

Intertemporal price discrimination: dynamic arrivals and changing values

Daniel F. Garrett*

Abstract

We study the profit-maximizing price path of a monopolist selling a durable good to buyers who arrive over time and whose values for the good evolve stochastically. The setting is completely stationary with an infinite horizon. Contrary to the case with constant values, optimal prices fluctuate with time. We argue that consumers' randomly changing values offer an explanation for temporary price reductions that are often observed in practice. (*JEL* D82, L12)

How should a firm price its goods over time? Does it maximize profits by occasionally reducing prices, or by holding them steady? This paper provides a new perspective on these questions by recognizing that buyers' values for many goods change randomly with time.

*Toulouse School of Economics, 21 allée de Brienne, 31015 Toulouse Cedex 6, France (e-mail: daniel.garrett@tse-fr.eu). I am grateful to four anonymous referees, as well as Alessandro Pavan, Daniel Bernhardt, Simon Board, Jeff Ely, Igal Hendel, Alexander Koch, Alessandro Lizzeri, George Mailath, Konrad Mierendorff, Aviv Nevo, Bill Rogerson, Bruno Strulovici, Miguel Villas-Boas and Rakesh Vohra. Thank you also to seminar participants at Aarhus University, Arizona State University, Bocconi University, HEC Paris, New York University, Northwestern University, Paris School of Economics, Toulouse School of Economics, the University of British Columbia, the University of California at Berkeley, the University of Pennsylvania, Washington University in St Louis, the Arne Ryde Memorial Lectures at the University of Lund, the Asia Meeting of the Econometric Society (2012) at the University of Delhi and the 2012 meeting of the SAET at the University of Queensland. The author declares that he has no relevant or material financial interests that relate to the research described in this paper.

We consider the classic durable goods pricing problem. Forward-looking buyers arrive to the market over time and must decide when to purchase a single unit. A well-known benchmark is where buyers’ values remain constant (see Conlisk, Gerstner and Sobel (1984)). In this case, the monopolist maximizes profits by committing to a constant price. If low-value buyers are excluded, they are excluded forever. The reason is simply that reducing prices to sell to low-value buyers “cannibalizes” earlier demand, reducing the prices the seller can charge to high-value buyers at earlier dates.

Contrary to the above benchmark, we consider buyers whose values change over time in response to idiosyncratic shocks. We show that the seller may then profit by occasionally dropping prices and selling to low-value buyers who are yet to purchase. The reason is that high-value buyers anticipate that their values may fall before they have the chance to buy at a discount. In this sense, high-value buyers are less able to take advantage of low future prices, which means they are willing to purchase at a higher price. We characterize the monopolist’s optimal choice of prices over time, and show how to determine the timing of price discounts. While a few papers already study the optimal (full-commitment) price path for buyers whose values change randomly with time (see Conlisk, 1984, Biehl, 2001, and Deb, 2014), ours is the first to demonstrate that repeated discounting can be optimal in such settings.

Our theory provides a natural explanation for temporary price discounts observed in practice (some other possible explanations are reviewed in the related literature section below). Temporary discounts are common among a broad range of goods, including durables (see, for example, Nakamura and Steinsson (2008) and Klenow and Malin (2010)). Figure 1 provides a summary of weekly prices for a model of blender sold by a large US retail chain. Prices jump frequently between the highest price (\$39.99) and lower discounted prices.¹ Warner and Barsky (1995) also provide detailed evidence of tempo-

¹Data are Retail Scanner Data from the Kilts-Nielsen Data Center (while durable goods are not the strength of their data set, they do have detailed data on some household items, such as blenders). Data are available at the weekly level. Figure 1 provides information on prices for one of many models of blender in the data. A large number of stores owned by the same retail chain offer this model for sale each week, and the figure displays upper

rary discounting for various household durables, including a camera and a food processor (see Figures 1 and 2 of their paper).

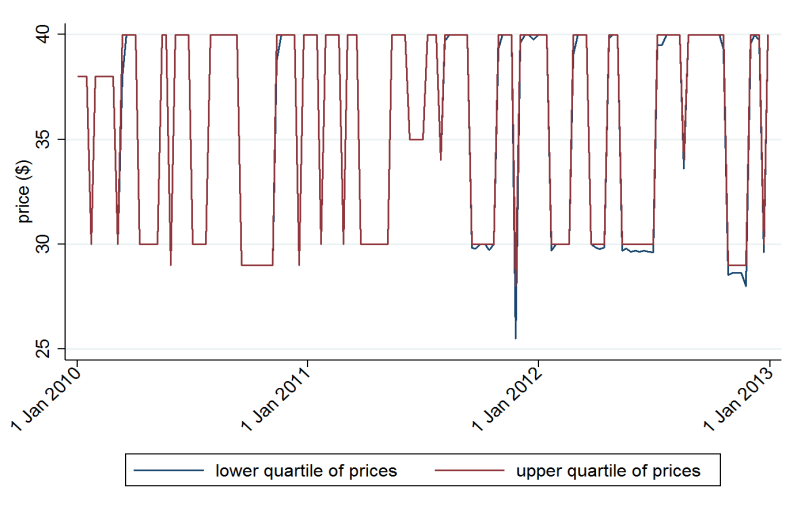


Figure 1: Lower and upper quartiles of prices across stores for a blender sold by a large US retailer. For clarity, prices are shown for the most recent three full years (2010-2012). Source: Kilts-Nielsen Data Center at the University of Chicago Booth School of Business.

Random changes in buyers' values should be expected for various reasons, such as their changing circumstances and unanticipated experiences they happen to encounter over time. Consider the goods mentioned above. A buyer's value for a food processor or a blender may change with the time he spends preparing healthful foods which benefit from their use. This, in turn, may vary with available free time (determined, say, by fluctuating pressures at work) or enjoyment of these foods. A buyer's value for a camera might vary with his interest in photography, in turn reflecting his random interactions with friends. His value may be high if he recently met a friend who enjoys photography, but he would also find it difficult to predict his future values (as he may later become more interested in alternative pursuits).

and lower quartiles of prices across all of this chain's stores selling the good in each week. That these quartiles are almost identical in each week suggests the chain follows a store-wide pricing policy.

The hypothesis of changing values might seem difficult to evaluate empirically. However, quite direct evidence would be available to a researcher who observes how individuals' purchases respond to personal experiences.² Shocks to values also seem a natural way to explain observed patterns in demand, such as purchases on dates when prices are temporarily high. This is the perspective taken by papers that use random utility models for dynamic demand estimation, for instance Gowrisankaran and Rysman (2012).

Our model features buyers who arrive over time and whose values then continue to change between two levels — low and high — according to a continuous-time Markov process. The preference changes are idiosyncratic, and there is no aggregate uncertainty. The environment is also *stationary*, in a sense we define carefully below. The key implication is that, although the values of individual buyers change, the distribution of values in the population does not. Stationarity allows us to focus on the role of individuals' uncertainty about their future values, rather than systematic variation in aggregate demand (see Biehl, 2001, for the same rationale).

The seller commits to a dynamic price path, with the possibility of setting a different price at every date. For a range of parameters, prices cycle. They decline gradually up to their lowest point (a “sale”) before jumping upwards. Buyers purchase immediately if their values are high but wait and purchase at the next sale if their values are low. Holding occasional sales is optimal for the reason explained above. The seller can exploit high-value buyers in between sales dates by charging them high prices, and these buyers are willing to purchase because they are concerned that their values may fall.

The fact that prices decline gradually up to each sale is to be expected, since they are chosen to keep high-value buyers indifferent between purchasing and waiting for the next sale. Waiting for the next sale becomes more tempting as it approaches, due both to discounting and the reduced probability that the value will switch from high to low. Of course, as has often been remarked

²For instance, Iyengar, Han and Gupta (2009) document the effects (both positive and negative) of recent social interactions on decisions to buy, effects which seem difficult to explain absent changing values.

(see Slade, 1998), continuous price changes are not observed in the data. This seems to suggest some costs of price adjustment, although there are other theories such as Kashyap’s (1995) idea that firms can maximize revenues by pricing at nominal thresholds or “price points”. We do not incorporate these considerations in the model.

The logic for occasionally discounting prices rests on downward switches in values; this is what makes high-value buyers reluctant to wait for the next sale. The implications of upward switches in values are different. Upward switches may mean that low-value buyers find it more attractive to delay their purchases, in which case they expect a positive rent. This occurs whenever there are future sales. If a low-value buyer delays his purchase until the next sale, then his value may switch from low to high by this date, and this implies an additional rent when the buyer makes his purchase. The seller can therefore limit rents by spacing sales further apart. Relatedly, upward switches in values make it more attractive for the seller *never* to sell to low-value buyers, by committing to a constant high price. We identify parameters for which repeated price discounting is optimal.

The literature. Our model is most closely related to Conlisk (1984), Biehl (2001) and Deb (2014), who study optimal price paths in durable goods models where buyers’ values change randomly with time. Conlisk and Biehl study two-period models whereas Deb considers an infinite horizon, but permits values to change at most once. In contrast, we study an infinite-horizon setting in which buyers arrive over time, and values continue to change indefinitely. More importantly, none of the earlier papers find repeated price fluctuations, and none describe our logic for delaying sales to low-value consumers. Both in Biehl’s and in Deb’s work, optimal prices are found to be non-decreasing over time. The intuition (which is also relevant in our paper) is that committing to high prices at later dates reduces the rents that must be left to buyers who purchase early on. Conlisk finds that optimal prices may fall, with the seller making sales to high-value buyers in the first period, and to all remaining buyers in the second. This result is driven by the finite horizon. Intuitively, there is a positive option value of not selling to low-valuers at date

1, but no such option value at date 2 (since the game ends at date 2).

The paper should be understood in light of Stokey's (1979) observation that, when buyers' values for a durable good are constant, the seller optimally commits to a constant price. This observation explains the aforementioned finding of Conlisk, Gerstner and Sobel (1984) when buyers arrive over time. Many other theories of time-varying durable goods prices can also be understood in relation to these benchmarks. For instance, Stokey herself shows that a seller may want to gradually reduce prices (and thus delay sales to the lowest types) if values change deterministically with time. She emphasizes that falling values can make high-value buyers effectively "impatient" to consume. In particular, high-value buyers may be reluctant to wait for low prices because their values will have fallen by the time these prices become available. This means the seller can profit by charging high prices to early purchasers. We show that a related intuition can apply when buyers' values are subject to idiosyncratic random shocks, even in an environment where demand is stationary.

Intuition like Stokey's is relevant in a few other settings as well. Landsberger and Meilijson (1985) consider a model in which buyers' values are constant but where buyers have a higher discount rate than the seller. Optimal prices may then decline over time. Low prices at later dates can be profitable because high-value buyers can still be charged relatively high prices early on. The reason is that high-value buyers heavily discount the rents they would earn by waiting for low prices. While our intuition is again related, our model instead posits that buyers share the same discount rate as the seller (for instance, they may have equal access to capital markets). Stokey's intuition also plays a role in Pesendorfer's (2002) work. In his model, high-value buyers are able to stay in the market for only one period, whereas low-value buyers can remain indefinitely. The seller occasionally sells the good to low-value buyers, but charges high-value buyers a price equal to their value at other times.

A number of other papers consider departures from the standard durable-goods environment and find time-varying prices. Board (2008) shows that optimal prices fluctuate when different cohorts of newly arriving buyers have

different demand. In contrast, demand is restricted to be stationary in our model, so price cycles here are not a result of systematic changes in aggregate demand. Board and Skrzypacz (forthcoming) and Gershkov, Moldovanu and Strack (2014) introduce limited capacity and deadlines, showing that optimal prices fall as the deadline approaches, but jump upwards when a unit is sold.³ In contrast, we study an infinite horizon and a seller facing no capacity constraints. Conlisk, Gerstner and Sobel (1984) and Sobel (1991) relax the assumption that the seller commits, and show how this can result in fluctuating prices (whereas the seller commits to the price path in our paper).

There is also a range of theories developed to explain sales, not specific to durable goods. For instance, Varian (1980) develops a theory based on competition and frictions in consumer search, while Maskin and Tirole (1988) develop a theory of dynamic pricing with competition and limited commitment, and show how this can generate “Edgeworth cycles”.

Structure of the paper. The rest of the paper unfolds as follows. Section I introduces the model, Section II examines the optimal price path, and Section III concludes. The Appendix collects omitted proofs.

I Model

Buyers in our model arrive over time. Once in the market, their values change stochastically. They may also leave the market stochastically due to a shock. We now specify our baseline restrictions on the relevant parameters (rates of arrival, exit, and the evolution of values), and postpone describing their role until after our main result (Proposition 1).

Buyers, arrivals and exits. Infinitesimal buyers can make purchases in a continuous-time setting with an infinite horizon. At date zero, a mass $I > 0$ of buyers enters the market. Thereafter, buyers arrive at a constant rate $\lambda > 0$.⁴ Having arrived, buyers receive exogenous shocks at rate $\rho > 0$ causing

³See Horner and Samuelson (2011) for a related environment, but where the seller cannot commit.

⁴Throughout, we take the “intuitive” approach to aggregating random variables over

them to leave the market. These shocks are permanent and render buyers unable to purchase the good forever after. Equivalently for the analysis, they render the good worthless to them. Throughout, we distinguish between buyers *leaving the market*, the result of an exogenous shock, and satisfying their unit demand by making a purchase. We assume $I = \frac{\lambda}{\rho}$. This means that the mass of buyers who have arrived to the market net of those who have left due to an exogenous shock remains constant over time.

Payoffs. The seller and the buyers are risk neutral and have a common discount rate $r > 0$. Buyers have unit demand and the seller's production cost is normalized to zero. If a buyer purchases at date t , his value from the good is $\theta_t \in \{\theta_L, \theta_H\}$, where $0 < \theta_L < \theta_H$. These payoffs might reflect the gratification of instantaneous consumption (for instance, buyers may want to experience consumption only once). Alternatively, they could reflect expected discounted flow payoffs over each buyer's lifetime, accounting for the possibility that these flow payoffs may continue to change in the future. In this case, buyers' expected payoffs from acquiring and holding the good would naturally depend on whether their current flow payoffs are low or high.

Process for values. Upon arrival at a date τ , buyers have a high value (θ_H) with probability $\gamma \in (0, 1)$ and a low value (θ_L) with probability $1 - \gamma$. Values evolve according to a time-invariant continuous-time Markov process. They switch from low to high at rate α_L and from high to low at rate α_H , where $\alpha_L, \alpha_H > 0$. We assume that $\gamma = \gamma^S \equiv \frac{\alpha_L}{\alpha_L + \alpha_H}$. Hence, the process for buyer values is stationary: the probability a buyer has a high value is equal to γ^S , irrespective of the date or time since arrival. Moreover, the stock of all buyers who have arrived and not yet left due to an exogenous shock is comprised at any moment of $I\gamma^S$ high valuers and $I(1 - \gamma^S)$ low valuers. (It is worth emphasizing that these statements concern the primitive environment; the mass of buyers who have arrived and not yet *purchased*, for example, will

infinitesimal buyers, which is standard in such models. While the choice to consider infinitesimal buyers is convenient for exposition, the optimal price path (where prices are permitted to be a function of time alone, as defined below) is identical if we instead take λ to be the Poisson arrival rate of buyers to the market and I to be the expected number of buyers in the market at date zero.

depend in equilibrium on the seller's choice of price path.)

II Analysis

A Price-path mechanisms

The seller commits at date zero to a price path $p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. If a buyer purchases the good at date t , he pays price $p(t)$. Thus there is no role for communication other than of the buyer's purchase decision. We restrict attention to deterministic price paths, but this does not harm profits.⁵

A buyer's problem of when to purchase is an optimal stopping problem. Let Σ be the set of Markov stopping rules, i.e. right-continuous functions $\sigma(\theta_L, \cdot), \sigma(\theta_H, \cdot) : \mathbb{R}_+ \rightarrow \{0, 1\}$ mapping time to a decision to purchase. The set Σ is taken to be the set of feasible buyer strategies. The restriction to Markov strategies is without loss of generality given that the process for buyer values is Markov (and hence does not depend on the arrival time, nor on the realization of past values). For any arrival date τ and any $\theta^{[\tau, \infty)} \equiv \{\theta_s : s \geq \tau\}$, a buyer's purchase date when using the strategy $\sigma \in \Sigma$ is $\mu_\sigma(\theta^{[\tau, \infty)}) \equiv \inf \{s \geq \tau : \sigma(\theta_s, s) = 1\}$, provided he is still in the market at that date.⁶ In case $\mu_\sigma(\theta^{[\tau, \infty)}) = +\infty$, the interpretation is that the buyer never purchases even if he remains in the market.

A "price-path mechanism" $\mathcal{M}_P = \langle p, x \rangle$ includes also the prescription $x \in \Sigma$ of whether buyers are to purchase for each value at each date.⁷ The interpretation is that a buyer is to purchase the good at date t paying $p(t)$ if and only if (a) he has arrived to the market by date t and has not left, (b) he has not yet purchased the good, and (c) his value for the good at date t is θ_t ,

⁵The reason the restriction to deterministic mechanisms is without loss of optimality relates to the linearity of the seller's problem of choosing an optimal allocation rule for conditionally optimal prices. The observation follows from the same arguments as in Strausz (2006).

⁶We find it notationally convenient to view buyers as continuing to draw values after leaving the market, keeping in mind that purchases after such a date are ruled out.

⁷Given that buyers are ex-ante identical and anonymous, it is without loss to consider a prescription x which is the same for all buyers.

with $x(\theta_t, t) = 1$.

Buyers' problem. Suppose a buyer uses the Markov strategy $\sigma \in \Sigma$, and consider his expected payoff at any date t at which he has not yet purchased. This is given, for each value $\theta_t \in \{\theta_L, \theta_H\}$, by

$$u_t^{\mathcal{M}_P}(\theta_t; \sigma) = \mathbb{E} \left[e^{-(r+\rho)(\tilde{\mu}_\sigma - t)} \left(\tilde{\theta}_{\tilde{\mu}_\sigma} - p(\tilde{\mu}_\sigma) \right) \mid \tilde{\theta}_t = \theta_t \right],$$

where $\tilde{\mu}_\sigma$ is the stopping time determined by σ . Here, the expectation is with respect to the evolution of the buyer's value, conditional on θ_t (the probability that the buyer survives in the market until the stopping time $\tilde{\mu}_\sigma$ is accounted for by the factor $e^{-(r+\rho)(\tilde{\mu}_\sigma - t)}$, which also accounts for discounting at rate r). Incentive compatibility of the price path mechanism $\mathcal{M}_P = \langle p, x \rangle$ is then the requirement that, for each t and θ_t ,

$$u_t^{\mathcal{M}_P}(\theta_t; x) = \nu_t^{\mathcal{M}_P}(\theta_t) \equiv \sup_{\sigma \in \Sigma} u_t^{\mathcal{M}_P}(\theta_t; \sigma).$$

Seller's problem. For an incentive-compatible mechanism $\mathcal{M}_P = \langle p, x \rangle$, the profit the seller expects to earn at date τ from a buyer who arrives at that date is

$$\pi^{\mathcal{M}_P}(\tau) = \mathbb{E} \left[e^{-(r+\rho)(\tilde{\mu}_x - \tau)} p(\tilde{\mu}_x) \right].$$

The expectation is now with respect to the unconditional evolution of the buyer's value. The present value of the seller's total profit is then

$$\Pi^{\mathcal{M}_P} = \int_0^\infty \lambda e^{-r\tau} \pi^{\mathcal{M}_P}(\tau) d\tau.$$

The seller's problem is to maximize $\Pi^{\mathcal{M}_P}$ by choice of an incentive-compatible price path mechanism \mathcal{M}_P .

B Optimal prices for a given allocation

To find the optimal price path, we formulate the seller's problem in terms of the allocation rule rather than directly in terms of prices. Analogous to the

textbook analysis of static mechanism design for agents with two types, we anticipate that buyers will behave efficiently if their values are high (i.e., they will purchase the good straight away), while they may not if their values are low (because their purchases may be delayed).

The rent that must be left to buyers will depend on the interaction between two effects. First, buyers earn positive rent if their values are high on any date at which prices are low enough that they are willing to purchase if their values are instead low. Second, buyers anticipate such opportunities to earn rent in advance. Because they expect their values to change, they can therefore expect positive rent even when their values are low.

Sales policies. We begin our analysis by defining a “sales policy”, the set of dates at which buyers purchase if their values are low. For an incentive-compatible mechanism $\langle p, x \rangle$, this is $A = \{t \in \mathbb{R}_+ : x(\theta_L, t) = 1\}$. $\mathcal{M}_{P,A} = \langle p_A, x_A \rangle$ denotes a mechanism with a sales policy A . Without loss of profits for the seller, we consider sales policies which are closed sets and which are countable unions of closed intervals (and points).⁸

A sales policy A can be completely described by a function $m_A : \mathbb{R}_+ \rightarrow \mathbb{R}_+ \cup \{+\infty\}$, such that, at any date t , the next date in A is⁹

$$(1) \quad m_A(t) \equiv \inf \{s \in A : s \geq t\}.$$

Thus, $m_A(t) = +\infty$ simply states that there are no “sales” at or after date t .

Optimal prices for a given sales policy. We now derive, for any sales policy A , the price path which maximizes profits while implementing A . We first deduce lower bounds for the value of a buyer’s problem $\left(\nu_t^{\mathcal{M}_{P,A}}(\cdot)\right)_{t \geq 0}$ in any price path mechanism implementing A . We then provide a mechanism in which the buyer’s payoffs coincide with these bounds, and where the buyer purchases immediately whenever his value is high. Subject to implementing

⁸The price-path mechanism we derive in this subsection provides an upper bound on profits for *any* sales policy. When it comes to characterizing the optimal sales policy in Proposition 1, we maximize over all sales policies, assuming that the bound on profits can be attained. Since the solution to the maximization problem satisfies the aforementioned restrictions, this will indeed be the case for the optimal policy.

⁹Note that because A is closed, it is measurable. Hence m_A is a measurable function.

A, this mechanism both maximizes efficiency and minimizes expected rents, so it maximizes the seller's profits.

We begin with the following observation.

Lemma 1 *The difference between expected rents for a high versus low-value buyer at date t satisfies*

$$(2) \quad \nu_t^{\mathcal{M}_{P,A}}(\theta_H) - \nu_t^{\mathcal{M}_{P,A}}(\theta_L) \geq e^{-(r+\rho+\alpha_L+\alpha_H)(m_A(t)-t)}(\theta_H - \theta_L).$$

Moreover, a low-value buyer's rent satisfies

$$(3) \quad \nu_t^{\mathcal{M}_{P,A}}(\theta_L) \geq \int_t^\infty \alpha_L e^{-(r+\rho+\alpha_L)(s-t)} \nu_s^{\mathcal{M}_{P,A}}(\theta_H) ds.$$

The lower bound on the difference in expected rents in (2) is determined by the difference in expected values for the good at the next sales date, $m_A(t)$, taking into account discounting and the probability of leaving the market by this date. It can be understood using the following thought experiment. Consider two buyers in the market at date t , one with a high value and one with a low value. A lower bound on the difference between their expected rents can be found as follows. Suppose the high-value buyer is not permitted to buy until either one of the two buyers' values changes, or until date $m_A(t)$, whichever is earlier. Clearly, this restriction weakly lowers the high-value buyer's expected rents, but has no effect on the low-value buyer's rents. First, note that, conditional on at least one of the buyer's values changing before $m_A(t)$, there is no difference in their expected rents (consider their expected continuation payoffs at the date of the first change in values). Second, conditional on no change in values before $m_A(t)$, the high-value buyer expects an additional (discounted) rent of at least $e^{-(r+\rho)(m_A(t)-t)}(\theta_H - \theta_L)$ (since the high-value buyer can choose to purchase, just as the low-value buyer does, at date $m_A(t)$). The probability of no change in values is equal to $e^{-(\alpha_L+\alpha_H)(m_A(t)-t)}$.

Equation (3) simply follows because a low-value buyer can do at least as well as to wait until his value turns high, and then follow an optimal continuation strategy. The two equations jointly yield lower bounds on the expected

rents obtained in any mechanism with sales policy A . In particular, solving (2) and (3) with equality, lower bounds on expected rents are given by

$$(4) \quad \nu_t^{\mathcal{M}_{P,A}}(\theta_L) \geq \int_t^\infty \alpha_L e^{-(r+\rho)(s-t) - (r+\rho+\alpha_L+\alpha_H)(m_A(s)-s)} (\theta_H - \theta_L) ds,$$

together with (2). These bounds will coincide with a buyer's actual expected payoffs for an optimal price path, which we give in the next result. This will mean, for instance, that (inspecting (4)) low-value buyers earn rents only because their values can become high in the future (at rate α_L), and because of future sales (i.e., if $m_A(s)$ is finite for $s > t$).

Lemma 2 *Let A be any sales policy. Define, for all t ,*

$$p_A^*(t) = \theta_H - \int_t^\infty \alpha_L e^{-(r+\rho)(s-t) - (r+\rho+\alpha_L+\alpha_H)(m_A(s)-s)} (\theta_H - \theta_L) ds \\ - e^{-(r+\rho+\alpha_L+\alpha_H)(m_A(t)-t)} (\theta_H - \theta_L).$$

Let $x_A^(\theta_L, t) = 1$ if $t \in A$ and $x_A^*(\theta_L, t) = 0$ otherwise, and let $x_A^*(\theta_H, t) = 1$ for all t . Then, $\mathcal{M}_{P,A}^* = \langle p_A^*, x_A^* \rangle$ maximizes the seller's expected profit conditional on implementing the sales policy A .*

The optimal choice of prices for a sales policy A can be completely specified from the following observations. First, at any date $t \in A$, if a buyer's value is low, then he is indifferent between a strategy of purchasing at that date and waiting and instead purchasing if and only if his value turns high. This means that, if there are future sales dates, the price at date t must be less than the low value θ_L . Second, at any date $t \notin A$ at which a buyer's value is high, he is indifferent between a strategy of purchasing at date t and waiting and purchasing with certainty at the next date in the sales policy $m_A(t)$, or never purchasing if there are no future dates (i.e., in case $m_A(t) = +\infty$). Thus, at dates $t \notin A$, prices fall gradually up to the next sales date if there is one, or otherwise remain constant and equal to θ_H .

C Optimal sales policy

Choosing the optimal sales policy. Given optimal prices are determined by Lemma 2, we can focus on deriving the optimal sales policy. We view the seller's problem as a sequence of sub-problems where, at each sales date a , the seller must determine the subsequent sales date (the validity of this approach is verified in the Appendix).

Our dynamic programming argument hinges on showing that the choice of sales policy after any sales date a is separable from the choice of sales policy before a . This follows from a property of the optimal prices found in Lemma 2. To illustrate, consider some sales date a , and suppose at first that the seller chooses no subsequent sales, so that a buyer with a low value at date a has payoff $\nu_a^{\mathcal{M}_{P,A}^*}(\theta_L) = 0$. Consider then the effect of including sales after a on the rent obtained by buyers arriving before a . If $\nu_a^{\mathcal{M}_{P,A}^*}(\theta_L)$ is increased to some value Δ , then the change in the expected payoff of a buyer who arrives at date $\tau < a$ is simply $\Delta e^{-(r+\rho)(a-\tau)}$, where recall r is the discount rate and ρ is the rate of leaving the market. The effect on date-zero profits from all buyers arriving before date a is then equal to

$$\Delta \left(I e^{-a(r+\rho)} + \int_0^a \lambda e^{-r\tau - \rho(a-\tau)} d\tau \right) = \Delta I e^{-ra},$$

where recall λ is the arrival rate and $I = \frac{\lambda}{\rho}$ is the constant population of buyers. Thus, the effect on profits from buyers arriving before a is proportional to the constant population size I , and is independent of the choice of sales dates before a .

This observation facilitates defining the seller's problem of choosing future sales dates after any sales date a . The value of this problem will be the profits earned from buyers arriving after a , less the reduction in profits from buyers arriving before a due to including sales after a . It can be stated recursively as follows. Given a sale at date a , the seller specifies the next sales date $a + z$. The value of the problem at date a comprises (i) the profits earned from buyers arriving between a and $a + z$ in case there are no sales after date $a + z$ (so that

a buyer with a low value at date $a + z$ has payoff zero; i.e. $\nu_{a+z}^{\mathcal{M}_{P,A}^*}(\theta_L) = 0$), (ii) the reduction in profits earned from buyers arriving before a due to including exactly one more sale, at date $a + z$, rather than including no additional sales, and (iii) the discounted continuation value of the problem at date $a + z$.

We now determine the profits the seller expects from a buyer who arrives a length s before the next sale, supposing that the next sale is the last. This is

$$(5) \quad R(s) = \gamma^S \theta_H + (1 - \gamma^S) \left[\theta_H \int_0^s \alpha_L e^{-t(\alpha_L + r + \rho)} dt + \theta_L e^{-s(\alpha_L + r + \rho)} \right] - \gamma^S e^{-s(r + \rho)} (\theta_H - \theta_L).$$

The first two terms are the expected surplus. The first term accounts for the surplus generated when the buyer's value is initially high (which occurs with probability γ^S). The second term accounts for the surplus generated when the buyer's value is initially low. Either the buyer's value jumps from low to high before the sale, precipitating a purchase (this happens at rate α_L and is accounted for by the first term in square brackets), or the buyer's value remains low and he buys at the sale (as accounted for by the second term in square brackets). The final term in (5) is the buyer's discounted expected rent. This is found by observing that (i) the buyer is willing to wait and purchase at the sale (he never has strict incentives to buy beforehand), and (ii) the probability that the buyer's value is high at the sale is γ^S , while, if his value is high at the sale and he does not leave the market due to an exogenous shock, he earns a discounted rent $e^{-rs} (\theta_H - \theta_L)$.¹⁰

Suppose then that there are sales at dates a and $a + z$, but no further sales. The date- a value of expected profits from buyers arriving between a and $a + z$

¹⁰Note that the simplicity of the expression for buyer rents is a result of our stationarity assumption. This implies that the probability a buyer's value is high at the sale is simply γ^S , which does not depend on how long the buyer has been in the market (nor directly on the rates of switching α_L and α_H).

is then

$$(6) \quad \phi(z) = \int_0^z \lambda e^{-r\tau} R(z - \tau) d\tau.$$

This expression simply integrates over the profits from buyers who arrive during this interval at rate λ , as given by (5).

Next, note that, if there is exactly one sale after date a , at date $a + z$, then the expected rent of a buyer with a low value at a is equal to

$$\nu_a^{\mathcal{M}_{P,A}^*}(\theta_L) = \Pr\left(\tilde{\theta}_{a+z} = \theta_H | \tilde{\theta}_a = \theta_L\right) e^{-z(r+\rho)} (\theta_H - \theta_L).$$

This follows because a low-value buyer at date a is indifferent between purchasing at date a and instead waiting and then purchasing at the next sale $a + z$ (if his value happens to be high at date $a + z$, and he has not exited the market, then he earns a discounted rent $e^{-rz}(\theta_H - \theta_L)$). Including the sale at date $a + z$ therefore reduces profits from buyers arriving before a by

$$I \Pr\left(\tilde{\theta}_{a+z} = \theta_H | \tilde{\theta}_a = \theta_L\right) e^{-z(r+\rho)} (\theta_H - \theta_L),$$

where the effect is measured in date- a terms.

We can now state recursively the seller's problem of choosing the next sale after a sales date a . In particular,

$$(7) \quad W^*(a) = \sup_{z>0} \left\{ \begin{array}{l} \phi(z) - I \Pr\left(\tilde{\theta}_{a+z} = \theta_H | \tilde{\theta}_a = \theta_L\right) e^{-z(r+\rho)} (\theta_H - \theta_L) \\ + e^{-rz} W^*(a+z) \end{array} \right\},$$

where $\phi(z)$ is given by (6). The stationarity of the seller's problem is thus clear. Since $W^*(a)$ is independent of a , it is given simply by

$$(8) \quad W^* = \sup_{z>0} \left\{ \frac{\phi(z) - I \Pr\left(\tilde{\theta}_z = \theta_H | \tilde{\theta}_0 = \theta_L\right) e^{-z(r+\rho)} (\theta_H - \theta_L)}{1 - e^{-rz}} \right\}.$$

Equation (8) can be used to determine the optimal time between sales dates,

if any.

Characterization of the optimal sales policy. It turns out that the form of the optimal sales policy can be divided broadly into three classes, depending on the ratio of the high to low values ($\frac{\theta_H}{\theta_L}$). In particular, there are thresholds $\underline{\theta}$ and $\bar{\theta}$, satisfying $1 < \underline{\theta} < \bar{\theta}$, such that the following result holds.¹¹

Proposition 1 *The optimal sales policy A^* is determined as follows.*

Case (i): If $\frac{\theta_H}{\theta_L} \leq \underline{\theta}$, then low-value buyers always buy; i.e., $A^ = \mathbb{R}_+$.*

Case (ii): If $\underline{\theta} < \frac{\theta_H}{\theta_L} < \bar{\theta}$, then low-value buyers buy periodically; there is a unique scalar $z^ > 0$ such that $A^* = \{iz^* : i \in \mathbb{N} \cup \{0\}\}$.*

Case (iii): If $\frac{\theta_H}{\theta_L} \geq \bar{\theta}$, low-value buyers either never buy, or they buy only at date zero; i.e., either $A^ = \emptyset$ or $A^* = \{0\}$.¹²*

The dependence on the ratio $\frac{\theta_H}{\theta_L}$ should be expected. If $\frac{\theta_H}{\theta_L}$ is no more than $\underline{\theta}$, the efficiency loss of delaying sales to low-value buyers is large relative to the reduction in buyer rents. The optimal sales policy A^* is all of \mathbb{R}_+ , which means the seller commits to a constant price below θ_L . Conversely, if $\frac{\theta_H}{\theta_L}$ is at least $\bar{\theta}$, the rents ceded to buyers by holding sales is large relative to efficiency gains, so the seller chooses either no sales date or a single sales date at zero (either possibility may arise, as explained below). In either case, the price is constant at θ_H at all dates after zero. If a sale is included at date zero, the price at this date is θ_L .

Case (ii) is intermediate between the two extremes, and is the focus of our attention. Sales then occur a constant length $z^* > 0$ apart, with the first sale

¹¹Let $\kappa = \frac{r+\rho+\alpha_L}{r+\rho+\alpha_L+\alpha_H}$. Then $\underline{\theta}$ and $\bar{\theta}$ satisfy

$$\underline{\theta} = \frac{\frac{\alpha_L}{\rho} + \gamma^S + \kappa(1 - \gamma^S)}{\frac{\alpha_L}{\rho} + \gamma^S + \kappa(1 - \gamma^S) \frac{\alpha_L}{r+\rho+\alpha_L}}$$

and

$$\bar{\theta} \in \left(\frac{\frac{\alpha_L}{\rho} + 1}{\frac{\alpha_L}{\rho} + \gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r+\rho+\alpha_L}}, \frac{1}{\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r+\rho+\alpha_L}} \right).$$

¹²In Case (i), the supremum in (8) is approached as $z \rightarrow 0$. In Case (ii), it is attained at z^* . In Case (iii), it is approached as $z \rightarrow +\infty$.

at date zero. Prices fluctuate over time as shown in Figure 2. Prices are at their lowest at date zero, but jump immediately thereafter. They then fall gradually up to each sale, and subsequently jump. As we have noted, the lowest price is below the low value θ_L (equal to 1 in the example). This ensures low-value buyers are willing to purchase, given that they can wait for their values to become high.




Figure 2: Optimal price path for the following parameters: $\theta_H = 1.65$, $\theta_L = 1$, $\alpha_H = \ln(4/3)$, $\alpha_L = (1/4) \ln(4/3)$, $\gamma^S = 1/5$, $r = \ln(3/2)$, $\rho = \ln(8/7)$, $\lambda = 1$. If time is measured in years, then a high-value buyer's value changes with probability one quarter within a year, while a low-value buyer's value changes with probability 0.07. The probability that a buyer leaves the market due to a shock (at rate ρ) in the course of a year is one eighth. One third of the good's value is destroyed due to discounting (at rate r) in the course of a year.

While the form of the optimal sales policy depends on the ratio $\frac{\theta_H}{\theta_L}$ as described in Proposition 1, the time between sales z^* in Case (ii) also depends on this ratio in a natural way. In particular, as $\frac{\theta_H}{\theta_L}$ increases, the seller finds frequent sales more costly in terms of buyer rents, and less valuable for improving efficiency, so that the following holds.

Corollary 1 *The time between sales z^* in Case (ii) of Proposition 1 is increasing in $\frac{\theta_H}{\theta_L}$ over the interval $(\underline{\theta}, \bar{\theta})$.*

Intuition for price fluctuations. Recalling the argument in the Introduction, consider why the seller may benefit from spacing sales apart, rather than selling to low-value buyers at all dates. By spacing sales apart, the seller can reduce the rents to buyers who arrive with high values between sales dates. The seller finds this particularly effective given that a high value may turn low by the time of the next sale. There is also another reason the seller benefits from this policy: It reduces the rents left to buyers who arrive with low values, owing to the possibility that their values become high in the future. Given the optimal prices described in Lemma 2, this in turn reduces the rents of all buyers who arrive at earlier dates, back to date zero.

This intuition can be reinforced by studying what happens as the rate of downward switching grows large (i.e., as $\alpha_H \rightarrow +\infty$). The next result considers this case while assuming that it is more efficient to sell to low-value buyers than to wait for their values to turn high. This means assuming $1 < \frac{\theta_H}{\theta_L} < \frac{r+\rho+\alpha_L}{\alpha_L}$.¹³

Corollary 2 *Fix $\lambda, \alpha_L, r, \rho, \theta_H, \theta_L > 0$ such that $1 < \frac{\theta_H}{\theta_L} < \frac{r+\rho+\alpha_L}{\alpha_L}$. Now, consider varying α_H , while letting the probability a buyer's value is initially high, γ , equal $\frac{\alpha_L}{\alpha_L+\alpha_H}$ (thus maintaining stationarity). For any $\varepsilon > 0$, there exists $\bar{\alpha}_H$ such that, if $\alpha_H > \bar{\alpha}_H$, then Case (ii) of Proposition 1 applies and $z^* < \varepsilon$.*

Corollary 2 simply states that, as α_H grows large, the seller optimally holds sales at regular intervals, and sales are very frequent. The reason for this result is the following. If sales are held periodically, and if the time between sales, z , is fixed to be small, then the efficiency loss due to delayed purchases is small. As $\alpha_H \rightarrow +\infty$, buyers' expected rents approach zero (i.e., there is full surplus extraction), and the price on sales dates approaches

¹³If instead $\frac{\theta_H}{\theta_L} > \frac{r+\rho+\alpha_L}{\alpha_L}$, then selling to a low-value buyer is inefficient. To see why this might occur in practice, suppose a buyer's date- t value θ_t represents the expected future flow payoffs from holding the good. Suppose the buyer's state alternates between "low", where he does not need the good, and "high", where he does need it (his value is θ_L in the low state and θ_H otherwise). Then, it is inefficient for the buyer to acquire the good in the low state if he incurs a storage cost and hence receives a negative flow payoff in this state.

the low value θ_L . To see why, note that a buyer with a high value at any time other than a sales date (or moments before a sales date) believes that his value will, with high probability, drop to low by the next sale. The additional rents a buyer expects due to a (temporarily) high value therefore vanish as $\alpha_H \rightarrow +\infty$ (recall Equation (2)). Buyers with low values then realize that, if their values become high, this is unlikely to occur at (or just before) a sale. The rent expected by low-value buyers hence vanishes as well (recall Equation (4)). In contrast, if $A = \mathbb{R}_+$, then rents do not vanish as $\alpha_H \rightarrow +\infty$. In fact, the optimal price path conditional on $A = \mathbb{R}_+$ (see Lemma 2) is a constant price, and this price is independent of α_H . In particular, it remains fixed below θ_L as $\alpha_H \rightarrow +\infty$. Hence, $A = \mathbb{R}_+$ cannot be optimal.

Corollary 2 stands in contrast to what happens if instead α_L grows large. Then, holding frequent sales cedes large rents to buyers (since low-value buyers believe it is likely their values will be high by the next sale). Moreover, since a low-value buyer's value will increase soon, precipitating a purchase, sales are less important for efficiency (in fact, they reduce surplus if α_L is large enough; see Footnote 13). Whenever α_L is sufficiently large, the seller therefore holds no sales and the optimal price path is constant at θ_H . Note that this conclusion also holds if *both* α_L and α_H grow large, while maintaining stationarity (i.e., $\gamma = \frac{\alpha_L}{\alpha_L + \alpha_H}$) and while holding the probability of a high value, $\frac{\alpha_L}{\alpha_L + \alpha_H}$, constant.

Constant values benchmark. It is now interesting to examine the role of the various elements in our model, which can be achieved by removing them, one at a time. First, suppose that values do not change (i.e., suppose $\alpha_L = \alpha_H = 0$), as in Conlisk, Gerstner and Sobel (1984). Set the probability of a high value equal to any $\gamma \in (0, 1)$. Then the optimal price is constant and either equal to θ_L or θ_H . A simple proof is as follows. If the seller can (counterfactually) condition her offer on each buyer's arrival date (consider any offer, not necessarily a price path), then the seller does no better than to offer the static monopoly price at that date, inducing the buyer to make any purchase immediately. But this policy does not depend on the arrival date, and the seller can therefore achieve the same outcome for all arrival dates by committing to a constant price path, with price equal to the static monopoly

price. The buyer then finds it optimal to make any purchase immediately, because his value does not change. This shows that a constant price is optimal among all possible mechanisms, not just price paths. While this logic is well known, it is worth pointing out that it applies also in our setting where buyers leave the market stochastically (at rate $\rho > 0$). It is therefore clear that price fluctuations are not driven in our model by random exits from the market.

Shutting down dynamic arrivals. We can instead shut down dynamic arrivals by assuming the initial mass of consumers I is positive, but that the arrival rate λ is zero. In this case, the seller optimally chooses at most one sales date. The profit from choosing a sales date a is $IR(a)$, with $R(\cdot)$ given by (5), since all buyers must wait exactly length a for a sale. It is easy to see that either $IR(a)$ obtains its maximum at $a = 0$,¹⁴ or $IR(a)$ approaches its supremum as $a \rightarrow +\infty$, so that $A^* = \emptyset$. Thus, under the optimal policy, either all buyers purchase at date zero, or low-value buyers never purchase. In particular, low-value buyers never *wait* to purchase, unlike what we find for our baseline model.¹⁵

This benchmark, where all buyers arrive on the same date, sheds light on the optimality of including a sale at date zero in our baseline model. If $IR(a)$ obtains its maximum at $a = 0$, then the optimal sales policy in Case (iii) of Proposition 1 is $A^* = \{0\}$ (otherwise, it is $A^* = \emptyset$). By selling to low-value buyers at zero, the seller maximizes profits obtained from the mass I of buyers arriving at zero, but it yields no rents to buyers who arrive later. The reason the seller finds it optimal to include a sale at date zero, but no subsequent sales, is simply that subsequent sales increase the rents of all earlier arrivals (whereas, at date zero, there are no earlier arrivals in our model). As explained above, the optimal price path for our baseline model therefore involves a low price (θ_L) at date zero and a constant high price (θ_H) thereafter. The form of

¹⁴This is true if and only if $\theta_L \geq \theta_H (\gamma^S + (1 - \gamma^S) \int_0^\infty \alpha_L e^{-t(\alpha_L + r + \rho)} dt)$.

¹⁵This finding is analogous to Biehl's (2001) result for a two-period model. Biehl finds that either all buyers purchase in the first period, or low-value buyers never purchase. In both Biehl's setting, and ours, stationarity of values plays an important role. If we instead assume $I > 0$, $\lambda = 0$, $\gamma > 0$, $\alpha_H > 0$ and $\alpha_L = 0$ (so that the evolution of values is nonstationary), then we may have $A^* = \{a\}$ for $a > 0$. In this case, a low-value buyer waits and purchases at date a .

the price path is thus as in Deb (2014), who studies a model where all buyers arrive at date zero, have a continuum of values, and where these values change at most once. Deb suggests this as a theory of “introductory pricing” (and points to its use by Amazon, iTunes and Google Play, among others). Finally, note that our observations here apply also in Case (ii) of Proposition 1. For parameters such that Case (ii) applies, $IR(a)$ necessarily obtains its maximum at $a = 0$, and we show it is indeed optimal to include a sale at date zero in the baseline model.

Our analysis also sheds light on what happens if there is a rush of buyers into the market on a particular date. For instance, consider our baseline model and suppose that Case (iii) of Proposition 1 applies, while $IR(a)$ obtains its maximum at $a = 0$. Then a large enough mass of buyers arriving at a given (foreseen) date implies the optimality of a sale at that date, with the price dropping to θ_L . This is simply because the seller optimally maximizes profits for the large mass of arrivals, and is relatively less concerned with lowering the rents of earlier arrivals.¹⁶ The observation arguably squares well with Warner and Barsky’s (1995) evidence that the prices of durables are low at times when shopping intensity is high (such as holidays). In particular, it may be reasonable to think that such dates are times when the number of arrivals to the market is large.¹⁷

Shutting down exogenous exits. Next, consider the role of stochastic departures from the market (at rate ρ). These are important for the stationarity of the seller’s problem (this stationarity, for instance, is what ensures price cycles have a constant length in Case (ii) of Proposition 1). If we instead set the exit rate ρ to zero, but suppose that the initial mass of buyers I is finite, then the total population of buyers (that is, all buyers, not only those who have not purchased) grows without bound and the seller’s problem is no

¹⁶Note that changing values are crucial for the finding of a price reduction in response to a large mass of arrivals: if values were instead constant over time, then optimal prices would be constant and independent of the timing of buyer arrivals.

¹⁷Warner and Barsky make a related, but different, argument. They suggest that price reductions often occur at times when consumers can economize on search costs by searching across a range of products. In particular, they suggest weekends and holidays as times when the “intensity of shopping activity” is exogenously high due to low search costs.

longer stationary. The original working paper version (Garrett, 2011) considers this case (setting $I = 0$) and shows that the optimal sales policy is always bounded. Hence, the optimal price eventually remains constant and equal to θ_H . The reason is that sales held at later dates becomes increasingly costly in terms of buyer rents, as compared with the efficiency gains from selling to low-value buyers. Intuitively, this is because all earlier buyers have the option to wait for the sale, and are at no risk of exiting the market due to an exogenous shock. However, for the same reasons discussed above, price fluctuations can still arise at early enough dates (i.e., before the price transitions to θ_H). This provides a sense in which stochastic exits are not crucial to our finding of cycling prices.

Conversely, one can consider the polar opposite case in which all buyers stay in the market only for an instant (this case has also been of interest in the literature; see, e.g., Gershkov and Moldovanu, 2009). That values change is then irrelevant and the optimal price remains constant over time.

Other processes and more than two values. Finally, it is important to consider the extent to which our finding of, and intuition for, fluctuating prices can be expected to extend to other stochastic processes for values. First, one can observe that stationarity is not essential. Many of the arguments in this paper can be easily extended to allow the probability a buyer's value is high at the arrival date, γ , to differ from the stationary probability $\gamma^S = \frac{\alpha_L}{\alpha_L + \alpha_H}$ (for instance, Lemmas 1 and 2 did not use $\gamma = \gamma^S$). Stationarity does simplify somewhat the formulation of the dynamic program (in (8)) and it simplifies the statement of our main result (Proposition 1).

For more general processes, for instance with more than two values, we can point out where our intuition for price fluctuations seems robust. Provided those buyers with higher values than others in the market do not expect these to persist for too long, the seller stands to gain by making them wait for low prices. Indeed, this is a way to reduce their rents, and hence also the rents of buyers who arrive earlier with relatively low values, but who expect that they may later have high values.

To be more precise, it is useful to recall the findings of the literature on

dynamic mechanism design with stochastic types. Unlike the present paper (which restricts the seller to offering a price path), this literature does not restrict the space of mechanisms or buyer communication. However, some of its insights remain relevant. For instance, Battaglini (2005) finds, in a two-value model, that low-value buyers initially receive inefficiently low allocations, close to the time of contracting. After enough time, however, buyers whose values remain low receive close-to-efficient allocations; Battaglini terms this the principle of “vanishing distortions at the bottom”. Case (ii) of Proposition 1 in the present paper reflects a similar principle. When this case applies, low-value buyers do not take the efficient action of purchasing the good immediately upon arrival, but they instead wait for a sale. Distortions vanish on the sales date, when all the low-value buyers (who are still in the market and have not yet purchased) buy the good. Garrett (2011) makes the connection between the optimal dynamic mechanism (of any form) and the optimal price path more formally. Pavan, Segal and Toikka (2014) consider the optimal mechanism with a continuum of agent types and show that vanishing distortions also arise provided that the impulse responses of later types to initial types vanish with time (for instance, consider a first-order autoregressive process with persistence parameter less than one). It thus seems reasonable to conjecture that, for the same reasons as in the present paper, the optimal price path should feature cycling prices when buyers have a continuum of values and arrive dynamically to the market, provided impulse responses vanish over time.

One reason the above observations are interesting is that they suggest cases where our ideas do not apply. Suppose impulse responses do not vanish, but are instead equal to one, as with a random walk. Then the key intuition of this paper should not be expected to apply. In this case, values are highly persistent, and a buyer with a high value therefore does not expect his value to systematically decline.

III Conclusion

We studied the profit-maximizing price path when buyers arrive over time and have values for the good which change stochastically. For a range of parameter values, optimal prices fluctuate over time. Prices gradually fall up to sales dates and jump thereafter.

Our results have both normative and positive implications. On the normative side, we suggest a new reason why firms may be justified in changing their prices over time: buyers' uncertainty about their own future values. Indeed, price discounting can be optimal even when the environment is completely stationary, as described above. Understanding optimal pricing in such environments seems important at a time when firms increasingly have the tools to inform themselves about patterns in customer preferences. This paper suggests the importance of understanding not only patterns in aggregate demand, but also how individual consumers expect their values for products to evolve. On the positive side, we provided a new explanation for temporary price discounts.

Appendix

This Appendix provides the proofs of all results.

Proof of Lemma 1. Follows from the arguments in the text. ■

Proof of Lemma 2. Let

$$\begin{aligned} w_t^A(\theta_L) &\equiv \int_t^\infty \alpha_L e^{-(r+\rho)(s-t) - (r+\rho+\alpha_L+\alpha_H)(m_A(s)-s)} (\theta_H - \theta_L) ds, \text{ and} \\ w_t^A(\theta_H) &\equiv w_t^A(\theta_L) + e^{-(r+\rho+\alpha_L+\alpha_H)(m_A(t)-t)} (\theta_H - \theta_L) \end{aligned}$$

be the lower bounds on buyer expected rents under the sales policy A , as derived in the text. As explained in the text, if buyers find the mechanism $\mathcal{M}_{P,A}^* = \langle p_A^*, x_A^* \rangle$ incentive compatible, then buyers expect rents equal to these

lower bounds, and so this mechanism must maximize expected profit subject to $x_A^*(\theta_L, t) = 1$ iff $t \in A$. Therefore, it is enough to verify x_A^* is indeed an optimal stopping rule for buyers given p_A^* , so that the value of a buyer's problem is indeed given by w_t^A .

For any stopping rule $\sigma \in \Sigma$, any t , and any $\theta_t \in \{\theta_L, \theta_H\}$, letting $\mathbf{1}_{\tilde{\theta}_s = \theta_s}$ be the indicator function for $\theta_s \in \{\theta_L, \theta_H\}$,

$$\begin{aligned}
& \mathbb{E}_t \left[e^{-(r+\rho)(\tilde{\mu}_\sigma - t)} \left(\tilde{\theta}_{\tilde{\mu}_\sigma} - p_A^*(\tilde{\mu}_\sigma) \right) \mid \tilde{\theta}_t = \theta_t \right] \\
& \leq \mathbb{E}_t \left[e^{-(r+\rho)(\tilde{\mu}_\sigma - t)} w_{\tilde{\mu}_\sigma}^A \left(\tilde{\theta}_{\tilde{\mu}_\sigma} \right) \mid \tilde{\theta}_t = \theta_t \right] \\
& \leq w_t^A(\theta_t) \\
& \quad + \mathbb{E} \left[\int_t^{\tilde{\mu}_\sigma} e^{-(r+\rho)(s-t)} \left(\mathbf{1}_{\tilde{\theta}_s = \theta_L} \begin{pmatrix} -(r+\rho)w_s^A(\theta_L) \\ + \frac{dw_s^A(\theta_L)}{ds} \\ + \alpha_L \begin{pmatrix} w_s^A(\theta_H) \\ -w_s^A(\theta_L) \end{pmatrix} \end{pmatrix} \right. \right. \\
& \quad \left. \left. + \mathbf{1}_{\tilde{\theta}_s = \theta_H} \begin{pmatrix} -(r+\rho)w_s^A(\theta_H) \\ + \frac{dw_s^A(\theta_H)}{ds} \\ + \alpha_H \begin{pmatrix} w_s^A(\theta_L) \\ -w_s^A(\theta_H) \end{pmatrix} \end{pmatrix} \right) ds \mid \tilde{\theta}_t = \theta_t \right] \\
& \leq w_t^A(\theta_t).
\end{aligned}$$

The first inequality follows by choice of $p_A^*(\cdot)$. The second inequality follows by applying Dynkin's formula, after observing that (i) $\frac{dw_s^A(\theta_H)}{ds}$ exists except on the countably many boundary points of A , ∂A , and (ii)

$$w_t^A(\theta_H) = w_0^A(\theta_H) + \int_0^t \frac{dw_s^A(\theta_H)}{ds} ds - \sum_{\{\tau \in \partial A: \tau < t\}} g(\tau),$$

where

$$g(\tau) = \left(1 - \lim_{s \searrow \tau} e^{-(r+\rho+\alpha_L+\alpha_H)(m_A(s)-s)} \right) (\theta_H - \theta_L)$$

takes only non-negative values on ∂A . The third inequality follows because

$$-(r + \rho) w_t^A(\theta_L) + \frac{dw_t^A(\theta_L)}{dt} + \alpha_L (w_t^A(\theta_H) - w_t^A(\theta_L)) = 0$$

and

$$-(r + \rho) w_t^A(\theta_H) + \frac{dw_t^A(\theta_H)}{dt} + \alpha_H (w_t^A(\theta_L) - w_t^A(\theta_H)) \leq 0$$

except for at most countably many points. When $\sigma = x_A^*$, all inequalities hold with equality. Thus w_t^A is the value function associated with the buyer's problem and x_A^* is an optimal strategy for the buyer. ■

Proof of Proposition 1. Let A be any left-closed subset of \mathbb{R}_+ .¹⁸ After integration by parts, the rent earned by all buyers who arrive after date zero is equal to

$$(9) \quad (\theta_H - \theta_L) \int_0^\infty \lambda e^{-r\tau} \left(\frac{\alpha_L}{\rho} (1 - e^{-\rho\tau}) + \gamma^S \right) e^{-(r+\rho+\alpha_L+\alpha_H)(m_A(\tau)-\tau)} d\tau,$$

where $m_A(\tau)$ gives the date of purchase for a buyer who arrives at date τ if his value remains low. The expected rent of a single buyer who arrives at date zero is equal to

$$(10) \quad (\theta_H - \theta_L) \int_0^\infty \alpha_L e^{-(r+\rho)s - (r+\rho+\alpha_L+\alpha_H)(m_A(s)-s)} ds \\ + \gamma^S e^{-(r+\rho+\alpha_L+\alpha_H)m_A(0)} (\theta_H - \theta_L).$$

Multiplying (10) by the mass of date-0 arrivals $I = \frac{\lambda}{\rho}$ and adding to (9), we find that the total rent across *all* buyers is

$$(\theta_H - \theta_L) \int_0^\infty \lambda e^{-r\tau} \left(\frac{\alpha_L}{\rho} + \gamma^S \right) e^{-(r+\rho+\alpha_L+\alpha_H)(m_A(\tau)-\tau)} d\tau \\ + \frac{\lambda}{\rho} \gamma^S e^{-(r+\rho+\alpha_L+\alpha_H)m_A(0)} (\theta_H - \theta_L).$$

¹⁸Thus A is any set of dates such that, for some $\sigma \in \Sigma$, $\sigma(\theta_L, t) = 1$ iff $t \in A$.

The seller's total expected profit is therefore

$$(11) \quad \int_0^\infty \lambda e^{-r\tau} b(m_A(\tau) - \tau) d\tau + \frac{\lambda}{\rho} \psi(m_A(0))$$

where, for all $y \geq 0$,

$$(12) \quad \begin{aligned} b(y) &= \theta_H \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r + \rho + \alpha_L} \right) \\ &\quad + e^{-(r+\rho+\alpha_L)y} (1 - \gamma^S) \left(\theta_L - \frac{\alpha_L}{r + \rho + \alpha_L} \theta_H \right) \\ &\quad - e^{-(r+\rho+\alpha_L+\alpha_H)y} \left(\frac{\alpha_L}{\rho} + \gamma^S \right) (\theta_H - \theta_L), \end{aligned}$$

and

$$\begin{aligned} \psi(y) &= \theta_H \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r + \rho + \alpha_L} \right) \\ &\quad + e^{-(r+\rho+\alpha_L)y} (1 - \gamma^S) \left(\theta_L - \frac{\alpha_L}{r + \rho + \alpha_L} \theta_H \right) \\ &\quad - \gamma^S e^{-(r+\rho+\alpha_L+\alpha_H)y} (\theta_H - \theta_L). \end{aligned}$$

The seller's problem is to maximize (11) by choice of A given that $m_A(\tau)$ is defined by (1) for all τ .

Case (i). It is easiest to dispose with Case (i) right away, hence suppose $\frac{\theta_H}{\theta_L} \leq \underline{\theta}$ (with $\underline{\theta}$ given in Footnote 11). In this case, it is easy to verify that b and ψ are maximized over \mathbb{R}_+ at zero. Hence, $m_A(\tau) = \tau$ for all τ is optimal, i.e. $A^* = \mathbb{R}_+$.

Remaining Cases. If instead $\frac{\theta_H}{\theta_L} > \underline{\theta}$, then Case (i) cannot apply. To see this, compute the derivative of b ,

$$\begin{aligned} &b'(y) \\ &= -e^{-(r+\rho+\alpha_L)y} \left(\begin{array}{c} (r + \rho + \alpha_L) (1 - \gamma^S) \left(\theta_L - \frac{\alpha_L}{r+\rho+\alpha_L} \theta_H \right) \\ -e^{-\alpha_H y} \left(\frac{\alpha_L}{\rho} + \gamma^S \right) (r + \rho + \alpha_L + \alpha_H) (\theta_H - \theta_L) \end{array} \right), \end{aligned}$$

and note that b is strictly quasi-concave. Either b is increasing on \mathbb{R}_+ , or there exists $y^\#$ such that b is increasing up to $y^\#$ and decreasing thereafter. In particular, $b'(y^\#) = 0$ or

$$(13) \quad y^\# \equiv \frac{1}{\alpha_H} \log \left(\frac{\left(\frac{\alpha_L}{\rho} + \gamma^S \right) (r + \rho + \alpha_L + \alpha_H) (\theta_H - \theta_L)}{(1 - \gamma^S) (r + \rho + \alpha_L) \left(\theta_L - \frac{\alpha_L}{r + \rho + \alpha_L} \theta_H \right)} \right).$$

It is thus easy to see that the policy $A = \mathbb{R}_+$ is not optimal. For instance, profits are increased by taking $A = \{iy^\# : i \in \mathbb{N} \cup \{0\}\}$: thus $m(0) = 0$, and $m(\tau) - \tau \leq y^\#$ for all $\tau > 0$.

Next note that our problem is separable in the sense described in the main text. If sales dates are determined up to date a , itself a sales date, then

$$W^*(a) = \sup_{A \setminus [0, a]} \int_a^\infty \lambda e^{-r(\tau-a)} b(m_A(\tau)) d\tau$$

is the value of the principal's continuation problem of choosing sales dates after a (as defined in (7) in the main text). Past sales choices do not enter this expression; nor does calendar time. Let

$$(14) \quad V(z) = \frac{\int_0^z \lambda e^{-r\tau} b(z - \tau) d\tau}{1 - e^{-rz}}.$$

It is easy to show that

$$W^*(a) = W^* \equiv \sup_{z > 0} V(z)$$

(where W^* is defined by (8)). Assuming that $\frac{\theta_H}{\theta_L} > \underline{\theta}$, continuity of $V(\cdot)$ and our previous observation implies that either $V(\cdot)$ obtains its maximum at some $z^* \in [y^\#, +\infty)$, or $W^* = \lim_{z \rightarrow +\infty} V(z)$. The former possibility corresponds to Case (ii) in the proposition, while the latter corresponds to Case (iii).

Case (ii). Consider now Case (ii), i.e., suppose $V(\cdot)$ is maximized at

$z^* \in [y^\#, +\infty)$. We now show that z^* is unique. Note that

$$V'(z) = \frac{\lambda b(z) - rV(z)}{1 - e^{-rz}}.$$

Hence the first-order condition for maximization of $V(\cdot)$ yields $b(z^*) = \frac{r}{\lambda}W^*$. Since $z^* \geq y^\#$, and since $b'(y) < 0$ for $y > y^\#$, there can exist only one value of z^* .

Next, we show that, when Case (ii) applies, the first sale optimally occurs at date zero. Since Case (ii) applies, there must be some initial sales date, say at date $k \geq 0$. The expected profit when the first sale is at k can be written

$$\begin{aligned} & \theta_H \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{\alpha_L + r + \rho} \right) \left(\frac{\lambda}{\rho} + \frac{\lambda}{r} \right) \\ & + (1 - \gamma^S) e^{-rk} \left(\frac{\lambda}{\rho} e^{-(\alpha_L + \rho)k} + \int_0^k \lambda e^{-(\alpha_L + \rho)(k - \tau)} d\tau \right) \left(\theta_L - \theta_H \frac{\alpha_L}{\alpha_L + r + \rho} \right) \\ & - e^{-rk} \gamma^S \frac{\lambda}{\rho} (\theta_H - \theta_L) \\ & + e^{-rk} \frac{\int_0^{z^*} \lambda e^{-r\tau} \left(b(z^* - \tau) - \theta_H \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r + \rho + \alpha_L} \right) \right) d\tau}{1 - e^{-rz^*}}. \end{aligned} \tag{15}$$

The first term is the surplus that would be generated if all buyers purchased only when their values become high. The second term represents the additional surplus generated because buyers arriving at or before k , whose values remain low until k , purchase at date k rather than waiting for their values to become high (such purchases generate an additional expected surplus equal to $\theta_L - \theta_H \frac{\alpha_L}{\alpha_L + r + \rho}$). The third term is the rent that would be earned by buyers arriving at or before date k if there were no sales after date k . In particular, this is the rent these buyers would earn if waiting until date k to purchase (they are willing to do so), and if there were no sales after date k . In this case, there would be $\gamma^S \frac{\lambda}{\rho}$ purchases by high-value buyers, realizing them a rent $(\theta_H - \theta_L)$. The final term is the discounted additional value of the continuation problem over and above what would be obtained by holding no sales after k (the value

of holding no sales after k is already accounted for by the first term). Note that the second term is positive and that $\frac{\lambda}{\rho}e^{-(\alpha_L+\rho)k} + \int_0^k \lambda e^{-(\alpha_L+\rho)(k-\tau)} d\tau$ obtains its maximum at $k = 0$. It is then easy to see that (15) is maximized by $k = 0$, which implies the result.

Case (iii). Case (iii) can be divided into two sub-cases. Either there is an initial sale at some date k , or there are no sales ($A^* = \emptyset$). In the first case, the (date-0 value of) additional surplus due to the sale is no greater than

$$e^{-rk} \frac{\lambda}{\rho} (1 - \gamma^S) \left(\theta_L - \theta_H \frac{\alpha_L}{\alpha_L + r + \rho} \right)$$

(since the mass of buyers with low values at date k , who have not left due to a shock and have not purchased the good, is weakly less than $\frac{\lambda}{\rho} (1 - \gamma^S)$; and since if there were no sales, these buyers would eventually purchase when their values become high), while the additional rent due to the sale is $e^{-rk} \frac{\lambda}{\rho} \gamma^S (\theta_H - \theta_L)$ (since buyers arriving before k earn a rent equal to that obtained by waiting and purchasing at k , and since the mass of purchases by high-value buyers is then $\frac{\lambda}{\rho} \gamma^S$). If $k = 0$, the additional profit from the sale is exactly equal to

$$\frac{\lambda}{\rho} \left((1 - \gamma^S) \left(\theta_L - \theta_H \frac{\alpha_L}{\alpha_L + r + \rho} \right) - \gamma^S (\theta_H - \theta_L) \right),$$

which is thus larger than the additional profit in case $k > 0$. This establishes that either there is a sale, and it occurs at date zero, or there is no sale.

Case (ii) versus Case (iii). It remains to consider the boundary between Cases (ii) and (iii). If Case (iii) applies then

$$W^* = \frac{\lambda}{r} \lim_{z \rightarrow +\infty} b(z) = \frac{\lambda}{r} \theta_H \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r + \rho + \alpha_L} \right).$$

The two cases are therefore mutually exclusive, since otherwise we must have $b(z^*) = \lim_{z \rightarrow +\infty} b(z)$ for $z^* \in [y^\#, +\infty)$, which is impossible because b is strictly decreasing on $[y^\#, +\infty)$.

Now, note that

$$(16) \quad \frac{\theta_H}{\theta_L} \leq \frac{\frac{\alpha_L}{\rho} + 1}{\frac{\alpha_L}{\rho} + \gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r + \rho + \alpha_L}}$$

implies

$$b(y) \geq \theta_H \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r + \rho + \alpha_L} \right),$$

with strict inequality whenever $y > 0$. Hence, if (16) holds, Case (iii) cannot apply.

It remains to show that there is a threshold $\bar{\theta}$ such that Case (iii) applies for $\frac{\theta_H}{\theta_L} \geq \bar{\theta}$, and to show that

$$\bar{\theta} \in \left(\frac{\frac{\alpha_L}{\rho} + 1}{\frac{\alpha_L}{\rho} + \gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r + \rho + \alpha_L}}, \frac{1}{\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r + \rho + \alpha_L}} \right)$$

(as specified in Footnote 11). That Case (iii) applies for

$$\frac{\theta_H}{\theta_L} \geq \frac{1}{\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{\alpha_L + r + \rho}}$$

is immediate, because then $R(s)$ (in (5)) approaches its supremum

$$\theta_H \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{\alpha_L + r + \rho} \right)$$

as $s \rightarrow +\infty$, and attains it if there are no sales. Moreover, complete absence of sales minimizes buyer rents. That Case (iii) applies also for somewhat smaller values follows from examining Equation (8).

Finally, we want to show that which of Case (ii) and Case (iii) applies is determined simply relative to a threshold value for $\frac{\theta_H}{\theta_L}$. To see this, note first that the optimal sales policy depends on the ratio $\frac{\theta_H}{\theta_L}$, but not the values θ_H and θ_L individually. Then note that, holding θ_L fixed, we have that, for all

$y \in [0, +\infty)$,

$$\begin{aligned} & b(y) - \theta_H \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{r + \rho + \alpha_L} \right) \\ = & e^{-(r+\rho+\alpha_L)y} \begin{pmatrix} (1 - \gamma^S) \left(\theta_L - \frac{\alpha_L}{r+\rho+\alpha_L} \theta_H \right) \\ - \left(\frac{\alpha_L}{\rho} + \gamma^S \right) e^{-\alpha_H y} (\theta_H - \theta_L) \end{pmatrix} \end{aligned}$$

is decreasing in θ_H . Hence, there must exist a threshold value for $\frac{\theta_H}{\theta_L}$, dependent on the other parameters of the model, above which Case (iii) applies. A simple continuity argument implies that the set of parameters for which Case (iii) applies must be closed. ■

Proof of Corollary 1. Since the optimal sales policy depends only on the ratio $\frac{\theta_H}{\theta_L}$, make the normalization $\theta_L = 1$. Then, let the functions defined in (12) and (14) depend explicitly on θ_H ; i.e., write $b(\cdot; \theta_H)$ and $V(\cdot; \theta_H)$. From the proof of Proposition 1, we know that the optimal time between sales $z^*(\theta_H)$, if Case (ii) applies, is given by

$$b(z^*(\theta_H); \theta_H) = \frac{r}{\lambda} V(z^*(\theta_H); \theta_H).$$

Fix $\theta'_H > 1$ such that Case (ii) of Proposition 1 applies. Then $b(z; \theta'_H) - \frac{r}{\lambda} V(z; \theta'_H)$ is strictly decreasing in z in a neighborhood of $z^*(\theta'_H)$, since, by optimality of $z^*(\theta'_H)$, $\frac{\partial V(z^*(\theta'_H); \theta'_H)}{\partial z} = 0$, while we argued in the proof of Proposition 1 that $\frac{\partial b(z^*(\theta'_H); \theta'_H)}{\partial z} < 0$. Thus, by the implicit function theorem, there is a unique local solution $z^*(\theta_H)$ for θ_H in a sufficiently small neighborhood of θ'_H . Continuity of W^* in θ_H then implies that the local solution must indeed give the optimal time between sales (since the optimal time between sales must satisfy $b(z^*(\theta_H); \theta_H) = \frac{r}{\lambda} W^*$ for all θ_H with $z^*(\theta_H) \geq y^\#$, where $y^\#$ is given

by (13), as argued in the proof of Proposition 1). Finally, observe that

$$\begin{aligned}
& b(z; \theta_H) - \frac{r}{\lambda} V(z; \theta_H) \\
&= e^{-(r+\rho+\alpha_L)z} (1 - \gamma^S) \left(1 - \frac{\alpha_L}{r + \rho + \alpha_L} \theta_H \right) \\
&\quad \times \left(\frac{r}{1 - e^{-rz}} \int_0^z e^{(\rho+\alpha_L+\alpha_H)\tau} d\tau - 1 \right) \\
&\quad \times \left(\frac{e^{-\alpha_H z} \left(\frac{\alpha_L}{\rho} + \gamma^S \right) (\theta_H - 1)}{1 - \frac{\alpha_L}{r+\rho+\alpha_L} \theta_H} - \frac{\frac{r}{1-e^{-rz}} \int_0^z e^{(\rho+\alpha_L)\tau} d\tau - 1}{\frac{r}{1-e^{-rz}} \int_0^z e^{(\rho+\alpha_L+\alpha_H)\tau} d\tau - 1} \right).
\end{aligned}$$

Thus, we find that $b(z^*(\theta'_H); \theta_H) - \frac{r}{\lambda} V(z^*(\theta'_H); \theta_H)$ is strictly positive for $\theta_H \in \left(\theta'_H, \frac{r+\rho+\alpha_L}{\alpha_L} \right)$, and strictly negative for $\theta_H \in (1, \theta'_H)$, and hence $z^*(\cdot)$ must be increasing over a sufficiently small neighborhood of θ'_H . ■

Proof of Corollary 2. Note first that $\underline{\theta}$ in Footnote 11 converges to 1 as $\alpha_H \rightarrow +\infty$; hence Case (i) of Proposition 1 does not apply for all α_H sufficiently large. Second, inspecting (12), we see that $W^* \rightarrow \frac{\lambda}{r} \theta_L$ as $\alpha_H \rightarrow +\infty$ and this is attained only by a policy with frequent sales, time z^* apart, where $z^* \rightarrow 0$ as $\alpha_H \rightarrow +\infty$. ■

References

- [1] Battaglini, Marco (2005), ‘Long-Term Contracting with Markovian Consumers,’ *American Economic Review*, 95, 637-658.
- [2] Biehl, Andrew R. (2001), ‘Durable-Goods Monopoly with Stochastic Values,’ *Rand Journal of Economics*, 32, 565-577.
- [3] Board, Simon (2008), ‘Durable-Goods Monopoly with Varying Demand,’ *Review of Economic Studies*, 75, 391-413.
- [4] Board, Simon and Andrzej Skrzypacz (forthcoming), ‘Revenue Management with Forward-Looking Buyers,’ *Journal of Political Economy*.

- [5] Conlisk, John (1984), ‘A peculiar example of temporal price discrimination,’ *Economics Letters*, 15, 121-126.
- [6] Conlisk, John, Eitan Gerstner and Joel Sobel (1984), ‘Cyclic Pricing by a Durable Goods Monopolist,’ *Quarterly Journal of Economics*, 99, 489-505.
- [7] Deb, Rahul (2014), ‘Intertemporal Price Discrimination with Stochastic Values,’ mimeo University of Toronto.
- [8] Garrett, Daniel (2011), ‘Durable Goods Sales with Dynamic Arrivals and Changing Values,’ working paper version, mimeo Northwestern University.
- [9] Gershkov, Alex and Benny Moldovanu (2009), ‘Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach,’ *American Economic Journal: Microeconomics*, 1, 168-198.
- [10] Gershkov, Alex, Benny Moldovanu and Philipp Strack (2014), ‘Revenue Maximizing Mechanisms with Strategic Customers and Unknown Demand: Name-Your-Own-Price,’ mimeo Hebrew University of Jerusalem, University of Bonn, and UC Berkeley.
- [11] Gowrisankaran, Gautam and Marc Rysman (2012), ‘Dynamics of Consumer Demand for New Durable Goods,’ *Journal of Political Economy*, 120, 1173-1219.
- [12] Horner, Johannes and Larry Samuelson (2011), ‘Managing Strategic Buyers,’ *Journal of Political Economy*, 119, 379-425.
- [13] Iyengar, Raghuram, Sangman Han and Sunil Gupta (2009), ‘Do Friends Influence Purchases in a Social Network,’ Harvard Business School Marketing Unit Working Paper No. 09-123.
- [14] Kashyap, Anil (1995), ‘Sticky Prices: New Evidence from Retail Catalogs,’ *Quarterly Journal of Economics*, 110, 245-274.

- [15] Kilts-Nielsen Data Center (2006–), ‘Retail Scanner Data,’ University of Chicago Booth School of Business. <https://research.chicagobooth.edu/nielsen/> (accessed July 1, 2015).
- [16] Klenow, Peter and Benjamin Malin (2010), ‘Microeconomic Evidence on Price-Setting,’ *Handbook of Monetary Economics*, 3, 231-284.
- [17] Landsberger, Michael and Isaac Meilijson (1985), ‘Intertemporal price discrimination and sales strategy under incomplete information’, *Rand Journal of Economics*, 15, 171-196.
- [18] Maskin, Eric and Jean Tirole (1988), ‘A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles’, *Econometrica*, 56, 571-599.
- [19] Nakamura, Emi and Jón Steinsson (2008), ‘Five Facts about Prices: A Reevaluation of Menu Cost Models,’ *Quarterly Journal of Economics*, 123, 1415-1464.
- [20] Pavan, Alessandro, Ilya Segal and Juuso Toikka (2014), ‘Dynamic Mechanism Design: A Myersonian Approach’, *Econometrica*, 82, 601-653.
- [21] Pesendorfer, Martin (2002), ‘Retail Sales: A Study of Pricing Behavior in Supermarkets’, *Journal of Business*, 75, 33-66.
- [22] Slade, Margaret (1998), ‘Optimal Pricing with Costly Adjustment: Evidence from Retail-Grocery Prices,’ *Review of Economic Studies*, 65, 87-107.
- [23] Sobel, Joel (1991), ‘Durable Goods Monopoly with Entry of New Consumers’, *Econometrica*, 1991, 1455-1485.
- [24] Stokey, Nancy L. (1979), ‘Intertemporal Price Discrimination’, *Quarterly Journal of Economics*, 93, 355-371.
- [25] Strausz, Roland (2006), ‘Deterministic versus Stochastic Mechanisms in Principal-Agent Models’, *Journal of Economic Theory*, 128, 306-314.

- [26] Varian, Hal R. (1980), 'A Model of Sales', *American Economic Review*, 70, 651-659.
- [27] Warner, Elizabeth J. and Robert B. Barsky (1995), 'The Timing and Magnitude of Retail Store Markdowns: Evidence from Weekends and Holidays', *Quarterly Journal of Economics*, 110, 321-352.