Deterministic Intertemporal Price Discrimination

Daniel Walton*

November 14, 2020

Abstract

I study the problem of a durable-goods monopolist with commitment who sells to strategic, dynamically arriving buyers. Since buyers are forward-looking, it may seem that the seller should randomize in order to not fully reveal the timing of price reductions. However, when buyer types are single-dimensional and buyer utilities are all risk-neutral, or linear, in price, the seller optimally chooses a deterministic price path, so that buyers know future prices with certainty. This is in contrast to stochastic results for related mechanism design problems and cases when buyers are risk-averse, and provides a justification for pure-strategy pricing used in many models of intertemporal price discrimination. I consider two applications of the framework, one where consumers are all present from the beginning of the market, and one where consumers arrive at a constant rate over time. In the first model, prices decline smoothly and the seller can perfectly separate types. In the second case, prices exhibit smooth and volatile phases, and are cyclical, with regularly-timed price reductions. Buyer types are not perfectly separated in this case, leading the seller to sometimes change price discontinuously.

1 Introduction and Overview

Many firms which sell goods change their prices over time. A common pattern observed among many price patterns are periodic or regularly timed price reductions, or sales. Often these sales in retail settings are widely publicized to potential customers and occur at traditional times of year, such as Black Friday shopping or Boxing Day (see Nakamura and Steinsson (2008), Pesendorfer (2002), and Klenow and Mailin (2010), for instance), that any strategic buyer will instinctively wait

*Stanford University. Email: dwalton@stanford.edu. I would like to thank my advisor Gabriel Carroll for a great deal of advice and many helpful conversations and ideas, as well as Matt Jackson, Ilya Segal, John Shoven, and Evan Storms for helpful comments and ideas. Additionally, I appreciate the feedback and participation from seminar audiences at Stanford and BYU.

to buy a good around the time of the sale because they anticipate a large price reduction. We are not the first to point out this phenomenon by any means; indeed, many theoretical and empirical studies document and give possible explanations for this fact, dating back to Stokey (1979), Varian (1980), Stokey (1981), Bulow (1982), and Conlisk et al. (1984). This paper presents a general framework under which a monopolist seller will choose a path of prices that is totally known to strategic buyers. In other words, there is no randomness in the level or timing of price changes from the point of view of the forward-looking buyers. Since these prices are nonrandom, we call such a pricing strategy *deterministic*. The framework is a closely related to Board and Skrzypacz (2016), confirming their result of a deterministic pricing strategy as well.

More specifically, we find that when buyers' utilities are linear, or risk-neutral, with respect to price, and buyer types are single-dimensional, in the sense that high-types buy when low-types buy and low-types don't buy when high-types don't buy, the seller finds it optimal to choose a deterministic pricing strategy. This framework encompasses many existing models of intertemporal price discrimination, such as Stokey (1979), Conlisk et al. (1984), Sobel (1991), Landsberger and Meilijson (1985), Hendel and Nevo (2013), Su (2007), Board (2008), White (2004), and Besanko and Winston (1990). In most cases, these models assume the seller chooses a deterministic pricing strategy for tractability reasons. Thus, the result given in this paper can be seen as a justification for making the simplification to deterministic pricing strategies.

At first glance, it may seem obvious that the seller can without loss of profits restrict their attention to deterministic prices when buyer utility is linear in prices, since buyers will only care about the expected value of future prices and not the variance of prices or any other higher moments. However, when types of the buyers are multidimensional, the seller may still find it optimal to use a mixed strategy–randomness over future prices, in order to best satisfy incentive constraints of buyers to whom the seller wishes to sell to at the current date. Furthermore, from the mechanism design literature, we observe that it is often the case that profit-maximizing mechanisms will offer lotteries, or randomized outcomes, as a type separation tool, when there multidimensional types, as in Hart and Reny (2015) and Haghpanah and Hartline (2016). Since a buyer's time of arrival to the market affects their payoffs as well as the type, the results about randomized mechanisms suggest that even in the case of single-dimensional buyers the seller could potentially benefit from non-deterministic prices. In Section 4.4, we give an example of a more general mechanism than posted prices under which the seller chooses randomized prices. This demonstrates a real restriction that is imposed on the seller when they are restricted to offering a single price to all buyers in each period. This means that the seller can not enact price discrimination *within periods* of time since all buyers have access to the good at the same price in the same period, but the seller can enact price discrimination *across periods* of time, by getting high-types to purchase early and low-types to purchase late. Time-discounting is a key element in both the seller and buyers' payoffs, and it implies that any delayed transactions destroys surplus for both the buyer and seller. Consequently, the seller prefers price discrimination within periods to across periods because this avoids the loss of surplus from time-discounting.

In the setting where buyers arrive over time, intuitively a deterministic price may not seem optimal from the seller's point of view, for the reason that types with high willingness-to-pay may arrive close to the timing of a price reduction and decide to wait to purchase at the low price. Instead, if the timing or level of the future price reduction were not certain to the high buyer, they may instead choose to purchase when they arrive to the market rather than wait. Randomness in prices thus creates more equality in incentives of buyers who arrive to the market at different times. This is the premise of the random price path used in Dilme and Garrett (2018), which has the property that past prices contain no information about future price reductions. They consider the scenario with risk-averse buyers, for which the seller has the additional benefit that randomizing prices is a deterrent for buyers to wait. However, the stationarity of the incentive structure is useful even in the risk-neutral setting, and indeed the random price path in their setting remains optimal under their specification with risk-neutral buyers (utility linear in price), and achieves the same profits as a deterministic one. This was further motivation for asking whether or not deterministic prices are optimal in the more general framework. There may be exogenous sources of randomness that induce prices to be random, as in Bitran and Caldentey (2003), Garrett (2016), Deb (2014), and Chen and Farias (2018). In reality, there are surely random sources of variation not controlled by the monopolist; this paper removes those sources to see if randomness is used purely for strategic reasons.

In the second part of the paper, we demonstrate the application of the optimality of deterministic prices to two models: first, the single-arrival model, where buyers are present at the beginning of the market and there are no subsequent arrivals, and the constant-arrivals model, where buyers arrive continuously to the market over time at a constant rate. The single-arrival model is closely related to Stokey (1979), with a key difference being varied discount rates among buyers and seller, making intertemporal price discrimination profitable in many cases, in contrast to Stokey's result. The constant-arrival model is most similar to Conlisk et al. (1984), Board (2008), and Su (2007). The constant-arrivals model extends these beyond two types, demonstrating the complexity of optimal prices possible with dynamic arrivals.

These two applications differ in a significant qualitative manner. The different price behavior is contrasted in Figure 1. The single-arrival model gives rise to smoothly decreasing prices, with the

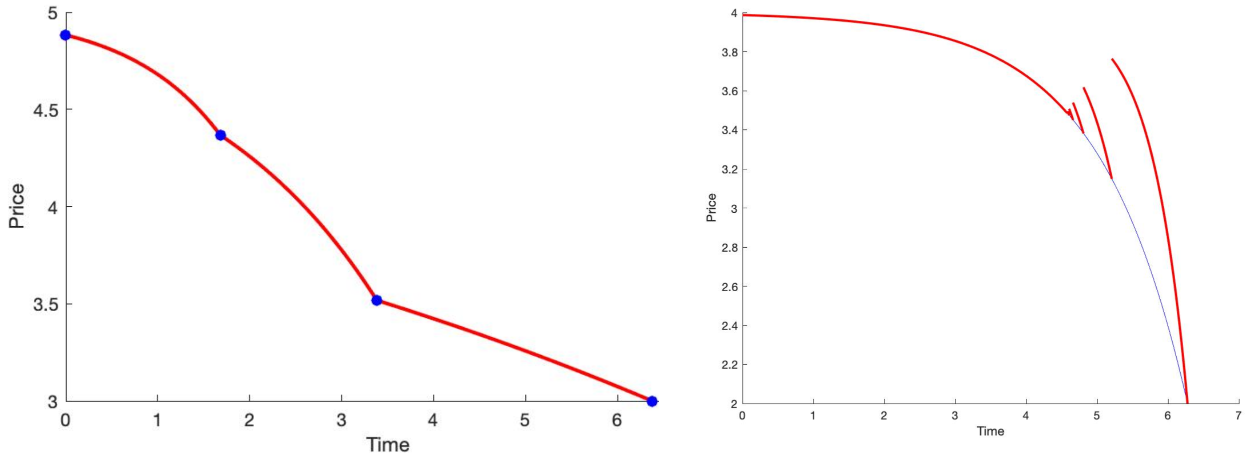

Figure 1: Left panel: an example optimal price path for the single-arrival model. Blue points represent times when a new buyer type purchases. Right panel: an example optimal price path for the constant-arrivals model. Close to the low-price time, prices experience jumps, deviating a large amount from the patient reserve prices.

buyer types cleanly separated by time of purchase. Intertemporal price discrimination is optimal as long as the market share of the lowest type is not too large. When the market share of the lowest type is large, the seller charges a price equal to the valuation of the lowest type and all types purchase immediately. The single-arrival model can also be interpreted as a dynamic arrivals model where all types enter the market at discrete moments that are far apart in time. In this case, prices are cyclical and follow the optimal single-arrival price path, repeated whenever a new set of buyers arrive to the market.

The constant-arrivals model is simple to specify: we consider buyers that vary along the dimension of their discount rates. The model is computationally tractable, though it gives rise to a somewhat complex seller-optimal pricing strategy. Buyer types are no longer separated by time of purchase at all moments in time; however, intertemporal price discrimination still occurs–there is some separation of types. In this model, prices are cyclical, having a regular period. The lowest-valuation type faces intertemporal price discrimination and buys at the end of a price cycle. Optimal price paths are smooth for an initial section of the price cycle, and then experience short intervals of price spikes before dropping to the lowest price of the cycle, after which prices reset to the beginning of an identical price cycle.

There are many adjacent literatures on dynamic pricing and durable goods monopolists. Coase (1972) points out that a seller without commitment power cannot credibly enact intertemporal price discrimination, and Bulow (1982), Gul et al. (1986), Kahn (1986), among others, examine the robustness of that result. There are other models that explain dynamic pricing not using fully strategic buyer, such as Varian (1980), Lazear (1986), and revenue management models that tend to assume exogenously determined demand, such as Gallego and van Ryzin (1994), McAfee and the Velde (2005), and Bitran and Caldentey (2003).

The organization of the body of the paper is as follows. In Section 2, we describe the basic general framework for a finite-horizon, discrete-time setting and the buyers' and seller's problems. Section 3 analyzes the profit-maximization problem of the seller and walks through the optimality of deterministic prices when buyer types are single dimensional with utility that is linear in prices. Section 4 concerns the first application of the deterministic pricing result to the single-arrival model, solving the seller's problem, examining comparative statics and welfare implications, and the mechanism design example briefed above. Section 5 describes and analyzes the constant-arrivals model with deterministic prices, also covering comparative statics and welfare implications. Section 6 concludes.

2 Model Description

In this section, I present a model of that defines the seller-optimal dynamic monopoly pricing problem in the presence of strategic forward-looking buyers with varying types.

Timing. Time is discrete and there is a finite time T after which the good can no longer be sold, thus we index time as $t = 0, 1, 2, \dots, T$. Within a period, the timing of the game is as follows:

(1) buyers of all types enter the market exogenously in a manner that will be described below.

(2) the price p_t for the period is set.

(3) all buyers observe the price p_t for the period.

(4) buyers make the choice to buy one unit of the good in time t at price p_t or to wait.

(5) buyers that choose to purchase get the good and leave the market.

The game then proceeds to the next period.

Price paths. The seller is restricted to posting a single price in each period. More specifically, the seller chooses a price path, which is a discrete stochastic process $\{p_t\}_{t=0}^T$ taking values in $\mathbb{R}$, adapted to some filtration $\mathcal{F}$. We drop the braces when discussing the stochastic process, since it should be clear from context if p_t is a single random variable or the whole process. The filtration $\mathcal{F}$ may be the filtration[1] generated by p_t but it may also be larger by some additional variation

[1]Recall that the filtration generated by p_t is the sequence of σ-algebras $F_t = \sigma(p_0, \dots, p_t)$, the σ-algebra generated by the history of p_t up to t.

introduced by the seller. We assume that the seller has the power to commit to a price path at time 0, which means that the seller chooses a fixed price path at the beginning of the game and does not alter it thereafter. In the discrete finite-horizon setting, we can think of a price path as a tuple of random variables $p = (p_0, p_1, \dots, p_T)$. The filtration is generated by the history of prices that will be commonly known to both the seller and all buyers. To make notation more compact, we will denote expectations conditional on the information generated by prices up to time t as $\mathbb{E}_t[\cdot] \equiv \mathbb{E}[\cdot|p_0, \dots, p_t]$.

Buyer utility. Each buyer has a type that is parameterized by $\theta \in \Theta \subset \mathbb{R}^M$ for some $M \geq 1$, Θ compact. The set of all types Θ will be further specified in both the proof of deterministic pricing as well as the analysis of the single-arrival and constant-arrival models, where it will be a finite set. For the purposes of the exposition, we write a summation over types, but if Θ is an interval, this would be substituted with an integral. A buyer's type is constant over time, and they have unit demand for the good. The set of possible valuations Θ of buyers entering the market is also constant over time, without loss of generality, since buyers of a given type not entering the market in a period can be assigned probability 0 of entering the market at that time. In this model, buyers are infinitesimal, so that a set of buyers of measure 0 do not influence the seller's profit with their purchasing decisions. A buyer's utility from purchasing the good at time t for price $p \in \mathbb{R}$ is a given function $u_t(p, \theta)$. We assume that u is strictly decreasing in p for any t, θ such that u is invertible in p for any utility in $\mathbb{R}$. The buyer's utility from not purchasing the good in any period is 0 and if they leave the market without purchase they get utility of 0, which can be thought of as an outside option with a normalized value.

Buyer stopping problem. We assume that all buyers have access to the same information about prices at all times, regardless of their time of arrival to the market. Each buyer i who is in the market at time t makes a purchasing decision that is a stopping time $\tau \geq t$, where the stopping time is defined relative to the filtration generated by prices. The buyer's optimal stopping problem at time t is to find a stopping time τ^* such that

$$\mathbb{E}_t[u_{\tau^*}(p_{\tau^*}, \theta)] = \sup_{\tau \geq t} \mathbb{E}_t[u_\tau(p_\tau, \theta)]$$

We then define the binary variable denoting type θ_i's purchase at time t: $b_t(\theta_i) = 1$ if $\tau^* = t$ is the unique solution to the optimal stopping problem. If the buyer is indifferent between stopping at t and some other stopping time, we allow the seller to break indifference arbitrarily and choose $b_t(\theta_i)$ to be 1 or 0. Otherwise, $b_t(\theta_i) = 0$. This assumption of full information independent of arrival time together with the independence of utility from arrival time, implies that all buyers of a given type will have identical stopping times if they are in the market at the same time. The stopping times clearly depend on the current price and the distribution of future prices, making the buyer's stopping problem difficult to analyze directly for arbitrary prices.

Demand Dynamics. An important feature of this model is the entry and exit of buyers in the market over time. In each period, buyers of type θ enter the market according to an exogenously specified time-dependent mass $\lambda_t(\theta) \geq 0$. To keep track of the mass of buyers of a given type that are in the market in any given period we denote the mass at time t as m_t, we have the transition equations (suppressing dependence on θ)

$$m_{t+1} = m_t(1 - b_t) + \lambda_{t+1}.$$

The most important part of the analysis for the seller regarding buyer behavior is the relationship between price and purchasing decisions. Let M_t be the total volume of purchases at time t, which is characterized as

$$M_t = \sum_{\theta \in \Theta} m_t(\theta) b_t(\theta).$$

Seller's Profit Objective. The seller faces the problem of maximizing profits. We assume that the marginal cost of producing the good is constant and normalized to 0. The seller also discounts time by $r > 0$, so that any profits accrued in time t are discounted to time 0 by e^{-rt}. Since cost is 0, the profits at any time are equal to sales volume times the price, which is $M_t p_t$ at time t. The sum of discounted expected profits is then

$$\Pi(p) = \sum_{t=0}^{T} e^{-rt} \mathbb{E}\left[M_t(p_t) p_t\right]. \tag{1}$$

The seller's objective is to maximize Π with respect to p, such that $p \in \Delta(\mathbb{R}_+^T)$, where $\Delta(\mathbb{R}_+^T)$ is the space of Borel-measurable distributions over $\mathbb{R}_+^T$. Any price path that maximizes (1) is called a *seller-optimal price path*. With the problem defined, our next task is to solve the seller's objective for the optimal price path.

Continuation Values. In order to solve for the optimal stopping times for each buyer type, we must first define the buyer's stopping problem as a function of the price path. Given a fixed price path p, we can solve the buyer's stopping problems by the backwards induction dynamic programming approach, due to the finite horizon of the game. We define the buyer's continuation value under a fixed price path $\{p_t\}_{t=0}^T$ at time t recursively as

$$V_t(\theta) = \max\{u_t(p_t, \theta), \mathbb{E}_t[V_{t+1}(\theta)]\}$$
$$V_{T+1}(\theta) = 0.$$

The continuation value are random variables, but are determined given the price history $p_0, \dots, p_t$. Intuitively, the continuation value recursive relationship captures the tradeoff that the buyer of type θ makes by either purchasing today or waiting until next period to revisit the purchasing decision.

3 Optimality of Deterministic Price Paths

The first step in analyzing the problem of the seller is to make a transformation of the choice set from prices to types, similar to Board (2008) and Board and Skrzypacz (2016). We note that, given a price path p, we have induced a distribution over which types θ buy in each period. While it is difficult to characterize this mapping, since it depends on the solution to the optimal stopping problems for each type at each period of time, it does suggest the possibility that we may start with a stochastic process over types and derive a price path p that is *consistent* with the distribution over types that buy in each period.

When types are single-dimensional it is straightforward to define a stochastic process over types that buy, because types are totally ordered in terms of their buying choices. For instance, if types are valuations or willingness-to-pay for the good, then if a buyer with a low valuation finds it optimal to buy the good in a given period at a given price, then the buyer with a high valuation must also find it optimal to buy the good in the same period at the same price. Given this intuition, we give the following definitions.

Definition 1. The ordering $\succeq$ on the set Θ is defined for $\theta, \theta' \in \Theta$ by

$$\theta \succeq \theta' \iff b_t(\theta, p) \geq b_t(\theta, p)$$

for all $p \in \Delta(\mathbb{R}_+^{T+1})$ and times $t \in [T]$. The partially ordered type set Θ is *single-dimensional* if the ordering $\succeq$ is linear[2].

While one must verify that $\succeq$ as defined indeed constitutes an ordering, the proof is straightforward and contained in the appendix for completeness. Single-dimensionality along with compactness of Θ implies that there is a minimum type (in the ordering $\succeq$) that buys in any given period. Thus,

[2]Recall that an ordering $\succeq$ on a set S is *linear* (also called *total*) if for $a, b \in S$, either $a \succeq b$ or $b \succeq a$.

single-dimensionality of Θ allows us to characterize the buying behavior of *all types* in a given period by specifying what the minimum buying type is.

Before moving on with the approach to solving the seller's maximization problem, we pause to give two examples of single-dimensional types, which we will use in the single-arrival and constant-arrival models where we do an in-depth analysis of price paths and welfare. Both examples use exponential time-discounting, which is common in the literature on intertemporal price discrimination and convenient to use, since for any given type, their incentives for purchase timing is independent of their arrival time to the market, so in the analysis, one does not need to keep track of the various cohorts of arrivals when computing stopping times or value functions. The first example has buyers differ on willingness-to-pay for the good, and the second example has buyers differ in discount rates over two different willingness-to-pay levels. Both examples feature utility that is linear in price, so Theorem 7 applies to any dynamic model with such a space of buyers. Another example of single-dimensional types is a combination of Examples 2 and 3 such that each ordering is preserved, so that for two types $(v, r), (v', r')$, $v \geq v'$ if and only if $r \leq r'$.

Example 2 (Varying buyer willingness-to-pay)**.** Let $\Theta = \{v_1, \ldots, v_N\} \subset \mathbb{R}_+$ with $v_1 < \cdots < v_N$. There is a fixed discount rate $r > 0$ by which buyers discount future surplus. Buyer utility of purchasing at price p at time t is given by

$$u_t(p, v) = e^{-rt}(v - p).$$

While it is intuitive that this Θ is a single-dimensional type space, since any higher type buyer should buy when a lower type buyer does so, it is instructive to show that this follows through the stopping times of the buyers and indeed satisfies single-dimensionality in our formal definition.

Take any types $v \geq v'$ from Θ. Take any time t and price path p (fixing some finite horizon T), where $b_t(v') = 1$ (at least one such price path exists: set $p_t = v_1$ for all t). Then type-v' has the stopping condition satisfied:

$$e^{-rt}(v' - p_t) \geq \sup_{\tau \geq t+1} \mathbb{E}[e^{-r\tau}(v' - p_\tau)]$$

which can be rearranged to

$$\begin{aligned} p_t &\leq \sup_{\tau \geq t+1} \mathbb{E}[(1 - e^{-r\tau})v' - e^{-r\tau}p_\tau] \\ &\leq \sup_{\tau \geq t+1} \mathbb{E}[(1 - e^{-r\tau})v - e^{-r\tau}p_\tau], \end{aligned}$$

thus the stopping condition for v is also satisfied, and $b_t(v) = 1$. To prove the converse by contrapositive, let $v > v'$, and we will find a price path and time t such that $b_t(v) = 1$ and $b_t(v') = 0$. Let $p_t = (1 - e^{-r})v + e^{-r}v_0$ and $p_{t'} = v_0$ for $t' \in \{t+1, \ldots, T\}$. Then

$$v - p_t = e^{-r}(v - v_0) = \sup_{\tau \geq t+1} \mathbb{E}[e^{-r\tau}(v - p_\tau)]$$

so $b_t(v) = 1$, but

$$\begin{aligned} v' - p_t &= v' - (1 - e^{-r})v - e^{-r}v_0 \\ &= (1 - e^{-r})(v' - v) + e^{-r}(v' - v_0) \\ &< e^{-r}(v' - v_0) \\ &= \sup_{\tau \geq t+1} \mathbb{E}[e^{-r\tau}(v' - p_\tau)]. \end{aligned}$$

Thus, we have formally shown that $v \geq v' \iff v \succeq v'$. The ordering on Θ is then a linear order since $\geq$ is a linear order on $\mathbb{R}$.

Example 3 (Varying buyer discount rates)**.** We now let $R = \{r_0, r_1, \ldots, r_N\}$ where $0 < r_0 < r_1 < \cdots < r_N$. There are two willingness-to-pay levels $0 < v_L < v_H$. The type space is then $\Theta = \{(v_H, r_0), (v_H, r_1), \ldots, (v_H, r_N), (v_L, r_0)\}$. We may make the type space be $\{\theta_L, \theta_H\} \times R$ rather than the given subset, but in the analysis it is without loss of generality to have the former structure. Defining Θ in this way also makes Θ single-dimensional under our definition, whereas the cross-product is only single dimensional restricting to prices that are bounded below by v_L (which they will be in any optimal price path).

In this set, it will turn out that $(v_H, r) \succeq (v_H, r') \iff r > r'$, or in other words, impatient types buy whenever more patient types buy. Also, $(v_H, r) \succeq (v_L, r_0)$: high willingness-to-pay types are also always more impatient than the low willingness-to-pay types, so the latter purchasing implies the former. As in the previous example, the stopping condition of type (v_H, r') is

$$\begin{aligned} p_t &\leq \sup_{\tau \geq t+1} \mathbb{E}\left[(1 - e^{-r'\tau})v_H - e^{-r'\tau}p_\tau\right] \\ &\leq \sup_{\tau \geq t+1} \mathbb{E}\left[(1 - e^{-r\tau})v_H - e^{-r\tau}p_\tau\right] \end{aligned}$$

when $r' < r$ since an optimal τ should only have support on t such that $p_t \leq v_H$, and similarly

$$\sup_{\tau \geq t+1} \mathbb{E}\left[(1 - e^{-r_0\tau})v_L - e^{-r_0\tau}p_\tau\right] \leq \sup_{\tau \geq t+1} \mathbb{E}\left[(1 - e^{-r\tau})v_H - e^{-r\tau}p_\tau\right]$$

since $r \geq r_0$, and the LHS is either equal to zero (the low type never purchases), in which the RHS also exceeds this value, or the LHS is not zero and there is a stopping time for which the low type buys, which will only occur at times t such that $p_t \leq v_L$. For such τ, the RHS exceeds the LHS since it weights v_H higher and $v_H > v_L \geq p_\tau$. Thus we have shown that $b_t(v_H, r) \geq b_t(v_H, r')$ if $r \geq r'$ and $b_t(v_H, r) \geq b_t(v_L, r_0)$ for all $r \in R$. If $r < r'$, we can show there exists a price path and time such that $b_t(v_H, r) = 1$ and $b_t(v_H, r') = 0$. Let $p_t = (1 - e^{-r})v_H + e^{-r}v_L$ and $p_{t'} = v_L$ for $t' \in \{t+1, \ldots, T\}$. Then

$$v_H - p_t = e^{-r}(v_H - v_L) = \sup_{\tau \geq t+1} \mathbb{E}[e^{-r\tau}(v_H - p_\tau)]$$

so $b_t(v_H, r) = 1$, but

$$v_H - p_t = e^{-r}(v_H - v_L) < e^{-r'}(v_H - v_L) = \sup_{\tau \geq t+1} \mathbb{E}[e^{-r'\tau}(v_H - p_\tau)]$$

so $b_t(v_H, r') = 0$. Here p_t is the reserve price for type (v_H, r) to buy at time t or wait until time $t+1$ to buy at price v_L. Since (v_H, r') is more patient, they would rather wait until $t+1$. Thus, we have the ordering $(v_H, r_N) \succeq \cdots \succeq (v_H, r_0) \succeq (v_L, r_0)$, which is linear.

Definition 4. If Θ is single-dimensional, a *type path* is a stochastic process $\vec{\theta} = (\theta_0, \ldots, \theta_T) \in \Delta(\Theta^{T+1})$, where θ_t is the minimum buying type in period t by the ordering $\succeq$. We refer to θ_t as the *threshold type*.

When types are multi-dimensional, we do not get the clean mapping between prices and types as in Proposition 1, which in general will create nonlinearities in continuation values and prices, and can lead to nondeterministic pricing. But before moving on, we note that while we can specify a type path, it is worth asking if it even makes sense. It only makes sense if there exists a price path that maps to the type path via the optimal stopping rules of the buyers, and it is only useful to analyze if we can characterize the relationship between the two.

Definition 5. Let $\vec{\theta}$ be a type path. We say that $\vec{\theta}$ is *feasible* if there exists a price path p such that p maps to $\vec{\theta}$ via optimal stopping rules for the buyers, almost surely[3]. For such a p, we say that p is *consistent* with $\vec{\theta}$.

In the single-dimensional case of Θ, any type path $\vec{\theta} \in \Delta(\Theta^{T+1})$ is feasible, as we will show. In other words, the mapping $\vec{p} \mapsto \vec{\theta}$ is onto $\Delta(\Theta^{T+1})$. In general, there may be many price paths

[3]Technically, $\vec{p}$ and $\vec{\theta}$ must be defined on the same underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and must lead to the same buyer outcomes with the exception for an event $E \in \mathcal{F}$ with $\mathbb{P}(E) = 0$, as well as satisfying the requirement that both are adapted to the given filtration. Cf. Section 7, Chapter V of Cinlar (2011).

that are consistent with a given type path, or in other words, the mapping $\vec{p} \mapsto \vec{\theta}$ is not one-to-one. In terms of solving the seller's optimality problem, we would like to solve for the best $\vec{p}$ that is consistent with $\vec{\theta}$, if it exists. Again, this is possible to do, allowing us to define an inverse mapping $\vec{\theta} \mapsto \vec{p}$.

As the optimal stopping problems for the buyers are amenable to dynamic programming methods, rather than specifying joint distributions in the price path or type path, it is useful to specify the sequence of conditional distributions: $p_0; p_1|p_0; p_2|p_0, p_1; \ldots; p_T|p_0, \ldots, p_{T-1}$, and the analogous conditioning for type paths.

Proposition 1. Given $\vec{\theta} = (\theta_0, \ldots, \theta_T) \in \Delta(\Theta^{T+1})$, $\vec{\theta}$ is feasible: there exists a consistent $\vec{p} = (p_0, \ldots, p_T) \in \Delta(\mathbb{R}_+^{T+1})$, defined recursively by

$$\begin{aligned} p_T &= u_t^{-1}(0, \theta_T) \\ p_t &= u_t^{-1}(\sup_{\tau \geq t+1} \mathbb{E}\left[u_\tau(p_\tau, \theta_t)|\theta_0, \ldots, \theta_t\right], \theta_t); \quad t \in [T-1]. \end{aligned}$$

Furthermore, this price path yields the maximum seller profits among all price paths consistent with $\vec{\theta}$.

Proof. Consistency of $\vec{p}$. We show that after each history of types up to a given time t, p_t is consistent, and therefore the entire path of prices $\vec{p}$ is consistent. Let h_{t-1} denote some history of types' purchases up to and including time $t-1$, and let θ_t be a threshold type in the support in $\vec{\theta}$ conditional on h_{t-1}. Then the price

$$p_t = u_t^{-1}(\sup_{\tau \geq t+1} \mathbb{E}[u_\tau(p_\tau, \theta_t)|h_{t-1}, \theta_t])$$

is the reserve price of type θ_t, so that $u_t(p_t, \theta_t) = \sup_{\tau \geq t+1} \mathbb{E}[u_\tau(p_\tau, \theta_t)|h_{t-1}, \theta_t]$. Thus the seller can choose $b_t(\theta_t) = 1$ since stopping immediately satisfies the buyer optimal stopping problem, and single-dimensionality implies $b_t(\theta) = 1$ for $\theta \succeq \theta_t$. On the other hand, if $\theta_t \succeq \theta$ for other types θ, the multi-valued arg max of the buyer stopping problem (either stopping at t or some stopping time $\geq t$ are both optimal), the seller can choose $b_t(\theta_t) = 0$, implying that it must be $b_t(\theta) = 0$ or that the seller can choose $b_t(\theta) = 0$ if θ is indifferent. Thus, we construct, history by history, a price history that corresponds to the type history. Thus, $\vec{p}$ is consistent with $\vec{\theta}$, showing that $\vec{\theta}$ is feasible.

Seller Optimality of $\vec{p}$. Suppose that $(\tilde{p}_0, \ldots, \tilde{p}_T)$ is another consistent price path (if another one does not exist then $\vec{p}$ is trivially optimal). First, we show that $\tilde{p}_t \leq p_t$ for all t and histories h_t except on a set of h_t's of measure 0. We do this by backward induction, starting from $t = T$. If

$\tilde{p}_T > p_T$, when by strict monotonicity of u_t in p, we have

$$u_T(\tilde{p}_T, \theta_T) < 0,$$

violating consistency for $\tilde{p}_T$, which can only occur on a set of measure 0 histories H_T. Now proceeding to the inductive step, assume the inductive hypothesis that $\tilde{p}_k \leq p_k$ for $k = t+1, \ldots, T$ except on a measure-zero set of histories H_{t+1}. Then for any partial history up to t,

$$\mathbb{E}[u_\tau(p_\tau, \theta_t)|h_t] \leq \mathbb{E}[u_\tau(\tilde{p}_\tau, \theta_t|h_t]$$

for all stopping times $\tau \geq t+1$, since $\tilde{p}_k \leq p_k$, $k \geq t+1$. Again we must exclude a set of (full) histories of measure 0, which are histories in H_{t+1}. Since this holds for all stopping times, it holds for the supremum, and therefore if $\tilde{p}_t > p_t$, we have

$$u_t(\tilde{p}_t, \theta_t) < u_t(p_t, \theta_t) = \sup_{\tau \geq t+1} \mathbb{E}[u_\tau(p_\tau, \theta_t)|h_t] \leq \sup_{\tau \geq t+1} \mathbb{E}[u_\tau(\tilde{p}_\tau, \theta_t|h_t]$$

which would violate consistency of $\tilde{p}$, and so can only happen on a set of measure 0 histories, which we combine with H_{t+1} to get a measure-zero set H_t. Applying induction yields the result.

Consider $\mathbb{E}[e^{-rt}M_t p_t]$, the seller's expected discounted profits from sales at t. Almost surely, $p_t \geq \tilde{p}_t$, and almost surely, M_t is the same under p_t and $\tilde{p}_t$. Thus, $\mathbb{E}[e^{-rt}M_t p_t] \geq \mathbb{E}[e^{-rt}M_t \tilde{p}_t]$. $\square$

The above result is intuitive: given a type path, just define the price path in terms of the reserve prices of the threshold type after any conditional distribution up to that point in time. All of the higher types will buy and the lower types will not buy, and it leaves demand unaffected while charging the highest price possible. Thus we have defined a mapping $\vec{\theta} \mapsto \vec{p}$ that maximizes seller profits while respecting consistency of prices. Proposition 1 has the fairly immediate consequence that prices are bounded in any seller-optimal price path.

Corollary 6. A seller-optimal price path has bounded prices:

$$p_t \in [u_t^{-1}(0, \underline{\theta}), u_t^{-1}(0, \overline{\theta})].$$

Proof. Take any type path $\vec{\theta}$, then the corresponding price path $\vec{p}$ via Proposition 1 has

$$\begin{aligned} p_t &= u_t^{-1}(\sup_{\tau>t} \mathbb{E}_t[u_t(p_\tau, \theta_t)], \theta_t) \\ &\leq u_t^{-1}(0, \theta_t) \\ &\leq u_t^{-1}(0, \overline{\theta}), \end{aligned}$$

so the upper bound is proved. For the lower bound, note that $p_T = u_t^{-1}(0, \theta_T) \geq u_t^{-1}(0, \underline{\theta})$. Under seller-optimal prices, $\sup_{\tau>t} \mathbb{E}_t[u_t(p_\tau, \underline{\theta})] = 0$, since $\underline{\theta}$ only buys at times when $\theta_t = \underline{\theta}$. Reserve utility of $\underline{\theta}$ to any point in time is then 0. Since $\theta_t \succeq \underline{\theta}$, it must be that the reserve price of θ_t at time t is greater than the reserve price of $\underline{\theta}$ at time t, $p_t \geq u_t^{-1}(0, \underline{\theta})$. □

In the case where buyer utility is of the form $u_t(p, \theta) = e^{-r(\theta)t}(v(\theta) - p)$, this means that p_t is bounded between $v(\underline{\theta})$ and $v(\overline{\theta})$, the lowest and highest valuations of the good. The main benefit of converting $\vec{p}$ to $\vec{\theta}$, however, is that demand M_t, as a function of $\vec{\theta}$, only depends on the history of threshold types up to time t, whereas M_t was dependent on the entire $\vec{p}$ since the distribution of future prices determined optimal stopping times. From this perspective, when the seller picks a type path, they are choosing the stopping times of the various buyer types, and prices correspond to the maximum prices such that the given stopping times are optimal for each buyer type. This is analogous to incentive compatibility restrictions in mechanism design. Since M_t only depends on information up to time t, $\mathbb{E}[M_t | \theta_0, \ldots, \theta_t] = M_t | \theta_0, \ldots, \theta_t$ is not random.

With the switch to the choice over type paths, we can also define the buyer's value function to be used in their optimal stopping problem. Using the standard finite-horizon approach to dynamic programming, we define the value function for the buyer of type θ, given a type path $\vec{\theta} = (\theta_0, \ldots, \theta_T)$ as follows:

$$\begin{aligned} p_T &= u_T^{-1}(0, \theta_T) \\ V_T(\theta) &= \max\{u_T(p_T, \theta), 0\} \\ p_t &= u_t^{-1}(\mathbb{E}_t[V_{t+1}(\theta_t)], \theta_t) \\ V_t(\theta) &= \max\{u_t(p_t, \theta), \mathbb{E}_t[V_{t+1}(\theta)]\}. \end{aligned}$$

Substituting the expression for p_t in the value function V_t, we obtain a recursion purely in terms of value functions for V:

$$V_t(\theta) = \max\{u_t(u_t^{-1}(\mathbb{E}_t[V_{t+1}(\theta_t)], \theta_t), \theta), \mathbb{E}_t[V_{t+1}(\theta)]\}$$

The price path $\vec{p}$ and the value function V have some monotonicity properties that are intuitive and used in the proof of the deterministic pricing theorem.

Proposition 2. If Θ is single-dimensional, given a type path $\vec{\theta} = (\theta_0, \dots, \theta_T)$, $\vec{p}$ and V have the following properties:

(a) The price p_t is increasing in the threshold type θ_t.

(b) $V_t(\theta)$ is decreasing in the threshold type θ_t for every $\theta \in \Theta$.

(c) V_t is increasing in θ (its argument).

(d) For each t and θ and type history $\theta_0, \dots, \theta_t$, $b_t(\theta) = 1 \iff V_t(\theta) = u_t(p_t, \theta)$.

Proof. We prove all of these properties simultaneously through induction. Consider T, $p_t = u_T^{-1}(0, \theta_T)$ is increasing in θ_T, $V_T = \max\{u_T(u_T^{-1}(0, \theta_T), \theta), 0\}$ has the first argument of the max that is decreasing in price, and since price is increasing in θ_T, the first argument is decreasing in θ_T. Thus the overall max is constant in θ when the second argument is greater, and decreasing in θ_T when the first argument is greater, so V_t is decreasing in θ_T. For fixed θ_T, V_T is increasing in θ since $u_T(p, \theta)$ is increasing in θ. Finally, type θ buys at T if and only if $V_T(\theta)$ equals its first argument, $u_T(p_T, \theta)$, since the continuation utility is 0.

Now assume the induction hypothesis. Then $p_t = u_t^{-1}(\mathbb{E}_t[V_{t+1}(\theta_t)], \theta_t)$. Let $\theta \succ \theta'$. If $p_t(\theta) \leq p_t(\theta')$, θ' would buy at $p_t(\theta)$ since the price would be lower than their reserve price. This violates consistency of the price path p_t, so $p_t(\theta) > p_t(\theta')$. To prove (b), note that the threshold type only affects the first argument of the maximum in the expression of V_t. By (a), p_t is increasing in θ_t, so $u_t(p_t, \theta)$ is decreasing in V_t, hence V_t is decreasing in θ_t. To prove (c), from induction, $V_{t+1}(\theta)$ is increasing in θ for all possible θ_{t+1}, from the induction hypothesis, and $u_t(p_t, \theta)$ is increasing in θ. Then the maximum of two increasing functions is increasing, so V_t is increasing in θ. To prove (d), note from the induction hypothesis that $u_{t+1}(p_{t+1}, \theta) \geq \sup_{\tau \geq t+2} \mathbb{E}_t[u_\tau(p_\tau, \theta)]$ if and only if $V_{t+1}(\theta) = u_{t+1}(p_{t+1}, \theta)$. This means that $V_{t+1}(\theta) = \sup_{\tau \geq t+1} \mathbb{E}_{t+1}[u_\tau(p_\tau, \theta)]$. Then $V_t(\theta) = u_t(p_t, \theta)$ if and only if $u_t(p_t, \theta) \geq \mathbb{E}_t[\sup_{\tau \geq t+1} \mathbb{E}_{t+1}[u_\tau(p_\tau, \theta)]] = \sup_{\tau \geq t+1} \mathbb{E}_t[u_\tau(p_\tau, \theta)]$, which is the condition for $b_t(\theta) = 1$. $\square$

We say that buyer utility is linear in price if there exist functions $\alpha, \beta : [T] \times \Theta \to \mathbb{R}$ such that $u_t(p, \theta) = \alpha_t(\theta) p + \beta_t(\theta)$. By our assumptions that u_t is invertible in p for every θ and is decreasing in p implies that $\alpha_t < 0$. The structure is now in place to state and prove the deterministic pricing theorem for single-dimensional type spaces.

Theorem 7. If Θ is single-dimensional and buyer utility is linear in price, then there exists a deterministic seller-optimal price path.

Proof. We first apply the linearity of utility to obtain the following linear recursion in V:

$$\begin{aligned} V_t(\theta) &= \max\left\{u_t(p_t, \theta), \mathbb{E}_t[V_{t+1}(\theta)]\right\} \\ &= \max\{F_t(\theta, \theta_t)\mathbb{E}_t[V_{t+1}(\theta_t)] + G_t(\theta, \theta_t), \mathbb{E}_t[V_{t+1}(\theta)]\} \\ &= \mathbb{E}_t\left[\tilde{F}_t(\theta, \theta_t)V_{t+1}(\min\{\theta, \theta_t\}) + \tilde{G}_t(\theta, \theta_t)\right] \end{aligned}$$

where

$$\begin{aligned} F_t(\theta, \theta_t) &= \frac{\alpha_t(\theta)}{\alpha_t(\theta_t)}, \\ G_t(\theta, \theta_t) &= \frac{\alpha_t(\theta_t)\beta_t(\theta) - \alpha_t(\theta)\beta_t(\theta_t)}{\alpha_t(\theta_t)} \\ \tilde{F}_t(\theta, \theta_t) &= F_t(\theta, \theta_t)\mathbb{1}\{\theta \geq \theta_t\} + \mathbb{1}\{\theta < \theta_t\} \\ \tilde{G}_t(\theta, \theta_t) &= G_t(\theta, \theta_t)\mathbb{1}\{\theta \geq \theta_t\}. \end{aligned}$$

The last equality is justified since $\tilde{F}$ and $\tilde{G}$ are determined with respect to information at time t and that V_{t+1} is increasing in its argument, so that $V_{t+1}(\theta_t)\mathbb{1}\{\theta \geq \theta_t\} + V_{t+1}(\theta)\mathbb{1}\{\theta < \theta_t\} = V_{t+1}(\min\{\theta, \theta_t\})$. It is convenient to use the notation

$$\theta_{t:t+k} = \min\{\theta_t, \theta_{t+1}, \ldots, \theta_{t+k}\}$$

for $k \in \mathbb{N}$, as it is the "running minimum" of threshold θ's that are passed on to future value functions in the recursion. We further expand the recursion on the RHS by substituting all the way to T for the argument θ_t:

$$\begin{aligned}
V_t(\theta_t) &= \mathbb{E}_t[V_{t+1}(\theta_t)] \\
&= \mathbb{E}_t\left[\mathbb{E}_{t+1}\left[\tilde{F}_{t+1}(\theta_t,\theta_{t+1})V_{t+2}(\theta_{t:t+1}) + \tilde{G}_{t+1}(\theta_t,\theta_{t+1})\right]\right] \\
&= \mathbb{E}_t\left[\tilde{F}_{t+1}(\theta_t,\theta_{t+1})V_{t+2}(\theta_{t:t+1}) + \tilde{G}_{t+1}(\theta_t,\theta_{t+1})\right] \\
&= \mathbb{E}_t[\tilde{F}_{t+1}(\theta_t,\theta_{t+1})\tilde{F}_{t+2}(\theta_{t:t+1},\theta_{t+2})V_{t+3}(\theta_{t:t+2}) + \tilde{F}_{t+1}(\theta_t,\theta_{t+1})\tilde{G}_{t+2}(\theta_{t:t+1},\theta_{t+2}) + \tilde{G}_{t+1}(\theta_t,\theta_{t+1})] \\
&= \cdots \\
&= \mathbb{E}_t\left[\sum_{k=1}^{T-t}\tilde{G}_{t+k}(\theta_{t:t+k-1},\theta_{t+k})\prod_{j=1}^{k-1}\tilde{F}_{t+j}(\theta_{t:t+j-1},\theta_{t+j})\right] \\
&= \mathbb{E}_t\left[\Xi_t(\theta_t,\ldots,\theta_T)\right]
\end{aligned}$$

since when we expand the recursion, $V_{T+1}(\theta) = 0$ for all θ.

The discounted expected seller profit from period t given the type path $\vec{\theta}$ is

$$\begin{aligned}
e^{-rt}\mathbb{E}[M_t p_t] &= e^{-rt}\mathbb{E}[M_t\alpha_t^{-1}(\theta_t)(\mathbb{E}_t[V_{t+1}(\theta_t)] - \beta_t(\theta_t))] \\
&= e^{-rt}\mathbb{E}[M_t\alpha_t^{-1}(\theta_t)(\mathbb{E}_t[\Xi_t(\theta_t,\ldots,\theta_T)] - \beta_t(\theta_t))] \\
&= \mathbb{E}\left[e^{-rt}\left(M_t\alpha_t^{-1}(\theta_t)(\Xi_t(\theta_t,\ldots,\theta_T) - \beta_t(\theta_t))\right)\right]
\end{aligned}$$

Summing the discounted expected profits over t, we have the seller's profit equal to

$$\mathbb{E}\left[\sum_{t=0}^{T} e^{-rt}\left(M_t\alpha_t^{-1}(\theta_t)(\Xi_t(\theta_t,\ldots,\theta_T) - \beta_t(\theta_t))\right)\right].$$

Since the expectation of a random variable is less than its maximum, if it exists, this means that

$$(\theta_0,\theta_1,\ldots,\theta_T) \in \arg\max_{\Theta^{T+1}}\sum_{t=0}^{T} e^{-rt}\left(M_t\alpha_t^{-1}(\theta_t)(\Xi_t(\theta_t,\ldots,\theta_T) - \beta_t(\theta_t))\right)$$

maximizes the seller's discounted expected profits. It remains to show that a maximum exists. If Θ is a finite set, the maximum exists since the set of values of the expression in the expectation is finite. If Θ is a general compact set, then continuity of α_t and β_t in θ ensures that F_t, G_t are continuous in all arguments and that $\tilde{F}_t, \tilde{G}_t$ are continuous: the only possible points of discontinuity are (θ_t, θ_t), but $\lim_{\theta\downarrow\theta_t}\tilde{F}_t(\theta,\theta_t) = 1$, and $\tilde{F}_t(\theta,\theta_t) = 1$ for $\theta_t \succeq \theta$, and $\lim_{\theta\downarrow\theta_t}\tilde{G}_t(\theta,\theta_t) = 0$, and $\tilde{G}_t(\theta,\theta_t) = 0$ for $\theta_t \succeq \theta$, so $\alpha_t^{-1}, \Xi_t, \beta_t$ are all continuous, and M_t is continuous, hence the seller profit is continuous over the compact set Θ^{T+1}, so there exists a maximum by the extreme value theorem. $\square$

The deterministic pricing theorem tells us that under broad conditions with buyers who have utility linear in prices, a monopolist who can commit to a price path should announce a nonrandom price path to buyers at the beginning of the game. The monopolist cannot encourage buyers to buy earlier or at higher prices by threatening random prices in the future, since utility as a function of price will only depend on the expectation of prices. Randomness per se does not help the seller by influencing continuation values, and thus the seller can simply choose a deterministic price path to maximize their discounted profits.

The deterministic pricing result as stated above can also be extended in several ways. In particular, the assumption of single-dimensional type space can be dropped. Dropping this assumption just makes the question of what type paths are feasible more complicated, and can be dealt with at the cost of a more complex maximization problem. This line of reasoning is developed in the appendix. The approach doesn't differ significantly in terms of method from what is given above. On the other hand, the deterministic pricing result can be extended to infinite-horizon problems, as well as continuous-time problems, subject to technical conditions. The infinite-horizon and continuous-time results extend to multi-dimensional types as well.

Theorem 8 (Infinite-Horizon Discrete Time Deterministic Pricing.)**.** In the infinite-horizon discrete time model, suppose Θ is single-dimensional and buyer utility is linear in price. Suppose that there exists t^* such that $t \geq t^*$ implies that $u_t(p, \theta) \leq \delta^t v(p, \theta)$ for some function v and $0 < \delta < 1$. Assume that M_t is bounded uniformly over t and all histories of types, and u_t^{-1} is uniformly bounded in t and θ. Then deterministic price processes approximate the supremum of the seller profits. If there exists a maximum, there exists a deterministic seller-optimal price path.

Proof. We apply Propositions A6.11 and A6.12 from Kreps (2013) to the value functions of the buyers for optimal stopping times. Fix a bounded price path $\{p_t\}_{t=0}^{\infty}$. We can approximate this problem with the finite horizon problem under the price path that follows the above up to T, after which the game ends. The function $U_t(\tilde{p}, \theta) = \sum_{k=t}^{\infty} \max\{0, u_k(\tilde{p}_k, \theta)\}$ where $\tilde{p}$ is one realized path of price. Then $U_t(\tilde{p}, \theta) \leq \sum_{k=t}^{t^*-1} \max\{0, u_k(\tilde{p}_k, \theta)\} + \sum_{k=t^*}^{\infty} \delta^k \max\{0, v(\tilde{p}_k, \theta)\}$ where the RHS is uniformly bounded (over $\tilde{p}$) for fixed θ since $u_t(\cdot, \theta)$ is uniformly bounded over $t = 0, \ldots, t^* - 1$ and v is bounded over possible values of $\tilde{p}$, so the second part of the sum is uniformly bounded. Hence $U(\cdot, \theta)$ is uniformly bounded. Then Proposition A6.11 implies that the finite-horizon utilities for any stopping strategy approximate the infinite horizon utilities from that stopping strategy as $T \to \infty$. Then Proposition A6.12 implies that the finite-horizon optimal stopping strategies approximate the infinite-horizon optimal stopping strategies as $T \to \infty$, thus we can approximate the infinite-horizon seller profits by the finite horizon problem

$$\mathbb{E}\left[\sum_{t=0}^{T} e^{-rt} M_t p_t\right] \tag{2}$$

where the M_t and p_t are defined in the discrete-time finite horizon problem. Since u_t^{-1} is uniformly bounded, we price p_t can be bounded by a constant, and M_t is uniformly bounded by a constant, therefore the sequence $M_t p_t$ is bounded, and in ℓ^∞, hence $\sum_{t=0}^{\infty} e^{-rt} M_t p_t$ exists for bounded sequences of p_t. Let $p^{(T)} \in \ell^\infty$ that maximizes $\mathbb{E}\left[\sum_{t=0}^{T} e^{-rt} M_t p_t\right]$, with $p_t^{(T)} = \bar{p}$ for $t > T$, where $\bar{p}$ is the uniform upper bound on u_t^{-1}. Let $\Pi^{(T)}$ be the profits up to time T from $p^{(T)}$. Then

$$\mathbb{E}\left[\sum_{t=0}^{\infty} e^{-rt} M_t p_t^{(T)}\right] = \Pi^{(T)} + e^{-r(T+1)} \mathbb{E}\left[\sum_{t=T+1}^{\infty} e^{r(t-T+1)} M_t \bar{p}\right]$$

since stopping times in $t = 0, \ldots, T$ are unaffected by the highest prices following T. Note that an upper bound on maximum profits over all stochastic price processes is then $\Pi^{(T)} + \frac{e^{-r(T+1)}}{1-e^{-r(T+1)}} \bar{M}\bar{p}$, and since $\frac{e^{-r(T+1)}}{1-e^{-r(T+1)}} \to 0$ as $T \to \infty$, there exists T^* for which $T \geq T^*$ implies that $\Pi^{(T)} \geq \sup_p \mathbb{E}\left[\sum_{t=0}^{\infty} e^{-rt} M_t p_t\right] - \epsilon$. Suppose there is a stochastic price process $\tilde{p}$ such that $\mathbb{E}\left[\sum_{t=0}^{\infty} e^{-rt} M_t \tilde{p}_t\right] > \mathbb{E}\left[\sum_{t=0}^{\infty} e^{-rt} M_t p_t\right]$ for all deterministic price processes p, by at least ϵ. Then for large T, $\mathbb{E}\left[\sum_{t=0}^{\infty} e^{-rt} M_t \tilde{p}_t\right] > \Pi^{(T)} + \epsilon \geq \sup_p \mathbb{E}\left[\sum_{t=0}^{\infty} e^{-rt} M_t p_t\right]$, a contradiction. Thus deterministic price processes are approximate the supremum, and if the supremum is achieved, it is done so by a deterministic process. □

4 Single-Arrival Model

4.1 Model Description

We now delve into the first application of the deterministic pricing theorem. Moving to continuous time over an infinite horizon, we assume that all agents arrive to the market at time 0, hence the term *single-arrival*. Although Theorem 8 is a discrete-time result, the optimality of deterministic price is preserved in the continuous-time limit. In terms of buyer types, we consider the two type spaces given in Examples 2 and 3 and their single-dimensional-preserving combination, where buyers differ in willingness-to-pay and patience levels, respectively. In both cases, buyer utility is linear in price and has exponential time-discounting: $u_t(p, \theta, r) = e^{-rt}(\theta - p)$. Arrival rates are given by $\lambda_0(\theta, r) > 0$ and no arrivals after time 0 means that $\lambda_t(\theta, r) = 0$ for $t > 0$. This model is a no-entry and no-production cost model in the case where the type space varies in the willingness-to-pay and patience dimensions (albeit in a single-dimensional manner), but for this exposition, we use a discrete type space. The deterministic pricing theorem is a justification for much of the intertemporal price discrimination literature to consider only price schedules that are deterministic, since in her model formulation, buyer utility is always linear in price. Consequently, the optimality results given in Stokey (1979), Sobel (1991), Board (2008), and others are robust to any ability of the seller to commit to a stochastic price schedule.

There are a few features of this version of the model that we can immediately take note of. First, any profit-maximizing price process should never take values less than the minimum willingness-to-pay by Corollary 6. Next, given any price process, because the type spaces we consider are single-dimensional in the sense of Definition 1, the optimal stopping times are by definition ordered by $\tau_0^*(\theta) \leq \tau_0^*(\theta')$ if and only if $\theta \succeq \theta'$. To ensure that optimal stopping times exist for the buyers, we require that any deterministic price path be lower semicontinuous, so that if price is trending downwards, the buyer is able to buy at the infimum of the trend. From the ordering we know that the highest type $\overline{\theta}$ will buy first, and then others follow. Therefore, if $\tau_0^*(\overline{\theta}) > 0$, there is a period of time for which no types buy. Such an outcome is suboptimal from the seller's point of view, since they could shift the price path to start at $p_{\tau_0^*(\overline{\theta})}$ at time 0 and follow the price path from there, so that $\tau_0^*(\overline{\theta}) = 0$ and all of the other stopping times are shifted back by the original amount of $\tau_0^*(\overline{\theta})$.

Since there are no arrivals after time 0, the game essentially ends when price reaches $\underline{\theta}$, since after that point, demand will always be 0. The dimension of stochastic price processes that a seller could implement is very large and the optimal stopping conditions for a buyer complex for a general stochastic process, thus the deterministic pricing theorem allows us to gain tractability on the seller's optimization problem. Under a deterministic price path, the buyer optimal stopping time is

$$\arg\max_{t \geq 0} e^{-rt}(\theta - p_t),$$

and if the arg max is multivalued, we break indifference by choosing the minimum time in the arg max. Instead of choosing a price path directly, as before the seller instead chooses threshold types that characterize purchasing decisions of each type at each point in time, and pricing is determined by the indifference curve of the threshold type. Since after a given type buys, demand clears for that type, there are points in time where no types are buying. For these times, prices can be determined by a threshold type, but since such a type has already bought, price can be anything that lies above the threshold price. Thus the only choices of threshold types that will have any effect on profits are the first times a given type is chosen as a threshold type, or in other words, the stopping times τ_0^*. Thus the problem of the seller is to choose the stopping times of the various buyer types to maximize profits. This leads to the following optimization program:

$$\max_{t_1,\ldots,t_N} \sum_{i=1}^{N} e^{-r_S t_i} \lambda_0(\theta_i) p_{t_i}$$
$$\text{subject to}$$
$$p_{t_i} = v_i - e^{-r_i(t_i - t_{i-1})}(v_i - p_{t_{i-1}}), \qquad i = 2, \ldots, N$$
$$p_{t_1} = v_1$$
$$0 = t_N \leq t_{N-1} \leq \cdots \leq t_1.$$

Note that we do not have to consider constraints on price other than the incentive restrictions that type i does not want wait to buy when type $i-1$ buys, as the other incentive constraints are implied by these due to the ordering of stopping times. The incentive constraints are binding for all types by the arguments laid out in the deterministic pricing section as well: the best price that the seller can do given a feasible path is to set them equal to the level that makes the threshold type indifferent between buying now and waiting. Equivalently, we can define the intervals $\Delta_i = t_{i-1} - t_i$ so that $t_i = \Delta_N + \ldots + \Delta_{i+1}$, and maximize over $\Delta_i \geq 0$ for each i. The objective function is concave in the Δ_i's, and while it has no closed form solution, the optimal price path can be solved for recursively, by optimizing over the lower prices that occur later in time independently of the higher earlier prices, and sequentially solving the higher earlier prices based on the optimal later prices.

4.2 Analysis of Optimal Price Paths

The approach begins with considering another transformation of the seller's problem. If we consider the choice variable of the seller as price once again, so that the choice is over price $(p_1, \ldots, p_N)$, we can determine the Δ_i intervals through the equations

$$e^{-r_i \Delta_i} = \frac{v_i - p_i}{v_i - p_{i-1}}$$
$$\Delta_i = \frac{1}{r_i}(\log(v_i - p_{i-1}) - \log(v_i - p_i)),$$

which says that the time interval between two prices p_i and p_{i-1} is set so the discount factor is equal to the ratio of surplus the buyer receives at price p_i to the surplus at price p_{i-1}. Thus, the interval Δ_i is approximately proportional to the percentage change in surplus from price p_i to p_{i-1}. This yields the insight that a higher difference in prices paid by two adjacent types means that a longer interval must pass between the times they purchase. It also says that the time interval between two prices only depends on those two prices, and not on prices elsewhere in the price path. This independence from nonadjacent prices suggests that the problem can be subdivided into subproblems, where optimality of in the subproblem is independent of other prices in the larger problem. This is indeed the case. Note that $\Delta_i \geq 0$ restriction says that $p_i \geq p_{i-1}$, so that $v_1 = p_1 \leq p_2 \leq \cdots \leq p_N$.

We first write out the seller problem in terms of prices. Since $t_i = \sum_{j=i+1}^{N} \Delta_j$, we have that the seller discount factor to time t_i is

$$e^{-r_S t_i} = e^{-r_S \sum_{j=i+1}^{N} \Delta_j} = \prod_{j=i+1}^{N} e^{-r_j \Delta_j (r_S/r_j)} = \prod_{j=i+1}^{N} \left(\frac{v_j - p_j}{v_j - p_{j-1}} \right)^{r_S/r_j}.$$

Thus, seller profits are

$$\Pi(p_1, \ldots, p_N) = \sum_{i=1}^{N} \lambda_i p_i \prod_{j=i+1}^{N} \left(\frac{v_j - p_j}{v_j - p_{j-1}} \right)^{r_S/r_j},$$

The beginning price p_N appears in all of the summation terms, since it not only determines profits at time 0, but also the time interval to p_{N-1} and consequently the time to any subsequent price. The next price p_{N-1} also appears in all terms but the Nth (since it has no discounting), and one can see that the price p_i appears in the first i terms for $1 \leq i \leq N$. In economic terms, this says that for any time at which a new type purchases, future prices beyond the immediately proceeding price do not affect profits for the current time. That is, future profits can be determined independently of the price choice today. This allows us to solve recursively for optimal prices p^*, where $p_1^* = v_1, p_2^*$ as a function of p_1^*, p_3^* as a function of p_2^* and p_1^*, and so on.

Here are the details of the recursion. Since we know that $p_1 = v_1$, we will turn attention to p_2, the price immediately preceding the lowest price. The only terms p_2 appears in are the first 2 terms:

$$\left[\lambda_1 v_1 \left(\frac{v_2 - p_2}{v_2 - v_1} \right)^{r_S/r_2} + \lambda_2 p_2 \right] \left(\frac{v_3 - p_3}{v_3 - p_2} \right)^{r_S/r_3} \prod_{j=4}^{N} \left(\frac{v_j - p_j}{v_j - p_{j-1}} \right)^{r_S/r_j},$$

and the objective of maximizing over p_2 is

$$R_2(p_1^*, p_2) = \frac{\lambda_2 p_2 + \lambda_1 p_1^* \left(\frac{v_2 - p_2}{v_2 - p_1^*} \right)^{r_S/r_2}}{(v_3 - p_2)^{r_S/r_3}} = \frac{\lambda_2 p_2 + (v_2 - p_2)^{r_S/r_2} R_1(p_1^*)}{(v_3 - p_2)^{r_S/r_3}}, \quad (3)$$

multiplied by a positive term that is constant with respect to p_2, so maximizing R_2 in (3) with respect to p_2 yields p_2^*, which is a function of primitives and p_1^*, with the restriction that $p_2 \geq v_1$. Recursively, define

$$R_i(p_1^*, \dots, p_{i-1}^*, p_i) = \frac{\lambda_i p_i + (v_i - p_i)^{r_S/r_i} R_{i-1}(p_1^*, \dots, p_{i-1}^*)}{(v_{i+1} - p_i)^{r_S/r_{i+1}}}. \tag{4}$$

As we will see, R_i is almost the profits from selling to types 1 to i, discounted to the time when i buys. So we can think of R_i as maximizing profits by selling only to the first i types. Note that R_i is a linear combination of profits from selling to i and the rest of the profits from selling to $i-1$ to 1:

$$R_i = \frac{1}{(v_{i+1} - p_i)^{r_S/r_{i+1}}} \lambda_i p_i + \frac{(v_i - p_i)^{r_S/r_i}}{(v_{i+1} - p_i)^{r_S/r_{i+1}}} R_{i-1}.$$

The first term is increasing in p_i and the second term is decreasing in p_i, highlighting the tradeoff between getting higher profits from charging a higher price to i and forcing all sales to lower types to come at a later time.

Proposition 3. The prices $(p_1^*, \dots, p_N^*)$ maximize Π if and only if p_i^* maximizes $R_i(p_1^*, \dots, p_{i-1}^*, p_i)$ over p_i for $2 \leq i \leq N$ and $p_1^* = v_1$.

Proof. We first show by induction that

$$\sum_{i=1}^{k} \lambda_i p_i \prod_{j=i+1}^{N} \left(\frac{v_j - p_j}{v_j - p_{j-1}} \right)^{r_S/r_j} = R_k(p_1, \dots, p_k)(v_{k+1} - p_{k+1})^{r_S/r_{k+1}} \prod_{j=k+2}^{N} \left(\frac{v_j - p_j}{v_j - p_{j-1}} \right)^{r_S/r_j}. \tag{5}$$

We have already shown this for the base case R_2. Assume the induction hypothesis for $k \geq 2$. Then

$$\sum_{i=1}^{k+1} \lambda_i p_i \prod_{j=i+1}^{N} \left(\frac{v_j - p_j}{v_j - p_{j-1}} \right)^{r_S/r_j} = \sum_{i=1}^{k} \lambda_i p_i \prod_{j=i+1}^{N} \left(\frac{v_j - p_j}{v_j - p_{j-1}} \right)^{r_S/r_j} + \lambda_{k+1} p_{k+1} \prod_{j=k+2}^{N} \left(\frac{v_j - p_j}{v_j - p_{j-1}} \right)^{r_S/r_j}$$

By substituting the induction hypothesis in for the first term and simplifying we have

$$
\begin{aligned}
&\left[R_k(p_1,\ldots,p_k)(v_{k+1}-p_{k+1})^{r_S/r_{k+1}}+\lambda_{k+1}p_{k+1}\right]\prod_{j=k+2}^{N}\left(\frac{v_j-p_j}{v_j-p_{j-1}}\right)^{r_S/r_j}\\
&=\frac{\lambda_{k+1}p_{k+1}+(v_{k+1}-p_{k+1})^{r_S/r_{k+1}}R_k(p_1,\ldots,p_k)}{(v_{k+2}-p_{k+1})^{r_S/r_{k+2}}}(v_{k+2}-p_{k+1})^{r_S/r_{k+2}}\prod_{j=k+2}^{N}\left(\frac{v_j-p_j}{v_j-p_{j-1}}\right)^{r_S/r_j}\\
&=R_{k+1}(p_1,\ldots,p_{k+1})(v_{k+2}-p_{k+1})^{r_S/r_{k+2}}\prod_{j=k+2}^{N}\left(\frac{v_j-p_j}{v_j-p_{j-1}}\right)^{r_S/r_j}\\
&=R_{k+1}(p_1,\ldots,p_{k+1})(v_{k+2}-p_{k+2})^{r_S/r_{k+2}}\prod_{j=k+3}^{N}\left(\frac{v_j-p_j}{v_j-p_{j-1}}\right)^{r_S/r_j},
\end{aligned}
$$

and by induction, equation (5) is proved.

First assume that $(p_1^*,\ldots,p_N^*)$ maximize Π. We already know that p_1^* must be equal to v_1 by Corollary 6. Considering $i\geq 2$, we know that only the first i terms of the profit function depend on p_i. Therefore the arg max of Π given the other optimal prices, with respect to p_i, is the same as the arg max of $R_i(p_1^*,\ldots,p_{i-1}^*,p_i)$ by equation (5) and the fact that the factors

$$(v_{i+2}-p_{i+2})^{r_S/r_{i+2}}\prod_{j=i+3}^{N}\left(\frac{v_j-p_j}{v_j-p_{j-1}}\right)^{r_S/r_j}$$

are constant with respect to $p_1,\ldots,p_i$. Thus the "only if" direction is proved.

Now assume $p_1^*=v_1$ and p_i^* maximizes $R_i(p_1^*,\ldots,p_{i-1}^*,p_i)$ over p_i for $i\geq 2$. Proceeding by induction, we know that $p_1^*=v_1$ maximizes Π given any other prices. We proceed by induction to show that $(p_1^*,\ldots,p_i^*)$ maximize Π given any prices $(p_{i+1},\ldots,p_N)$. Assuming $(p_1^*,\ldots,p_i^*)$ are optimal for any other prices, we know that p_{i+1}^* is optimal since maximizing Π with respect to p_{i+1} is equivalent to maximizing R_{i+1}, by (5), and since prices up through p_i^* are optimal for any other prices, $(p_1^*,\ldots,p_{i+1}^*)$ is optimal for any prices $(p_{i+1},\ldots,p_N)$. Then by induction, $(p_1^*,\ldots,p_N^*)$ is maximizes Π.

□

Proposition 3 tells us that past optimal prices do not influence future optimal prices from the seller's perspective. This result mirrors Stokey (1979), which notes that changing the optimal pricing strategy for some initial length of time doesn't change the optimal pricing strategy over the remaining length of time. This is intuitive for the seller: given that they face a certain composition of buyers in the market, their strategy for intertemporal price discrimination on that set of buyers is not influenced if the time is t or t', due to multiplicative discounting.

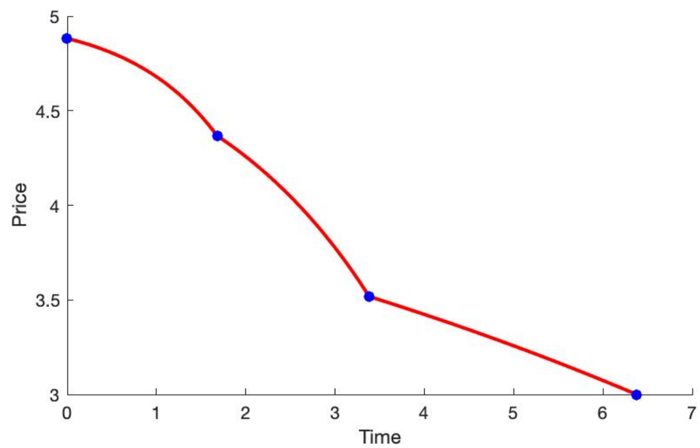

Figure 2: An optimal price path in the single-arrival model, with 4 single-dimensional types in varying discount rates. Blue points represent prices where sales actually occur to a new type. The model parameterization used is $r = .01$, $R = (.1, .1, .5, 1)$, $v_H = 5$, $v_L = 3$, $\lambda = (10, 1.5, .5, 3)$.

This result also gives rise to an interesting asymmetry in the importance of different buyer types in the market. Since we are assuming single-dimensionality in order to cleanly order the buying times of the types, we know that either the highest willingness-to-pay types or the most impatient types buy first (high types), and the lowest willingness-to-pay types buy last (low types). Variation in the level or amount of the high types does not lead to price changes for the lower types, but it may change the amount of time they have to wait until they buy the good. On the other hand, variation in the level or amount of low types can lead to both pricing and purchase timing changes for higher types.

4.3 Comparative Statics and Welfare

We now consider what happens in the model when we vary parameters and what influence this has on other related variables of interest. While dynamic price patterns and explaining their variation over time is interesting in its own right, economists and policymakers are often more interested in what welfare effects these price patterns have, and whether or not efficiency is achieved. We investigate the comparative statics and welfare now in the single-arrival model with single-dimensional types.

From the previous analysis, we know that, in general, not all types buy at the beginning of the game, and that the purchase times of types are ordered, with higher types buying before lower types. The lowest type always receives utility equal to the outside option, which is normalized to be zero. Prices at times which a new (lower) type buys are also decreasing over time. We also established the asymmetry in pricing for the various buyer types: changes to optimal pricing for the higher types only affects the lower types that purchase later in terms of their purchasing times, not in their purchasing price, whereas changes in the optimal pricing for lower types affects both the price and the timing at which higher types will buy.

Consider some type (v_i, r_i). Consider what happens to the optimal price path when we increase or decrease the mass of buyers of this type. If we increase the mass of buyers (v_i, r_i), this means that selling more aggressively to these types is now more attractive to the seller, which means on the optimal price path, types (v_i, r_i) now pay a higher price. In order to get type (v_i, r_i) to pay a higher price, the lower future prices must get pushed out to later dates. So lower types are unambiguously worse off, since they buy at the same price at a later date. With a higher λ_i, higher prices are offered to type (v_i, r_i), but do they imply a higher or lower surplus for higher types? On the one hand, it seems that it may lead to lower prices for the higher types so that the seller can sell at an earlier time to type i, and on the other hand by charging higher prices to i, for the same timing it looks like the seller can charge higher prices to higher types. The change in timing due to the price change in p_i is entirely due to incentives of type $i+1$, the next highest type. Treating the R_i and R_{i+1} functions defined in (4) as partial profits, we have the insight that the profit R_i is now higher, leading to a higher profit R_{i+1} for any value of p_{i+1}, and the equilibrium p_{i+1} is lower, as long as we have an interior solution, i.e., $p_{i+1} > p_i$. Continuing recursively, all R_j are now higher for j a higher type than i, meaning p_j in equilibrium is lower. As λ_i gets larger, however, eventually the constraint $p_{i+1} \geq p_i$ binds, and when that occurs, p_{i+1} increases with λ_i, and eventually higher types will have binding price constraints and will increase with λ_i as well. We have the following result.

Proposition 4 (Comparative Statics in Market Shares)**.** In the single-arrivals model with single-dimensional types, for a given type (v_i, r_i), as λ_i increases,

(a) p_i is increasing, and t_i is decreasing.

(b) p_j is fixed for $(v_i, r_i) \succeq (v_j, r_j)$, and t_j is increasing.

(c) p_j is decreasing for $(v_j, r_j) \succeq (v_i, r_i)$ as long as the constraint $p_j \geq p_{j-1}$ doesn't bind. If $p_j = p_{j-1}$, then p_j is increasing. The timing of pricing t_j is decreasing.

The economic dynamics of this problem thus give rise to the following intuition: as lower buyer types increase in the share of the market, profits from selling to the low types increase, thus making it more attractive for the seller to sell to the lower types at an earlier time. Earlier sales times means that the seller needs to charge lower prices to the high types, up until the lower price bounds on the high types are binding, at which point they become "lumped in" with the low types and face rising prices as the share of the low types continues to grow. Types lower than the type with the growing market share are unambiguously always worse off; they face the same prices but receive the good at later times.

In terms of welfare, when types purchase at earlier times in the game, there is a larger amount of total surplus available, since both the buyers and seller lose payoffs from the transaction as time wears on. From the comparative statics result, we can directly tell that lower types are losing surplus as there gets to be more higher types, but it is not immediately clear that higher types lose or gain surplus when lower types increase in market share. It turns out that strictly higher types will gain as long as their price constraint does not bind, since their price and time of purchase are both decreasing. Once the pricing constraint constraint binds, however, they become grouped with the type with a growing market share and face a price that is increasing but an earlier purchase time. This leads to a distinction between the models where types vary along the willingness-to-pay dimension or the discount rate dimension. In the former, the price charged to a lower willingness-to-pay type maxes out at a value strictly less than the higher type's willingness-to-pay, leading to a minimum level of surplus for the high type given that they buy at the same price as the lower type. Thus, in this case, their surplus is increasing even as they are priced together with lower types, since they purchase at earlier times. In the case where types differ in discount rates, the minimum level of surplus for a higher type is zero, since their willingness to pay is the same, as long as the type with the increasing share is not the lowest type, which has a strictly smaller willingness-to-pay than the other types. In this case, their surplus from buying the good at a price lower than their willingness-to-pay shrinks faster than their benefit of receiving the good earlier.

Finally, in this section, we comment on the efficiency of the seller's pricing scheme. As long as there are some types that purchase at a time strictly later than the start of the game, there is loss of total surplus. Both the seller and the buyers all would strictly prefer to make transactions at time 0, given a fixed price. Under most sets of parameters, the seller imposes some intertemporal price discrimination, and some types buy at a time later than 0, thus creating inefficiency. Thus, the inability of the seller to differentially price based on type at time 0, whether to information asymmetry or restrictions on available pricing schemes, coupled with the ability of the seller to credibly commit to keeping prices high for a long time, destroys surplus. This is in contrast to the well-known result from Coase (1972) on efficient sales of durable good, the difference being due to the commitment power of the seller. This paper is not the first to point out this fact, and it is highlighted in Stokey (1979) under certain conditions as well as Besanko and Winston (1990), Su (2007), and Board (2008). We also note that with a sufficiently high market share of the lowest type when types vary in either dimension leads to all types purchasing at time 0 for a price of v_1.

This is the highest possible total surplus in the game, as well as the highest possible consumer surplus for each buyer type. Even though there is no intertemporal price discrimination, the seller still exercises their monopoly power by charging a price above marginal cost (zero in this case). We conclude that it is the presence of many high types that cause loss of total surplus because the seller is tempted to implement intertemporal price discrimination to destroy some surplus in order to extract higher rents from the high types. The high types suffer from higher prices and the low types suffer from having to wait a long period to buy the good.

4.4 Mechanism Design in the Single-Arrival Model

In the setting of a monopolist with commitment power, it is worth asking why we would restrict the seller to offering a posted price in each period, since there are other conceivable pricing schemes that a seller could use to implement price discrimination and extract rents. There are a few reasons I consider the restriction in this model to posted prices. The first is observational. In practice, it is often seen that, absent information about buyers' types, such as web browser cookies or geographical residence or other information that may be correlated with a buyer's type, we often observe the same posted price offered to all buyers in the market at any given point in time. Another reason is that the seller may have legal restrictions or other considerations that cause them to offer a posted price versus a more general price mechanism. For instance, the seller may be legally required to provide the good at the same price and at the same time to all buyers who purchase at the same time, or the seller may not have the technology or ability to implement something more complex than a posted price, or the seller or buyers may even have a preference for posted prices, since they are easier to understand or seem more fair. The other main reason for considering the restriction to posted prices is for tractability in modeling and analysis. There are various ways to pose the mechanism design problem for a profit-maximizing monopolist in this setting, but most lead to difficulties that do not have a tractable analytical framework and are therefore difficult to draw theoretical conclusions from.

Optimal mechanisms in the intertemporal price discrimination model turn out to be complex and involve randomness in general, even when buyers have utility that is linear in price. This is in stark contrast to the deterministic pricing theorem, which offers a clean restriction of intertemporal price discrimination to a tractable space of pricing schemes that are easy to interpret and understand. Consequently, the question of whether or not the space of price processes achieves the same profits as general mechanisms do is answered in the negative. This contrast between tractability of posted prices and general mechanisms is part of the motivation for investigating whether or not posted prices should follow a deterministic path; because of this difference, it is not immediately obvious that one should rule out stochastic price paths as seller-optimal.

To demonstrate this contrast, I offer a mechanism design framework that applies to the single-arrival setting and show that optimal mechanisms can involve randomness in a way that is not replicable with a price process, even with a minimal number of buyer types with utility that is linear in price.

Consider the setting with constant arrivals and 3 types that vary in patience, so that $\Theta = \{(v_H, r_2), (v_H, r_1), (v_L, r_0)\}$, with $v_H > v_L > 0$ and $r_2 > r_1 \geq r_0 \geq r_S > 0$, with $\lambda_2, \lambda_1, \lambda_0$ the masses of buyer types who are present in the initial time of the market, respectively. We refer to type (v_H, r_2) as *impatient*, (v_H, r_1) as *patient*, and (v_L, r_0) as *low-value*. A buyer's type is their private information and unknown to the seller, as before. A *mechanism* is a pair of functions (τ, p), that map a space of messages M to random variables over the positive reals $\mathbb{R}_+$ and the reals $\mathbb{R}$, respectively. The random variables[4] τ and p may be jointly dependent. The interpretation of such a mechanism is that the buyer reports a message and the seller then provides the good at time τ at which time the buyer pays the random price p to the seller. It is helpful to point out that since there is a mismatch between the seller's rate of time preference and at least one buyer in price, the seller can do better by having the buyer make any payments at time 0, but for the sake of making this model comparable to the dynamic pricing model, we require that payment is only made upon delivery of the good. Both the mechanism design and dynamic pricing models can be recast with a single rate of time preference for all sellers and buyers with buyers differing in their rate of time preference for the good and the same qualitative results in this entire section hold.

Given a mechanism, ex-ante buyer utility from reporting a message $m \in M$ when their type is (v, r) is

$$\mathbb{E}\left[e^{-r\tau(m)}(v - p(m))\right].$$

We can apply a version of the revelation principle for Bayesian incentive compatibility (including a participation constraint) and restrict to a message space equal to Θ plus an outside option of 0 and impose the incentive compatibility restrictions all types maximize utility by reporting their own type to the mechanism.

A first observation is that we can write buyer expected utility as

$$\mathbb{E}\left[e^{-r\tau(v,r)}(v - p(v,r))\right] = \mathbb{E}\left[e^{-r\tau(v,r)}(v - \mathbb{E}[p(v,r)|\tau])\right]$$

[4]Assume restrictions on the joint distribution of (τ, p) such that $\mathbb{E}[e^{-r\tau}(v - p)]$ exists for any type (v, r).

so we can without loss of generality let p be deterministic, conditional on τ. We call (τ, p) a *deterministic mechanism* if $\tau(v, r)$ takes a single value. It is clear that an incentive-compatible deterministic mechanism can be implemented as a deterministic price path: simply set the prices in the price path $p_{\tau(v,r)}$ and incentive-compatibility ensures that each type (v, r) buys at time $\tau(v, r)$ for price $p_{\tau(v,r)}$. Thus, maximizing profits over deterministic incentive-compatible mechanisms is equivalent to maximizing profits over deterministic price paths.

Using the machinery developed in Section 4.2 we can solve for the best deterministic mechanism. A consequence of Theorem 7 is that if there is a mechanism that achieves higher profits than the best deterministic mechanism, it must not be implementable as a (stochastic) price path.

The optimal deterministic mechanism has the form $\tau(v_H, r_2) = 0$, $\tau(v_H, r_1) = t_1 \geq 0$, and $\tau(v_L, r_0) = t_2 \geq t_1$, with $p_0 \geq p_{t_1} \geq p_{t_2} = v_L$. The seller's profits are

$$\lambda_2 p_0 + e^{-r_S t_1} \lambda_1 p_{t_1} + e^{-r_S t_2} \lambda_0 v_L.$$

How could this be improved by randomizing over price times? One such way is to offer the good for a price that is close to v_H at time 0 for one of the high-value types with high probability, and with low probability offer the good for free at a later date. The timing of getting the good for free at a later date differentiates the patient and impatient high types in valuing the prospect of getting the good for low cost then. Under certain parameter conditions, we get distinct deterministic prices in the mechanism: $v_H > p_0 > p_{t_1} > v_L$ with $0 < t_1 < t_2$. The incentive constraints that bind will be for the impatient type reporting as patient and the patient type reporting as low-value and the low-value type taking the outside option, which are the typical downward-binding constraints in mechanism design. Thus,

$$\begin{aligned} v_H - p_0 &= e^{-r_2 t_1}(v_H - p_{t_1}) \\ v_H - p_{t_1} &= e^{-r_1(t_2 - t_1)}(v_H - v_L). \end{aligned}$$

Giving the mechanism the same outcomes for the impatient and low-value types, we modify the outcomes for the patient types as follows: let $\tau(v_H, r_1) = t^*$, where $t^* = \tilde{t} > 0$ with probability $\pi \in (0, 1]$ and equal to 0 with probability $1 - \pi$. Define prices

$$p^*(t^*) = \begin{cases} v_H & \text{if } t^* = 0 \\ 0 & \text{if } t^* = \tilde{t} \end{cases}$$

In order to preserve incentive compatibility, certain conditions on π and $\tilde{t}$ must be met. In particular, the downward incentive constraints must hold: impatient should not report as patient, and patient should not report as low-value. Profits will not depend on $\tilde{t}$, but the seller will want to make π as small as possible, meaning $\tilde{t}$ will be as small as possible as well. The downward incentive constraints then must bind, and we obtain the system of equations

$$v_H - p_0 = \pi e^{-r_2 \tilde{t}} v_H \qquad \text{Impatient binding IC}$$
$$\pi e^{-r_1 \tilde{t}} v_H = e^{-r_1 t_2}(v_H - v_L) \qquad \text{Patient binding IC}$$

This system has a unique solution in $(\pi, \tilde{t})$. Solving for π directly, since $\tilde{t}$ doesn't affect seller profits, we have

$$\pi = e^{-\frac{r_1 r_2}{r_2 - r_1} t_2} \frac{1}{v_H} \left(\frac{(v_H - v_L)^{r_2}}{(v_H - p_0)^{r_1}} \right)^{1/(r_2 - r_1)}$$

At least under some parameters, $\pi \in (0, 1)$, making it a valid probability, as well as large enough so that $\tilde{t} > 0$. The downward incentive constraints bind, and the original incentive constraints for the impatient and low-value types are unchanged, so the only other constraint to check is that the patient type would not prefer to report as impatient, which simply follows from the downward constraint of the impatient type: $v_H - p_0 = \pi e^{-r_2 \tilde{t}} v_H < \pi e^{-r_1 \tilde{t}} v_H$.

What we would like to show is that the seller can make higher profits under this mechanism than the deterministic mechanism. Since profits between the two only differ for the impatient types, whether or not this stochastic mechanism leads to higher profits comes down to the difference

$$\begin{aligned} &(1 - \pi) v_H - e^{-r_S t_1} p_{t_1} \\ = v_H - e^{-\frac{r_1 r_2}{r_2 - r_1} t_2} &\left(\frac{(v_H - v_L)^{r_2}}{(v_H - p_0)^{r_1}} \right)^{\frac{1}{r_2 - r_1}} - e^{r_S t_1}(v_H - e^{r_2 t_1}(v_H - p_0)). \end{aligned}$$

By picking parameters $v_H = 7, v_L = 3$, $r_S = .5, r_1 = 3, r_2 = 8$, $\lambda_0 = 4, \lambda_1 = 1, \lambda_2 = 10$, we obtain an optimal deterministic mechanism with $t_1 = .21, t_2 = .93$, $p_0 = 6.91$, $p_1 = 6.53$, with discounted to time-0 profits of 82.57. The stochastic mechanism corresponds to $\pi = 0.066$ and $\tilde{t} = t_1$, with discounted to time-0 profits of 83.23, yielding the small increase of 0.67. Of course, the difference may be made arbitrarily large by scaling.

The intuition that can be offered here is that there is a greater amount of total surplus to be had when the good is traded earlier. When the good is traded at time 0, both the seller and buyer are better off than trading at a later time at the same price, due to the discounting losses. A randomized mechanism allows the seller to price discriminate at time 0 with high probability and not need to rely on intertemporal discrimination to differentiate the different buyer types. In other words, intertemporal discrimination is inefficient because it burns surplus through time discounting.

The non-deterministic mechanism described yields higher profits than any deterministic mechanism, since we started with the optimal deterministic mechanism. However, the given mechanism is not necessarily optimal overall. We do know that any non-deterministic mechanism that performs better than the deterministic mechanism must not be implementable as a price path, due to Theorem 7. So what is it about this mechanism that is non-implementable as a price path? The main issue is that the mechanism offers different prices to different types at the same time. At time 0, there is high probability that the patient type buys at time 0 at a different price from the impatient type. In order for the mechanism to differentiate types with different patience levels, they must have some time discounting component, but the analysis shows that this should be done in a minimal manner, leading to profit increases over the approach that only allows intertemporal price discrimination.

5 Constant-Arrivals Model

In this section, we consider dynamic arrival of buyers and the optimal *deterministic* price path, which by Theorem 8 (extended to continuous-time) will be optimal among stochastic price paths as well. We assume the type space from Example 3 where types vary in patience levels. We restrict to the case where $\Theta = \{(v_H, r_2), (v_H, r_1), (v_L, r_0)\}$ where $v_H > v_L > 0$ and $r_2 > r_1 \geq r_0 \geq r > 0$, r being the seller's discount rate. We denote the arrival rates of the various types as λ_2 being the rate of arrival for (v_H, r_2), λ_1 the rate of arrival for (v_H, r_1), and λ_0 the rate of arrival for (v_L, r_0), which are constant over time. Thus we refer to this model as *constant arrivals*. For now, assume that the exit rate $1 - \mu_t = 0$ at all times.

In this case, deterministic price path is a function $p : [0, \infty) \to \mathbb{R}_+$ that is chosen at time 0 and known to all players at all moments in time. The only other restriction we impose is that p is lower semicontinuous. The left-continuity restriction is used to guarantee that a sequence of decreasing prices actually achieves the infimum, which may be a value at which a given buyer type would purchase. Given that p is deterministic, the stopping problem for each type of buyer is simple: buy at the time τ_t such that

$$\tau_t(v, r_i) = \min\left(\arg\max_{t' \geq t} e^{-r_i t'}(v - p(t'))\right)$$

where the buyer's type is (v, r_i) and t is their time of arrival to the market. If argmax is empty, then they never buy and get utility 0, and we set $\tau_t(v, r_i) = \infty$. Note that if the maximum is achieved at several times, then we break indifference by choosing the smallest t.

A useful quantity in the analysis will be the mass of each type of buyer in the market at any given time t. Indeed, these quantities are the state variables of the dynamic process. These quantities are dependent on the price process p as the stopping times τ are defined as

$$m_t(v, r_i) = \int_0^t \lambda_i \mathbb{1}\{\tau_{t'}(v, r_i) > t\} \, dt'$$

In words, this is the integral of the inflows of buyers of a given type λ_i conditional on these buyers not having purchased yet $(\tau_{t'}(v, r_i) > t)$.

5.1 Properties of Optimal Processes

A first and rather intuitive fact is that an optimal deterministic price process never drops below v_L. This follows from Corollary 6, but the intuition is simple: if $p(t) < v_L$ for some t, consider the process $\tilde{p}(t) = \max\{v_L, p(t)\}$. For all buyers purchasing at a time t such that $p(t) \geq v_L$, they will also do so under $\tilde{p}(t)$, since the argmax is unchanged. For buyers purchasing at $p(t) < v_L$, they will purchase at some $\tilde{t} \leq t$ under $\tilde{p}$. For any $\tilde{t} > t$, $\tilde{p}(t) \leq \tilde{p}(\tilde{t})$, so $e^{-rt}(v - v_L) \geq e^{-r\tilde{t}}(v - \tilde{p}(\tilde{t}))$. Since $v \geq v_L$, the buyer is at worst indifferent between purchasing at v_L or never buying, and we break indifference to buy.

We now split into 2 cases: (i) $p(t) > v_L$ for all t, and (ii) $p(t) = v_L$ for some t. As in the other model variants, we search for profit-maximizing strategies for the monopolist, and call those optimal. We find the optimal deterministic price within each of these cases and compare. Splitting into these cases simplifies the analysis: in case (i), buyers with type $v = v_L$ will never buy since prices always exceed their valuation. Given that low types never buy, the monopolist should simply set prices to equal the valuation of the high types $p(t) = v_H$ for all t. This leads to a discounted profit of

$$\int_0^\infty e^{-rt}(\lambda_1 + \lambda_2) v_H \, dt = \frac{1}{r}[\lambda_1 + \lambda_2] v_H.$$

The more interesting analysis lies in case (ii), where the monopolist sells to the low types, at least sometimes. An optimal price process (if it exists) here has a cyclicality property, which will aid in solving the problem, essentially turning the infinite-horizon problem maximizing profits over all time into a finite-horizon problem of maximizing profits over a finite cycle.

Proposition 5 (Cyclicality of prices in constant arrivals)**.** Suppose p is an optimal deterministic price process and has $p(t') = v_L$ for some t'. Let $T = \inf\{t \geq 0 : p(t) = v_L\}$. There exists an optimal price process $\tilde{p}$ such that $\tilde{p}(t) = \tilde{p}(t + nT)$ for all $n \in \mathbb{N}$ and $t \in [0, T]$.

The intuition for this result is simple: if T is the first time of price hitting its minimum v_L, then everyone remaining in the market at T buys and the market "clears," so that all buyers in the market after that point will be new arrivals. This makes the market look like time 0 again. Since p is optimal over the interval $[0, T]$, the same price path on $(T, 2T]$ should also be optimal.

Proof. Since p is right-continuous and the set $\{p(t) = v_L\}$ nonempty, the infimum is achieved at T: $p(T) = v_L$. Define the price process pointwise by $\tilde{p}(t) = p(t)$ for $t \in [0, T]$, $\tilde{p}(t) = p(t - T)$ for $t > T$. Then $p(t) = \tilde{p}(t)$ for $t \in [0, T]$. Clearly, each buyer who arrives in $[0, T]$ will have the same stopping times, since all will buy at or before T when $p(T) = \tilde{p}(T) = v_L$. Hence profits from the time interval $[0, T]$ will be the same for both p and $\tilde{p}$. Denote profits over a given interval I (discounted to the beginning of the interval) as $\pi(I)$ for p and $\tilde{p}(I)$ for $\tilde{p}$. By optimality, we must have that

$$\pi([0, T]) + e^{-rT}\pi((T, \infty)) \geq \tilde{\pi}([0, T]) + e^{-rT}\tilde{\pi}((T, \infty))$$

which implies $\pi((T, \infty)) \geq \tilde{\pi}((T, \infty))$. Suppose to the contrary that

$$\begin{aligned}\pi((T, \infty)) &> \tilde{\pi}((T, \infty)) \\ &= \tilde{\pi}((T, 2T]) + e^{-rT}\tilde{\pi}((2T, \infty)) \\ &= \tilde{\pi}([0, T]) + e^{-rT}\tilde{\pi}((T, \infty)) \\ &= \pi([0, T]) + e^{-rT}\pi((T, \infty)).\end{aligned}$$

The third and fourth equalities follow directly from the construction of $\tilde{p}$. Hence, using the price $\widehat{p}(t) = p(t + T)$ produces a greater profit than p, contradiction. So $\pi((T, \infty)) = \tilde{\pi}((T, \infty))$ and we can replace p with $\tilde{p}$. Repeating this argument we can replace p on the interval $(nT, (n + 1)T]$ for any $n \in \mathbb{N}$ with p on $(0, T]$ and preserve optimality. □

Letting π_T be the (discounted to time 0) profit from the interval 0 until the first low price time ($p(T) = v_L$) from the above proposition, we have that for an optimal price process, profits equal to

$$\pi([0, \infty)) = \sum_{n=0}^{\infty} e^{-rnT}\pi_T = \frac{1}{1 - e^{-rT}}\pi_T$$

by summing the geometric series. We then take the following approach to solving for maximal profits: (1) first fix $T > 0$, (2) compute maximal profits subject to $p(T) = v_L$ as the first time p hits v_L, (3) finally maximize $\frac{1}{1-e^{-rT}}\pi_T$ over T.

5.2 Recursive Exponential Sales

Fix T. At time T, we have $p(T) = v_L$. The low types v_L will not buy until time T, therefore contributing $e^{-rT}[\lambda_1 + \lambda_2]v_L$ to profits. This leaves the high types (v_H, r_1) impatient and (v_H, r_2) patient that may buy at a date earlier than T. Suppose that a buyer of type (v_H, r_2) is deciding between buying at time t or T. The price at which they are indifferent between the two is

$$e^{-r_2 t}(v_H - p_2(t)) = e^{-r_2 T}(v_H - v_L)$$

and solving for $p_2(t)$ yields

$$p_2(t) = v_H - e^{-r_2(T-t)}\overline{v},$$

where $\overline{v} = v_H - v_L$. Similarly, for the patient type (v_H, r_1), the indifference condition yields

$$p_1(t) = v_H - e^{-r_1(T-t)}\overline{v}.$$

The curves p_1 and p_2 are concave, smooth, and decreasing, with $p_1(t) < p_2(t)$ for all $t \in [0, T)$ and $p_1(T) = p_2(T) = v_L$. These paths will help us solve for the optimal price path. An optimal price path will vary between these two curves in some sense, respecting the strategic timing of purchase of the impatient types. We first show that a certain class of prices, which we call *recursive exponential sales* (RES), are optimal.

Definition 9. A price path $p : [0, T] \to \mathbb{R}_+$ is called a *recursive exponential sale* if there exists a partition of $[0, T]$, $0 = t_0 < t_1 < \cdots < t_n = T$ such that within any interval $(t_i, t_{i+1}]$, for $t \in (t_i, t_{i+1}]$, one of the following holds:

(i) (high price) $p(t) = v_H - e^{-r_2(t_{i+1}-t)-r_1(T-t_{i+1})}\overline{v}$;

(ii) (medium price) $p(t) = p_1(t)$.

A decreasing exponential sale is a price process that respects strategic timing of purchasing decisions of different types. In particular, if p is a "high price" in the definition, the impatient types will be just willing to buy the good at each time t in the interval $(t_i, t_{i+1}]$, and patient types will wait until the end of the interval to buy at $p(t_{i+1}) = p_1(t_{i+1})$. If p is a "medium price," then both impatient and patient types will buy at all instants in the interval $(t_i, t_{i+1}]$. The prices given are the maximum possible prices that retain these properties. Thus a recursive exponential sale follows the reserve prices of one of the types at any given point in time.

Before proving that RES processes are optimal, let us briefly investigate some of its properties. A basic property of a RES is that it clears the market of all types v_H at the end of every interval, while all types v_L accumulate until time T when they buy.

Proposition 6. Let p be a RES with associated partition $t_0 < \cdots < t_n$. Then for $0 \leq i \leq n$,

$$
\begin{aligned}
m_{t_i}(v_H, r_2) &= 0 \\
m_{t_i}(v_H, r_1) &= 0 \\
m_{t_i}(v_L, r_0) &= t_i \lambda_0 \mathbb{1}\{i \neq n\}
\end{aligned}
$$

Proof. At time $t_0 = 0$ the market has not has time to accumulate any mass of buyers, so the result is trivial. Consider therefore $i \geq 1$. The price at time t_i, regardless of the interval type high or medium, is $p(t_i) = v_H - e^{-r_1(T-t)}\overline{v}$. All future prices of the RES t' until time T will lie at or above $p_1(t') = v_H - e^{-r_1(T-t')}\overline{v}$. We need not consider prices beyond time T, since $p(T) = v_L$ is the minimum price from Corollary 6.

First consider the high impatient types $(1, 2)$. In terms of utility discounted to t_i,

$$
\begin{aligned}
v_H - p(t_i) &= e^{-r_1(T-t)}\overline{v} \\
&> e^{-r_2(t'-t)-r_1(T-t')}\overline{v} \\
&= e^{-r_2 t'}(v_H - p_1(t'))
\end{aligned}
$$

since $e^{-r_2(t'-t)} < e^{-r_1(t'-t)}$. The discounted utility at time t' cannot exceed the last expression on the RHS, hence the high impatient types should buy at t if they haven't already.

Now consider the high patient types $(1, 1)$. The utility comparison for them is

$$
v_H - p(t_i) = e^{-r_1(T-t)}\overline{v} = e^{-r_1(t'-t)-r_1(T-t')}\overline{v} = e^{-r_1 t'}(v_H - p_1(t')).
$$

Again, discounted utility of type $(1, 1)$ at time t' cannot exceed the last expression, and the minimum buying stopping time (if they haven't purchased already) is t.

For a low type (v_L, r_0), $p(t) > v_L$ for $t \in [0, T)$, so clearly all of these buyers will wait to purchase until T, when $p(T) = v_L$. This achieves 0 discounted utility, which is the best they can do since the minimum price is v_L. Integrating the arrival rate λ_0 from 0 to t_i gives the result until $t_i = T$, at which point all purchase. This yields the desired expression. $\square$

Proposition 7. (Profits with a RES) Fix a RES. Ignoring profits from selling to (v_L, r_0) at T:

(i) A high price interval $(t_i, t_{i+1}]$ will induce only (v_H, r_2) types (impatient) to buy on the interior of the interval, and a mass $(t_{i+1} - t_i)\lambda(1,1)$ of (v_H, r_1) types (patient) to buy at t_{i+1}, leading to profits (discounted to the beginning of the interval t_i)

$$\pi_H(t_i, t_{i+1}) = \lambda(1,2)G_2(t_i, t_{i+1}) + \lambda(1,1)G_0(t_{i+1})$$

where

$$G_2(t_i, t_{i+1}) = \frac{v_H}{r}(1 - e^{-r(t_{i+1}-t_i)}) - \frac{\overline{v}}{r_2 - r}e^{-r_1(T-t_{i+1})}(e^{-r(t_{i+1}-t_i)} - e^{-r_2(t_{i+1}-t_i)})$$
$$G_0(t_{i+1}) = (t_{i+1} - t_i)e^{-r(t_{i+1}-t_i)}(v_H - e^{-r_1(T-t_{i+1})}\overline{v});$$

(ii) A medium price interval $(t_i, t_{i+1}]$ will induce both patient and impatient high types to buy over the whole interval, leading to profits (discounted to the beginning of the interval t_i)

$$\pi_M(t_i, t_{i+1}) = [\lambda_1 + \lambda_2]G_1(t_i, t_{i+1})$$

where

$$G_1(t_i, t_{i+1}) = \frac{v_H}{r}(1 - e^{-r(t_{i+1}-t_i)}) - \frac{\overline{v}}{r_1 - r}e^{-r_1(T-t_{i+1})}(e^{-r(t_{i+1}-t_i)} - e^{-r_1(t_{i+1}-t_i)}).$$

Proof. We first show that patient and impatient buyers decide to purchase at the times claimed. Then we derive the expressions for profits.

Consider a high price interval. For the high impatient types, they face at time $t \in (t_i, t_{i+1}]$ $v_H - p(t) = e^{-r_2(t_{i+1}-t)-r_1(T-t_{i+1})}\overline{v} = e^{-r_2(t'-t)}e^{-r_2(t_{i+1}-t')-r_1(T-t_{i+1})}\overline{v} = e^{-r_2(t'-t)}(v_H - p(t'))$ for $t' \in (t, t_{i+1}])$, so they are indifferent between buying at t or any time in the interval $(t_i, t_{i+1}]$. After t_{i+1}, price decays at most according to p_1, which decreases slower than the discounting by r_2. Hence the stopping time of a buyer with type $(1,2)$ is their arrival time in the interval $(t_i, t_{i+1}]$. For the high patient types,

$$v_H - p(t) = e^{-r_2(t_{i+1}-t)-r_1(T-t_{i+1})} < e^{-r_1(t_{i+1}-t)-r_1(T-t_{i+1})}\overline{v} = e^{-r_1(t_{i+1}-t)}(v_H - p(t_{i+1}),$$

so the high patient types should wait until the end of the interval t_{i+1} to buy, by Proposition 6 the high patient types accumulate until t_{i+1} when they buy at $p_1(t_{i+1})$. Low types will clearly not buy until price reaches v_L.

Now consider a medium price interval. The price path follows the indifference curve of the high impatient types, therefore they buy at their earliest arrival time in the interval. The high impatient types will also buy at their earliest arrival, because the impatient type's indifference implies their strict preference to buy immediately.

Now we compute profits discounted to the beginning of the interval. Consider the high price interval first. We have a flow of profits from the high impatient types over the entire interval and the mass of high patient types that accumulate over the interval that buy at t_{i+1}. This results in

$$\begin{aligned}\pi_H(t_i, t_{i+1}) = &\int_{t_i}^{t_{i+1}} \lambda(1,2)e^{-r(t-t_i)}(v_H - e^{-r_2(t_{i+1}-t)-r_1(T-t_{i+1})}\overline{v})\, dt \\ &+ (t_{i+1}-t_i)\lambda(1,1)e^{-r(t_{i+1}-t_i)}(v_H - e^{-r_1(T-t_{i+1})}\overline{v}).\end{aligned}$$

Simplifying the integral and pulling out the λ terms yields the expression for π_H. For the medium price interval, we have a flow of profits from both the high types over the entire interval. This results is

$$\pi_M(t_i, t_{i+1}) = \int_{t_i}^{t_{i+1}} [\lambda(1,1) + \lambda(1,2)]e^{-r(t-t_i)}(v_H - e^{-r_1(T-t)}\overline{v})\, dt$$

and again pulling out the λ terms and simplifying the integral yields the desired expression. □

Let us examine the expressions for profit from setting the high or medium price over a given interval. First, the profits from selling high over an interval are split into two pieces: one from selling to the flow of impatient buyers at the high price over the interval, and one from selling to the mass of patient buyers that accumulate over the interval at the end of the interval. The profits from selling to the patient buyers is their reserve price at the end of the interval, and the profits from selling to the impatient buyers is less than v_H/r, the value from selling at v_H for all time, by an amount that depends on the relative rate of discounting between the impatient type and the seller $r_2 - r$ and the difference in high and low valuations $\overline{v} = v_H - v_L$, discounted at different intervals, one at rate r_2 until t_{i+1}, and one at rate r_1 between t_{i+1} to T. This reflects the reserve price of the patient types at the end of the interval, and the reserve price of the impatient types in the middle of the interval. The profits from pricing at the medium level over the interval yield similar intuition, though the only reserve prices that determine prices are that of the patient types, and only their discount rates appear.

5.3 Optimality of High or Medium Pricing in an Interval

We now address, at least partially, the question of optimality within the class of RES price paths. For fixed T, consider some subinterval of $[0, T]$, say $(a, b]$. If this interval is one of the intervals of a RES, the profits over this interval are independent of the behavior of the price path of the rest of the RES, since prices on the rest of the interval $[0, T]$ lie on or above the indifference curve of the patient type with respect to a price $p_T = v_L$. Therefore, the question of what type of RES pricing to use on this interval can be answered independently, and only depends on Δ_1, the length of the interval $(a, b]$, and Δ_2, the length of the interval $(b, T]$, which is the length of time until the price v_L is reached. Given Δ_1 and Δ_2, then, what is optimal, medium or high pricing? Answering this question is the first step in computing an approximately optimal RES.

Let us reparameterize the profits, by choosing $\Delta_1 = b - a$ and $\Delta_2 = T - b$. In this way, we can rewrite G_0, G_1, G_2 as

$$
\begin{aligned}
G_0(\Delta_1, \Delta_2) &= \Delta_1 e^{-r\Delta_1}\left(v_H - \bar{v}e^{-r\Delta_2}\right) \\
G_1(\Delta_1, \Delta_2) &= \frac{v_H\left(1 - e^{-r\Delta_1}\right)}{r} - \frac{\bar{v}e^{-r_1\Delta_2}\left(e^{-r\Delta_1} - e^{-r_1\Delta_1}\right)}{r_1 - r} \\
G_2(\Delta_1, \Delta_2) &= \frac{v_H\left(1 - e^{-r\Delta_1}\right)}{r} - \frac{\bar{v}e^{-r_1\Delta_2}\left(e^{-r\Delta_1} - e^{-r_2\Delta_1}\right)}{r_2 - r}.
\end{aligned}
$$

First, fix Δ_2 and consider interval profits as a function of Δ_1. As $\Delta_1 \downarrow 0$, profits approach zero for both high and medium pricing, as the mass of buyers over the interval shrinks to zero. On the other hand, as Δ_2 gets large, the difference between G_2 and G_1 goes to zero, reflecting the fact that when a discount is far out in the future, the patient and impatient types have similar reserve prices. However, the difference in G_1 and G_0 does not go to zero, since G_0 reflects sales all at later times than in G_1, albeit at a very similar price when Δ_2 is large. The difference in profits from pricing high or medium is

$$\pi_M(\Delta_1, \Delta_2) - \pi_H(\Delta_1, \Delta_2) = \lambda_1[G_1(\Delta_1, \Delta_2) - G_0(\Delta_1, \Delta_2)] + \lambda_2[G_1(\Delta_1, \Delta_2) - G_2(\Delta_1, \Delta_2)].$$

Since G_0 is the profit from selling to all the patient types at the end of the interval at the patient price p_1, and G_1 is profit selling to the patient types over the course of the interval at the patient price (which exceeds the patient price at the end of the interval), $G_0 < G_1$. Also, since G_1 sells at a price lower than G_2, $G_1 < G_2$. There are a few dynamics of interest.

First, using L'Hopital's rule, we have

$$\begin{aligned}
\lim_{\Delta_1 \to 0} \frac{G_1(\Delta_1, \Delta_2) - G_0(\Delta_1, \Delta_2)}{G_2(\Delta_1, \Delta_2) - G_1(\Delta_1, \Delta_2)} &= \lim_{\Delta_1 \to 0} \frac{\frac{\partial^2}{\partial \Delta_1^2}[G_1(\Delta_1, \Delta_2) - G_0(\Delta_1, \Delta_2)]}{\frac{\partial^2}{\partial \Delta_1^2}[G_2(\Delta_1, \Delta_2) - G_1(\Delta_1, \Delta_2)]} \\
&= \frac{e^{r_1 \Delta_2}\left((r_1 - r)\,\bar{v} + r v_H e^{-r_1 \Delta_2}\right)}{e^{-r_1 \Delta_2}\,(r_2 - r_1)\,\bar{v}} \\
&= \frac{(r_1 - r)\bar{v} + r e^{r_1 \Delta_2} v_H}{(r_2 - r_1)\bar{v}}.
\end{aligned}$$

This quantity is useful because of the fact that

$$\pi_M(\Delta_1, \Delta_2) - \pi_H(\Delta_1, \Delta_2) \geq 0 \iff \Gamma(\Delta_1, \Delta_2) = \frac{G_1(\Delta_1, \Delta_2) - G_0(\Delta_1, \Delta_2)}{G_2(\Delta_1, \Delta_2) - G_1(\Delta_1, \Delta_2)} \geq \frac{\lambda_2}{\lambda_1}.$$

This quantity Γ, which I call the *gain ratio*, is the ratio of profit gains from selling to the patient types over the entire interval at their reserve price instead of at the end of the interval $G_1 - G_0$, to the profit gains from selling to the impatient types at their reserve price rather than the patient type's reserve price $G_2 - G_1$. In this limit as $\Delta_1 \downarrow 0$, we see that this ratio is increasing in Δ_2, reflecting the intuition that pricing at the high and medium levels is a small difference when Δ_2 is large, but the gap $G_2 - G_1$ remains regardless of how large Δ_2 is. Thus, when the gain ratio reaches the threshold ratio of arrivals λ_2/λ_1, pricing to the patient becomes more profitable than pricing to the impatient, at least over small intervals. More generally, the gain ratio is increasing in Δ_2 for all values of Δ_1, though the point at which is crosses the threshold ratio will be lower.

Proposition 8. The gain ratio Γ is smooth and convex in (Δ_1, Δ_2) and increasing in each argument.

The proof is a straightforward but computationally tedious verification that Γ has positive partial derivatives and a positive-definite Hessian matrix, whose proof is in the appendix. An immediate consequence of this proposition is that the threshold ratio λ_2/λ_1 is the maximum cutoff that a high price can be optimal. That is, letting Δ_2^* be implicitly defined by $\lim_{\Delta_1 \to 0} \Gamma(\Delta_1, \Delta_2^*) = \lambda_2/\lambda_1$, we have that $\Gamma(\Delta_1, \Delta_2) \geq \lambda_2/\lambda_1$ for any $\Delta_1 > 0$ and $\Delta_2 \geq \Delta_2^*$. The main use of the proposition is in determining an optimal RES: given a specific subdivision of the interval $[0, T]$ for an RES, we can compute whether high or medium pricing for each interval is optimal with Γ.

5.4 Approximately Optimal Recursive Exponential Sales

We now investigate what an approximately optimal RES looks like. The first question is what the partition if the RES should be. What happens as we make the partition of an RES finer and finer? When we split an interval of an RES into 2 subintervals, we effectively do 2 things: (1)

create an intermediate point at which patient types will buy, and (2) shrink the high price curve in either subinterval. The effect of (1) potentially increases profits, while (2) depresses profits. The interesting behavior of optimal RES paths is that we can divide an optimal RES into two phases, the middle pricing interval $[0, T-\Delta_2^*]$ and the high pricing interval $(T-\Delta_2^*, T]$. On the middle pricing interval, prices follow the patient type's indifference curve $p_1(t)$ and are smoothly decreasing. Thus, on this interval, both the patient and impatient types are purchasing the good continuously. When Δ_2^* is reached, high pricing over small intervals dominates medium pricing over small intervals, so that high pricing is used on small intervals, with each subsequent interval larger than the previous. In fact, no RES with a finite partition can be optimal, because any small interval $(T-\Delta_2^*, T-\Delta_2^*+\epsilon]$ should be again subdivided to into two high pricing intervals.

In order to quantify the tradeoff in profits from subdividing an interval, define the function

$$H(\gamma, \Delta_1, \Delta_2) = \pi_H(\gamma\Delta_1, (1-\gamma)\Delta_1 + \Delta_2) + e^{-r\gamma\Delta_1}\pi_H((1-\gamma)\Delta_1, \Delta_2)$$

for the parameter $\gamma \in [0, 1]$. The function H computes profits of dividing the interval of length Δ_1 into subintervals of size $\gamma\Delta_1$ and $(1-\gamma)\Delta_1$ and charging a high price over those intervals. Note that $H(0, \Delta_1, \Delta_2) = H(1, \Delta_1, \Delta_2) = \pi_H(\Delta_1, \Delta_2)$, the profits on the original interval. Thus if H is maximized for a value of $\gamma \in (0, 1)$, a subdivision is optimal. Furthermore, if $\frac{\partial}{\partial\gamma}H(0, \Delta_1, \Delta_2) > 0$, this is enough to ensure an interior solution γ because of the equality of endpoints of H. When $\Delta_1 + \Delta_2 \geq \Delta_2^*$, it must be the case that H satisfies this condition. The intuition for this is that, if we make a small subdivision from $T - \Delta_2^*$, this small interval optimally charges the high price, due to the fact that the gain ratio selects high pricing over medium pricing, but only over a small interval, since the indifference price to the patient type is close to the indifference price for the impatient type, making the loss from waiting to get the patient type to purchase relatively worse.

Since there is always a subdivision of an interval that contains $T - \Delta_2^*$, in the limit the optimal price process has an infinite number of intervals that the seller charges a high price on, with the size of the intervals shrinking towards zero as you approach $T - \Delta_2^*$ from the right. This implies that Δ_2^* is the cutoff point for high prices. Any price that occurs earlier in time must lie on $p_1(t)$, and any later price follows the high pricing for that interval. Since profits over any small interval are small, charging the high price on that interval is not a big change in profits. Therefore, a RES with a finite partition can approximate the optimal price path by charging p_1 before $T - \Delta_2^*$, charging the high price on the interval $[T-\Delta_2^*, T-\Delta_2^*+\epsilon)$, and charging high prices on subsequent intervals until T. Profits on the interval $[T-\Delta_2^*, T-\Delta_2^*+\epsilon)$ are less than the maximum profits by at most[5]

[5]This is not the tightest bound, but it is derived by knowing that $(\lambda_1 + \lambda_2)v_H\epsilon$ is an upper bound on profit, and

$(\lambda_1 + \lambda_2)(v_H - v_L)\epsilon$, so we can approximate the optimal profits with arbitrary precision by making ϵ small.

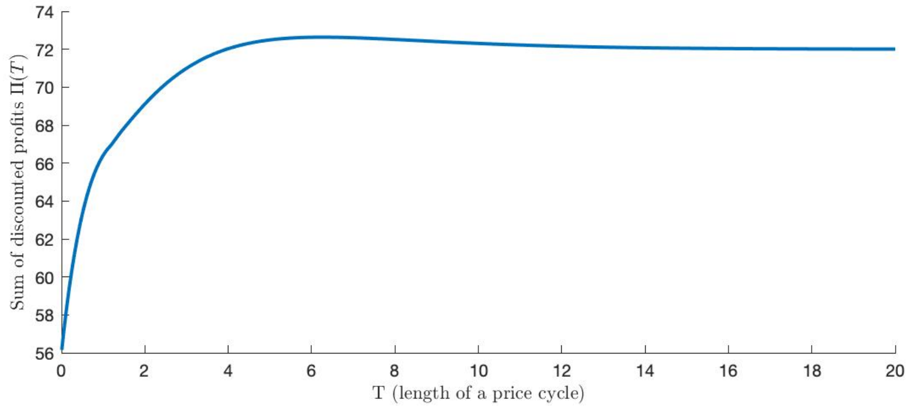

Figure 3: The discounted profit function Π for the same parameterization as Figure 4. The function is single-peaked and asymptotes to the profits $(\lambda_1 + \lambda_2)v_H/r$ from selling at a constant price v_H.

By fixing an ϵ, I can proceed to give a recursive method for computing the optimal RES. Let Δ_2^* be the profit gain threshold value of Δ_2, and choose ϵ ensuring that ϵ is smaller than Δ_2^*, and let $N\epsilon = \Delta_2^*$ for some $N \in \mathbb{N}$, by making ϵ smaller if needed. The algorithm proceeds as follows:

(1) Compute maximum profits for $T = \epsilon$ by setting $\pi^*(\epsilon) = \pi_H(\epsilon, 0) + e^{-r\epsilon}\lambda_0 \epsilon v_L$.

(2) For $k = 2, \ldots, N$: solve for maximum profits for $T = k\epsilon$ by maximizing over possible endpoints of the first interval of the RES, choosing high pricing:

$$\pi^*(k\epsilon) = \max_{l=1,\ldots,k-1} \pi_H(l\epsilon, (k-l)\epsilon) + e^{-rl\epsilon}\pi^*((k-l)\epsilon) + e^{-rk\epsilon}\lambda_0 l\epsilon v_L.$$

(3) For $T > \Delta_2^*$, medium pricing is always optimal on a new interval, so

$$\pi^*(T) = \pi_M(T - \Delta_2^*, \Delta_2^*) + e^{-r(T-\Delta_2^*)}\pi^*(\Delta_2^*) + e^{-rT}\lambda_0(T - \Delta_2^*)v_L$$

Using this method, we can compute approximately optimal π^* for as large of T as desired. The objective by which we evaluate total profits, we recall, is

$$\Pi(T) = \frac{1}{1 - e^{-rT}}\pi^*(T).$$

$(\lambda_1 + \lambda_2)v_L\epsilon$ is a lower bound for pricing high over the entire interval.

The profits π^* are continuous in T. To see this, take any sequence of $\{T_i\}$ that converges to some T will eventually lie in some interval $[T-\delta, T+\delta]$. The maximum possible difference in profits π^* between any two points in this interval is less than $e^{2r\delta}2\delta(\lambda_0+\lambda_1+\lambda_2)v_H$, since the maximum difference in T is 2δ, so the difference in sold to types is a max mass of $2\delta(\lambda_0+\lambda_1+\lambda_2)$, which all must buy at a price lower than v_H, and the difference in discounting is bounded by $e^{2r\delta}$. If the difference were greater, it would imply a greater profit obtained on the types in the optimal price paths that overlap in time, which violates the fact that $\pi^*(T)$ are the maximum profits on the path $[0, T]$. Thus $\pi^*(T_i)$ must converge to $\pi^*(T)$. Furthermore, $\pi^*(T)$ is dominated by the $\pi_M(T-\Delta_2^*, \Delta_2^*)$ term as T gets large, since the latter two terms approach 0 for large T. For large T, π_M is then approximately $(\lambda_1+\lambda_2)\frac{v_H(1-e^{-rT})}{r}$, so we see that $\lim_{T\to\infty}\Pi(T) = v_H(\lambda_1+\lambda_2)/r$, which are the profits from selling only to high-valuation types. This limit is never the best profits achievable by the seller.

Consider a RES that consists of a single interval that sells to the patient type over $[0, T]$, so $p(t) = p_1(t) = v_H - e^{-r_1(T-t)}\overline{v}$. Profits discounted to time zero from this price path are

$$(\lambda_1+\lambda_2)\left[\frac{v_H(1-e^{-rT})}{r} - \frac{\overline{v}}{r_1-r}(e^{-rT}-e^{-r_1T})\right] + e^{-rT}\lambda_0 T v_L.$$

Dividing through by $(1-e^{-rT})$, we have

$$(\lambda_1+\lambda_2)\frac{v_H}{r} + \frac{Ce^{-rT}T - D(e^{-rT}-e^{-r_1T})}{1-e^{-rT}}$$

for $C = \lambda_0 v_L$, $D = \overline{v}/(r_1-r)$. Note that in the limit as $T\to\infty$, this also approaches $(\lambda_1+\lambda_2)\frac{v_H}{r}$. The question is whether the second term is ever positive. This is equivalent to

$$\frac{T}{1-e^{-(r_1-r)T}} > D/C$$

for some T. The limit on the left as $T\to\infty$ is ∞ as long as $r_1 \geq r$, so the second term is positive for large enough values of T. Thus, we conclude that never selling to the low-valuation type is not optimal, since it is strictly dominated by this RES. A further implication is that $\Pi(T)$ is maximized at a finite value of T.

Now that we have analyzed the approximately optimal price paths, let us take a step back and recount what has been accomplished. First, we can compute approximately optimal recursive exponential sales for any price cycle length T algorithmically. The seller optimally splits the price cycle into two sections, where the first section consists of a smoothly decreasing price path along which both the impatient and patient types purchase. In the later section, there are price jumps,

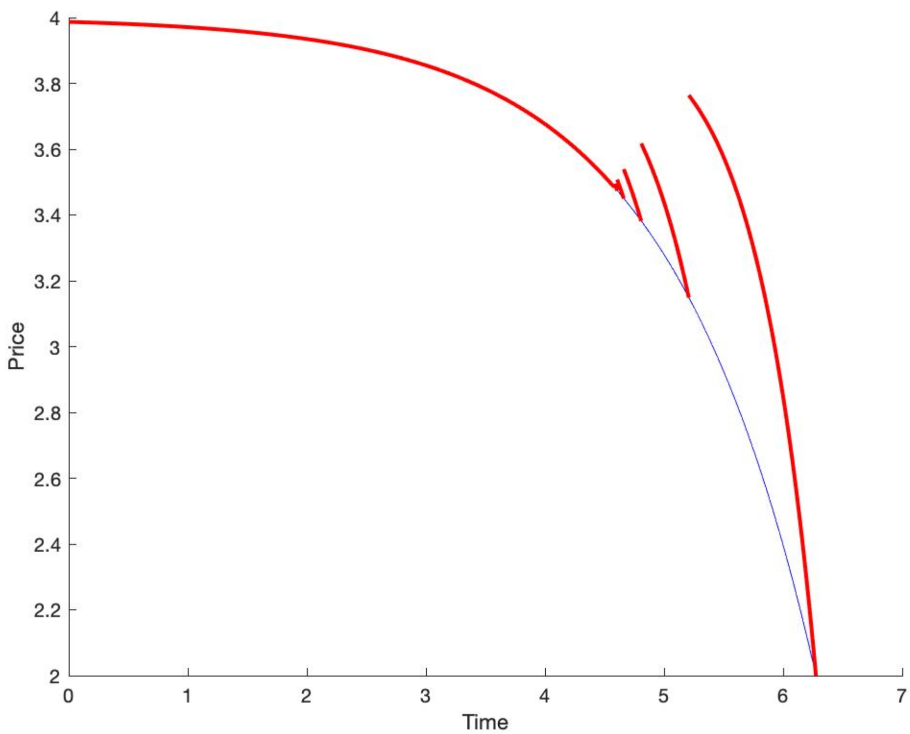

Figure 4: One cycle of the approximately optimal price path under parameters $\lambda_0 = 5$, $\lambda_1 = 2$, $\lambda_2 = 7$, $r = .5$, $r_1 = .8$, $r_2 = 2$, $v_H = 4$, $v_L = 2$. The red lines are the prices, and the blue line traces out the patient type's indifference curve.

along which the impatient types are always purchasing, and the patient types purchase at the end of a decreasing segment, right before the next price jump. The dividing point between smooth and erratic prices is determined by the model primitives and the ratio of impatient arrivals to patient arrivals λ_2/λ_1, and is independent of the length of price cycle length. Thus, as the price cycle gets long, the optimal price path is mostly smooth until near the end of the price cycle, when prices erratically jump. In Figure 4, we see that prices can actually increase over previous price highs following a price jump. This counterintuitive point follows from the fact that separating the patient and impatient types becomes more valuable as the price drop to v_L nears, since the impatient types have a much higher reserve price than the patient types at this point. Pricing to the impatient types can result in a large jump in price if their discount rate is much higher than the patient type's discount rate.

Contrast the constant-arrivals optimal price paths with the resulting optimal price paths from the single-arrivals model. We see that the seller can smoothly separate types in the single-arrivals model by enacting price discrimination, so that each type buys at a separate time for a separate price. This separation is not possible in the constant-arrivals model. The continuous arrival of buyers of all types constrains the seller from extracting the reserve price from every type that purchases at a given moment in time, such as when the seller chooses the medium price, the impatient type receives a strictly higher utility than they would from waiting any longer. The mixed mass of types that buy at different points in time limit the price discrimination abilities of the seller, and it leads to an interesting reversal of the ordering in which types purchase. In the single-arrivals model, the impatient types purchased, followed by the patient types, and the low-valuation types came last. In the constant-arrivals model, the patient types and the impatient types are purchasing continuously for a period of time, and close to the sale the patient types become excluded for small intervals, until the sale to the low-valuation types. The driving force for the erratic prices is the continual arrival of the impatient types, then. If the impatient types all arrived at some early date, the seller could quickly screen them out, then run the constant-arrivals pricing on the patient and low-valuation types without issue, without the need to resort to erratic pricing late in the price cycle.

The differences between constant-arrivals and single-arrival pricing raises an interesting observational question: can we infer the arrival rates, i.e., the dynamic demand, based on observed prices? Indeed, the qualitative differences between constant-arrivals pricing and single-arrival pricing make it possible to distinguish between the two in many cases. The more general question asks whether or not a given optimal price path corresponds to a single set of arrival rates for the model. The answer to this question is no, since the single-arrival pricing may be optimal, for instance, when the low-valuation types arrive over a short period of time (before the time of sale v_L), since this preserves all incentives of participants. In other words, the identifiability of the arrival rates isn't perfect, though it exists to some extent. The question of how well we can identify demand given prices is thus more subtle, and is left for future work.

5.5 Comparative Statics and Welfare in Constant-Arrivals

We now investigate some comparative statics and welfare implications in the constant-arrivals model, with the use of some theoretical analysis as well as simulations. Since there are many different times at which types purchase and exponential discounting is in effect, there is somewhat of a question of how to define consumer surplus. Simply integrating all of the utilities over time neglects the fact that consumers who arrive later will simply have a smaller amount of surplus to gain, due to time discounting. To correct for this, we consider surplus that is time discounted to the time of arrival for a buyer. We leave the arrival rates out since they are a constant multiplier of the surplus. With a RES, the surplus is computable over each interval of the partition:

$$CS_2^H(\Delta_1, \Delta_2) = \int_0^{\Delta_1} (v_H - v_H + e^{-r_2 t} e^{-r_1 \Delta_2} \overline{v}) \, dt = \frac{\overline{v}}{r_2} e^{-r_1 \Delta_2} (1 - e^{-r_2 \Delta_1}) \tag{6}$$

$$CS_1^H(\Delta_1, \Delta_2) = \int_0^{\Delta_1} e^{-r_1 t} (v_H - v_H + e^{-r_1 \Delta_2} \overline{v}) \, dt = \frac{\overline{v}}{r_1} e^{-r_1 \Delta_2} (1 - e^{-r_1 \Delta_1}) \tag{7}$$

$$CS^M(\Delta_1, \Delta_2) = \int_0^{\Delta_1} (v_H - v_H + e^{-r_1 t} e^{-r_1 \Delta_2} \overline{v}) \, dt = \frac{\overline{v}}{r_1} e^{-r_1 \Delta_2} (1 - e^{-r_1 \Delta_1}) \tag{8}$$

where (6) is the surplus for the impatient types under high pricing, (7) for the patient types under high pricing, and (8) is the surplus of both patient and impatient types under medium pricing.

Next, we move onto the intuition of surplus for each type. First, we note that the low-valuation types never achieve any positive amount of surplus, since a price cycle never drops below their valuation. Thus, the timing at which they purchase the good does not influence their utility. On the other hand, it makes sense to impose a weak preference for receiving the good sooner at their willingness-to-pay for the low-valuation types. Second, note that the patient, high-valuation types always buy at a point along their indifference curve, with indifference between buying the good at the current price, or waiting until T to buy the good for price v_L. To this point, the patient type's welfare is solely determined by the length of the price cycle T. Turning attention to the impatient, high-valuation types, we note that their surplus is linear decreasing in the price at every point in time, since they continuously purchase the good.

The main set of comparative statics to investigate are changes in the arrival rates $\lambda_0, \lambda_1, \lambda_2$ of the various buyer types. The first of these, λ_0, is straightforward to analyze. We first note that in the analysis of an approximately optimal RES of a fixed price cycle length T, the value λ_0 does not come into play; whatever RES is chosen, the seller profits from the low-valuation types is constant $e^{-rT}\lambda_0 T v_L$. Thus, λ_0 only determines the cycle length T, and not the optimal price path conditional on T. Clearly, as λ_0 increases, the optimal cycle length T decreases, since selling to the low-valuation types becomes relatively more attractive to the seller. This makes low-valuation types weakly better off in the sense that they receive the good earlier, and the patient types that arrive before T strictly better off. It does not make all patient types better off. If T was the original length of the price cycle, and $T' < T$ is the new length, then there are the patient types arriving between T' and T that would have purchased at a lower price when they arrive to the market because they arrived right before a low price, instead of right after under T'. On the whole, though, more patient types are better off than not. A similar logic applies to the impatient types: most are better off, but some are worse off due to the changed timing of the price cycle. As λ_0 gets large, prices prices converge towards the constant low price v_L, but before then, prices converge to the single-arrivals prices when the patient and low-valuation types are grouped together, due to the last high-pricing interval that occurs in the optimal pricing scheme.

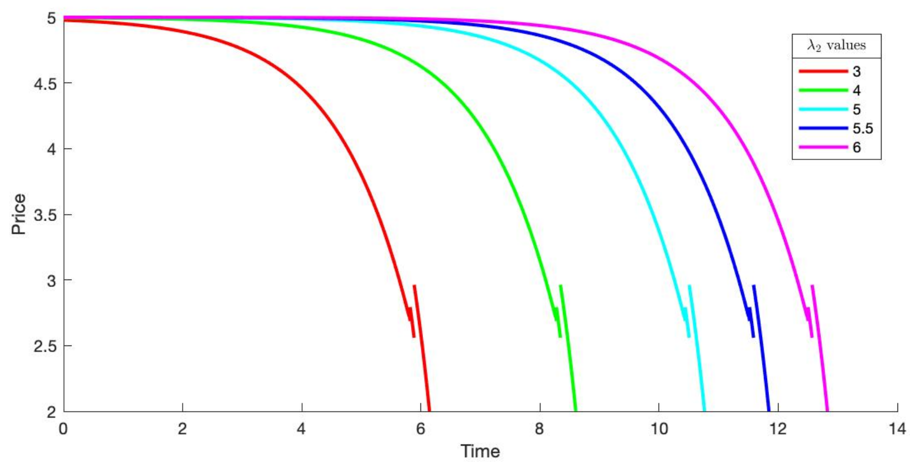

Figure 5: Comparative statics on price paths when λ_1/λ_2 is fixed, with the relative share of low-valuation types decreasing. The threshold Δ_2^* is unchanged, but optimal cycle length is increasing.

Changes in arrival rates of the impatient and patient types with high valuations affect optimal prices in two ways: (1) through changes in the ratio λ_2/λ_1, which changes the Δ_2^* cutoff for medium and high pricing, and (2) through changes in the relative rates of arrival between the high-valuation and low-valuation types, which affect T. The effect of a higher ratio λ_2/λ_1 is a higher Δ_2^*, meaning that the erratic phase of the price cycle is longer, since there are relatively more impatient types present than patient types. On the other hand, higher levels of λ_1 and λ_2 push T to be larger. Some simulations of approximately optimal price paths in Figure 5 show what happens as we keep the ratio λ_2/λ_1 fixed but increase the rates of λ_2 and λ_1 relative to λ_0. Since the ratio λ_2/λ_1 is fixed, the interval of erratic pricing is preserved, we observe the same result as making λ_0 small. We also have results from simulating approximately optimal price paths in Figures 6 and 7 where we vary λ_2/λ_1 but keep one of the relative rates λ_1/λ_0 or λ_2/λ_0 fixed. Fixing λ_1/λ_0 and increasing relative to λ_2, we observe a decreasing T, due to higher λ_0, and a smaller interval of high pricing, due to decreased λ_2/λ_1. Fixing λ_2/λ_0 and increasing relative to λ_1, we still observe a decreasing T, and larger intervals of high pricing, due to increased λ_2/λ_1.

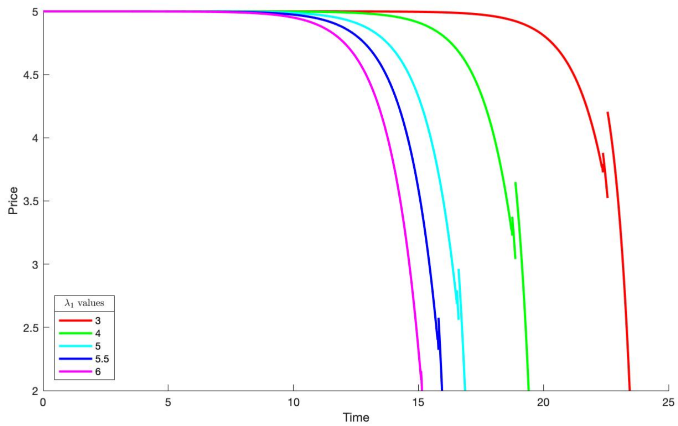

Figure 6: Comparative statics of price paths when λ_1/λ_0 is fixed, with the relative share of the impatient type decreasing.

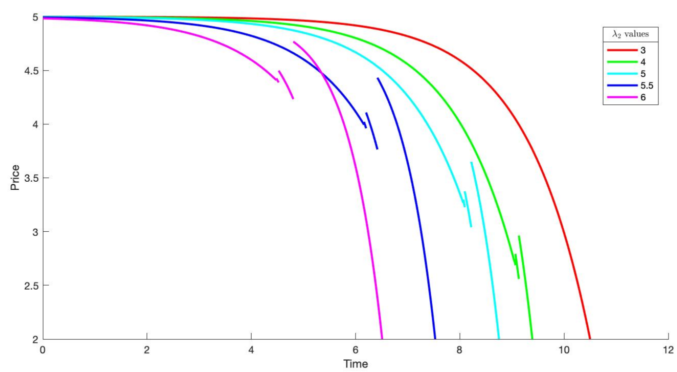

Figure 7: Comparative statics of price paths when λ_2/λ_0 is fixed, with the relative share of the impatient type decreasing.

6 Conclusion

In this paper, we consider the problem of optimally selling a good to a heterogeneous group of buyers who arrive dynamically and are strategic about purchase times. When the types of buyers can be totally ordered by their buying preferences and all buyers have utility functions that are linear in price, then the seller will optimally choose to offer a deterministic price schedule to buyers. With such a set of buyers, the seller has a clear way of pricing the goods to extract maximal rents, which depend on buyer types in a linear way. The linearity in prices assumption means that buyers are risk-neutral, so the seller cannot enforce incentive constraints of the buyer through any moments other than the mean price. Coupled with the ordering on buyer types, this implies that the seller can choose a deterministic price path that maximizes a profit function, which does not depend on any moments of the path of prices. This result points out both the strength and weaknesses of single-dimensional and linear utility assumptions. On the one hand, many price patterns observed in practice are predictable on some level to buyers, such as seasonal sales and Black Friday and Cyber Monday. From a practical perspective, the ability to restrict to deterministic pricing is crucial for tractability in many models of intertemporal price discrimination. This result gives a justification for why optimization over deterministic prices makes sense. On the other hand, many observed price paths seem to exhibit some level of irregularity, as documented empirically by Nakamura and Steinsson (2008) and Klenow and Mailin (2010). The single-dimensional linear-utility model falls short in producing such results, and requires a relaxation of at least one of those assumptions. Indeed, Dilme and Garrett (2018) show a simple version of this framework with buyer risk-aversion produces a memoryless optimal random price path, so that there is essentially no predictability in prices for buyers. Furthermore, this demonstrates a real restriction that is imposed by optimizing over price paths rather than a more general set of pricing mechanisms. We give an example that demonstrates that a seller wishes to price discriminate *within periods* as opposed to *across periods* as much as possible, but price paths allow no within-period price discrimination, and the best across-period price discrimination is deterministic.

The second part of this paper concerns two modeling applications of the deterministic pricing theorem. We consider a model with no dynamic arrivals of buyers, or block arrivals with long periods of no arrivals, called the single-arrivals model, similar to Stokey (1979). We also analyze a model with constant, continuous arrivals of buyers, similar to Conlisk et al. (1984), Sobel (1991), and Board (2008). These models are closely related to many common dynamic price discrimination models, and are tractable as a result of deterministic pricing. We give efficient methods for solving for optimal price paths, and consider the welfare implications and comparative statics of the resulting optimal price paths. The models differ qualitatively in terms of price patterns. Both feature decreasing price paths towards a low price, but the single-arrivals model does so smoothly and separates buyer types across time perfectly. The constant-arrivals model has unevenly spaced price jumps along the way to the low price, and is cyclical with a regular period. Given the differences across these price paths, it seems possible to partially identify the dynamic demand, given an optimal price path. This is left as a question for future research. There are a range of factors that one would wish to consider outside the framework given in this paper in order to estimate demand or operationalize price discrimination, and this paper can be viewed as both a justification and complement to many other investigations of intertemporal price discrimination, offering a more precise insight into the common assumptions in many pricing models. The main differences between these models are (1) they feature qualitatively different price paths: one is smooth, the other containing many jump discontinuities, demonstrating a wide pattern of price behavior that can be captured in the framework, and (2) welfare results, where intermediate types are totally separated from high types in the single-arrival case, but purchase along with high types for most of the price cycle in the constant-arrivals case.

References

Besanko, David and Wayne Winston, "Optimal Price Skimming by a Monopolist Facing Rational Consumers," *Management Science*, 1990, *36* (5), 555–567.

Bitran, Gabriel and Rene Caldentey, "An Overview of Pricing Models for Revenue Management," *Manufacturing & Service Operations Management*, Summer 2003, *5* (3), 203–229.

Board, Simon, "Durable-Goods Monopoly with Varying Demand," *Review of Economic Studies*, 2008, *75* (2), 391–413.

_ **and Andy Skrzypacz**, "Revenue Management with Forward-Looking Buyers," *Journal of Political Economy*, August 2016, *124* (4), 1046–1087.

Bulow, Jeremy, "Durable-Goods Monopolists," *Journal of Political Economy*, April 1982, *90* (2), 314–332.

Chen, Yiwei and Vivek Farias, "Robust Dynamic Pricing with Strategic Customers," *Mathematics of Operations Research*, November 2018, *43* (4).

Cinlar, Erhan, *Probability and Stochastics*, Springer, 2011.

Coase, Ronald, "Durability and Monopoly," *Journal of Law and Economics*, 1972, *15*, 143–149.

Conlisk, John, Eitan Gerstner, and Joel Sobel, "Cyclic Pricing By A Durable Goods Monopolist," *The Quarterly Journal of Economics*, August 1984, *99* (3), 489–505.

Deb, Rahul, "Intertemporal Price Discrimination with Stochastic Values," 2014. Working Paper.

Dilme, Francesc and Daniel Garrett, "A Dynamic Theory of Random Price Discounts," *Working Paper*, September 2018.

Gallego, Guillermo and Garrett van Ryzin, "Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons," *Management Science*, August 1994, *40* (8), 999–1020.

Garrett, Daniel, "Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values," *American Economic Review*, November 2016, *106* (11), 3275–3299.

Gul, Faruk, Hugo Sonnenschein, and Robert Wilson, "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 1986, *39*, 155–190.

Haghpanah, Nima and Jason Hartline, "Multi-dimensional Virtual Values and Second-degree Price Discrimination," *Proceedings of the 16th ACM Conference on Electronic Commerce EC'15*, 2016.

Hart, Sergiu and Philp Reny, "Maximal revenue with multiple goods: Nonmonotonicity and other observations," *Theoretical Economics*, 2015, *10*, 893–922.

Hendel, IIgal and Aviv Nevo, "Intertemporal Price Discrimination in Storable Goods Markets," *American Economic Review*, December 2013, *103* (7), 2722–2751.

Kahn, Charles, "The Durable Goods Monopolist and Consistency with Increasing Costs," *Econometrica*, March 1986, *54* (2), 275–294.

Klenow, Peter and Benjamin Mailin, "Microeconomic Evidence on Price-Setting," *Handbook of Monetary Economics*, 2010, *3* (6), 231–284.

Kreps, David, *Microeconomic Foundations I: Choice and Competitive Markets*, Princeton University Press, 2013.

Landsberger, Michael and Isaac Meilijson, "Intertemporal price discrimination and sales strategy under incomplete information," *The RAND Journal of Economics*, Autumn 1985, *16* (3), 424–430.

Lazear, Edward P., "Retail Pricing and Clearance Sales," *American Economic Review*, March 1986, *76* (1), 14–32.

McAfee, Preston and Vera the Velde, "Dynamic Pricing in the Airline Industry," 2005. Working Paper, California Institute of Technology.

Nakamura, Emi and Jon Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *The Quarterly Journal of Economics*, November 2008, *123* (4), 1415–1464.

Pesendorfer, Martin, "Retail Sales: A Study of Pricing Behavior in Supermarkets," *The Journal of Business*, 2002, *75* (1), 33–66.

Sobel, Joel, "Durable Goods Monopoly With Entry Of New Consumers," *Econometrica*, September 1991, *59* (5), 1455–1485.

Stokey, Nancy, "Intertemporal Price Discrimination," *The Quarterly Journal of Economics*, August 1979, *93* (3), 355–371.

_ , "Rational expectations and durable goods pricing," *The Bell Journal of Economics*, Spring 1981, *12* (1), 112–128.

Su, Xuanming, "Intertemporal Pricing with Strategic Customer Behavior," *Management Science*, May 2007, *53* (5), 726–741.

Varian, Hal, "A Model of Sales," *American Economic Review*, September 1980, *70* (4), 651–659.

White, Mark, "A Simple Model of Intertemporal Price Discrimination," *Eastern Economic Journal*, Summer 2004, *30* (3), 487–492.