev.* 99(5):2022–2049.

Lippman SA, Mccall JJ (1976) The economics of job search: A survey. *Econom. Inquiry* 14(2):155–189.

Liu Q, van Ryzin GJ (2008) Strategic capacity rationing to induce early purchases. *Management Sci.* 54(6):1115–1131.

Liu Q, Zhang D (2013) Dynamic pricing competition with strategic customers under vertical product differentiation. *Management Sci.* 59(1):84–101.

Moon K, Bimpikis K, Mendelson H (2017) Randomized markdowns and online monitoring. *Management Sci.* 64(3):1271–1290.

Reinganum JF (1979) A simple model of equilibrium price dispersion. *J. Political Econom.* 87(4):851–858.

Salop S (1977) The noisy monopolist: Imperfect information, price dispersion and price discrimination. *Rev. Econom. Stud.* 44(3):393–406.

Salop S, Stiglitz J (1977) Bargains and rip-offs: A model of monopolistically competitive price dispersion. *Rev. Econom. Stud.* 44(3):493–510.

Sinitsyn M (2017) Pricing with prescheduled sales. *Marketing Sci.* 36(6):999–1014.

Sobel J (1984) The timing of sales. *Rev. Econom. Stud.* 51(3):353–368.

Su X (2007) Intertemporal pricing with strategic customer behavior. *Management Sci.* 53(5):726–741.

Varian H (1980) A model of sales. *Amer. Econom. Rev.* 70(4):651–659.

Weitzman ML (1979) Optimal search for the best alternative. *Econometrica* 47(3):641–654.

Wu J, Li L, Xu L (2014) A randomized pricing decision support system in electronic commerce. *Decision Support Systems* 58(2):43–52.

Wu J, Zhang D, Liu Y (2020) Sales and price guarantees under Markovian pricing. Preprint, submitted July 20, https://dx.doi.org/10.2139/ssrn.3637161.

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.