

Durable Goods Sales with Dynamic Arrivals and Changing Values^{*†}

Daniel Garrett

Department of Economics
Northwestern University
danielgarrett2013@u.northwestern.edu

November, 2011

JOB MARKET PAPER

Abstract

I study a profit-maximizing monopolist selling a durable good to buyers who arrive over time. Buyers are forward looking and their values for the good evolve stochastically. I first suppose that the seller commits to a path of posted prices in a stationary environment. Contrary to the case where values remain constant, optimal prices fluctuate over time. This is consistent with the use of sales or periodic discounts. If the arrival rate varies with time, the seller favors discounting when the rate is at its fastest, suggesting an explanation for the common empirical finding of countercyclical markups. A path of posted prices, however, is not the most profitable mechanism. Because values evolve over time, higher profit can be obtained with more sophisticated mechanisms. In particular, the unrestricted optimal mechanism can be implemented by selling options to purchase the good at future dates. As is the case for optimal posted prices, buyers who arrive later expect to earn less rent and expect to wait longer to obtain the good. Contracts agreed at later dates are therefore less efficient.

JEL classification: D82. *Keywords:* changing types, durable goods, dynamic mechanism design, price cycles, sales, countercyclical markups

^{*}This paper has benefited from detailed comments from my advisor, Alessandro Pavan, as well as the helpful suggestions of Jeff Ely, Igal Hendel, Alexander Koch, Konrad Mierendorff, Aviv Nevo, Bill Rogerson, Bruno Strulovici and Rakesh Vohra. Thank you also to seminar participants at Aarhus University, Northwestern University and at the Arne Ryde Memorial Lectures at the University of Lund. I am solely accountable for any errors.

[†]This paper previously circulated under the title ‘Durable Goods Sales with Changing Values and Unobservable Arrival Dates’.

1 Introduction

This paper examines the problem of a monopolist selling a durable good to buyers who arrive over time and whose values evolve stochastically. Buyers are forward looking and both values and arrival times are their private information. Sellers of a wide range of durable goods face buyers with these characteristics. One reason is simply that the private circumstances of consumers are always changing, leading them to consider new goods that they have not considered in the past, and implying that their values for these goods subsequently evolve with the passage of time. Somebody who develops an interest in photography may start considering high-end cameras; from this time on, the value of owning a camera depends on the extent of opportunities to take interesting photos. A family that newly acquires a beach house may consider purchasing a boat; from then on, their value for a boat depends on the amount of time that they have for vacation. Since consumers often find it difficult to predict how much they will use a good, or how much they will enjoy doing so, they face private uncertainty about their future values. Since values and arrival times reflect personal circumstances, they are typically private information.

I study profit-maximizing mechanisms in the environment described above. One question that continues to be debated in the economics literature as well as in the business press is “under what conditions should firms hold sales in order to maximize profits”?¹ In terms of what is observed empirically, there are firms such as Apple Inc. which have reputations for rarely discounting their products.² On the other hand, price fluctuations and sales are observed across a broad range of markets (see, e.g., Warner and Barsky (1995) for appliances and other consumer durables, and Bils and Klenow (2004) for a more extensive data set). Although there are many theories of sales in the literature (I attempt a review below), the optimality of a constant price is nonetheless a common prediction in environments where the seller can fully commit and where buyers’ values are constant. Supposing that firms can find ways to credibly commit to future actions (e.g., through reputation, as in the case of Apple), why they should choose to vary their prices over time, at least in stationary environments, remains something of a puzzle.³

Motivated by these questions, the first part of this paper studies the profit-maximizing path of posted prices when the seller can fully commit. The main reason for studying posted prices is that they are the predominant institution in retail markets. I do not try to explain why posted prices are so common.

¹See Kadet (2004) for an example of detailed discussion in the business press. For the economics literature, see the review provided below.

²See Sutter (2010, November 29).

³See McAfee and te Velde (2006) for a discussion in this regard.

Possible explanations seem to include the cost of repeated interactions between firms and customers required to implement more sophisticated mechanisms, as well as the costs of contractual complexity, especially the cost of keeping records. Instead, I examine the tradeoffs a seller faces in the environment described above when selling goods according to the prevailing institution.

Assuming buyers' values switch stochastically between two levels ("low" and "high"), the profit-maximizing price path often exhibits a distinctive pattern of cycles. The price is high at the beginning of a cycle, but it falls gradually over time before jumping at the start of a new cycle. After arriving on the market, the buyer purchases immediately whenever his value is high but waits until the end of the cycle before purchasing if his value remains low. Following others in the literature (e.g., Conlisk et al. (1984)), I interpret the end of each cycle as a "sale". Contrary to some existing explanations for sales (e.g., Board (2008), where prices fluctuate because different kinds of buyers arrive on the market at different times), the environment which I consider is completely stationary. Sales in my theory are therefore not driven by systematic changes in the environment, but are rather a response to using the posted price institution in a setting where buyers have private uncertainty about their future values.

Because buyers face private uncertainty, the seller can obtain greater profits by using more sophisticated mechanisms than posted prices. In order to understand the institution which maximizes seller profits in the environment described above, I characterize the unrestricted optimal mechanism. The unrestricted optimum can be implemented by selling options to buy the good at prices which depend on both the date of and on the time since the option's sale. The fee paid for the option contract also depends on the date of sale. Option contracts expire after finite time, after which buyers can no longer obtain the good. Buyers execute their options whenever their values are high, or, if their values remain low, at the expiration date.

This paper is the first to characterize profit-maximizing mechanisms in an environment where agents' values change stochastically and where their arrival is dynamic. Since buyers' values and arrival times are their private information, a key difficulty is incentivizing them to reveal themselves at the moment of arrival by purchasing an option.⁴ In particular, buyers with low values must be discouraged from waiting and participating only when their values become high. Whether buyers find it optimal to participate as soon as they arrive depends on what they can expect to obtain by participating upon arrival as well as at each date into the infinite future.

⁴Contracting before the arrival date is impossible simply because buyers do not read the seller's advertisements or pay any attention to her offers. That is, the search technology is such that buyers come into contact with the seller's offer only after developing an interest in the product.

Several papers study sellers facing buyers who arrive over time (for a discussion of the work on selling to dynamic populations see Bergemann and Said (2010) as well as the literature review below). However, the work on profit maximization does not consider the possibility that buyer values change stochastically. At the same time, the literature on profit-maximizing dynamic mechanisms for agents with stochastically changing types does not consider the possibility of dynamic arrival. Allowing for both elements simultaneously raises new challenges for the design of optimal mechanisms.

In contrast to the above literature, buyers expect at the time they participate to earn positive rent even if their values are low. Often, buyers who arrive sufficiently early receive a subsidy to take the option (rather than paying a positive fee for it). In return for the subsidy, buyers give up the right to participate in the mechanism later (i.e., to purchase future option contracts or to receive any future subsidies). That buyers are subsidized can mean that they pay less for the good than its expected value, even when their values are low at the time of acquisition. These effects arise because the seller must dissuade buyers with low initial values from waiting to participate only when their values become high.

For the purposes of deriving the optimal mechanism, the most convenient way to understand the problem is that buyers face endogenous outside options which depend both on their values upon arrival and on the nature of the mechanism they face if they decide to delay participation. The approach is then to consider a “relaxed program” that takes into account the endogenous participation constraints only for buyers with low initial values. The approach yields a complete solution, with closed-form expressions for the relevant prices. A companion paper, Garrett (2011), shows how a similar approach can be applied in an environment where buyers have a continuum of values and where values can evolve according to a rich class of stochastic processes. However, it only achieves a partial characterization and can confirm the validity of the relaxed program only in particular instances. The analysis in the companion paper suggests nonetheless that the principles and findings described here apply more generally.

Two important observations about the dynamics of the optimal contracts are as follows. First, the buyer always receives the good eventually, even if his value remains low. This reflects a general principle for contracts with agents who have partially persistent private information: contractual relationships become more efficient as time passes. This is the “principle of vanishing distortions” explained, for example, by Battaglini (2005). Second, the later the buyer purchases the option contract, the longer he waits to obtain the good if his value remains low. That is, contracts become more distorted the later they are signed. This

is a new principle relative to the two strands of literature discussed above. Indeed, imposing tougher terms at later contracting dates, which require buyers to wait longer to obtain the good, is the seller's key tool for punishing those who hide their presence and participate later than their arrival.

Now return to considering the optimal path of posted prices. This restricted mechanism is conceptually simpler and so is treated first in the analysis. Rather than selling individualized price paths to buyers, as in the unrestricted optimum, the price buyers face on each date is constrained to be the same regardless of their moment of arrival.

To study the optimal path of posted prices, I focus on the allocation rule. Recognizing that optimal prices must induce immediate purchase whenever buyers have high values, this requires considering the set of dates at which buyers purchase if their values are low. This set is termed the "sales policy".

I first determine the price path that is optimal conditional on a given sales policy. Consider an interval over which buyers purchase only if their values are high, followed by a date at which they also purchase if their values are low (a "sale"). The price declines over the interval since the sales date must have a relatively low price, and because the option of waiting until this date to purchase given a high value becomes more valuable as the date approaches. That the option of waiting becomes more valuable reflects both discounting and the reduced probability that the value switches to low whilst waiting.

Now consider the optimal choice of sales policy. The finding of occasional sales or price cycles corresponds to a sales policy that includes intermittent dates. A crucial consideration explaining the optimality of price cycles is that, if a buyer's value is high at a given date, it may not remain high at future dates in the sales policy. Since it is opportunities to buy at dates in the sales policy when their values are high that give rise to the rent earned under the optimal price path, a "sparse" sales policy reduces the rents buyers can expect to earn. Importantly, because buyers understand that their values may become low when they are high, they are effectively "impatient" to obtain the good quickly, and this impatience can be exploited most effectively if the opportunities to purchase at a low price in the future are limited. On the other hand, delayed purchase when buyers' values are low is inefficient. The optimal policy reflects a trade-off between the concern for efficiency and the concern for rent extraction.

The optimality of price cycles is not driven by any particular assumption on the distribution of buyer values upon arrival. My focus is on showing that price cycles are optimal in case the evolution of buyer values is stationary, which will be the case provided buyers' initial values are drawn from the stationary

(i.e., ergodic) distribution. Considering a large population of buyers, the assumption implies (at least approximately) that the proportion of buyers with a high value in the population remains constant with time; in the words of Biehl (2001, p 568), there is “*individual* demand uncertainty” whilst the composition of “*aggregate* demand is stationary”.

Although the evolution of values is stationary, the probability that a given buyer has arrived (equivalently, the size of the population that has arrived) increases with the date. As a result, the trade-offs involved in determining whether to hold a sale change with time. In particular, since a low price at any date increases the option value of delaying purchase at all previous dates, it becomes less desirable to hold sales later on. This reflects the same principle as described above for the fully-optimal mechanism: the seller optimally punishes buyers for late arrival in order to reduce the rents left to early arrivers. As a result, after a sufficiently long time, optimal prices remain at a constant level such that buyers purchase only if their values are high. A related finding, at least in numerical examples, is that the duration between sales increases over time. The last two observations imply that, although the optimal price path is non-monotone, there is an upward trend.

The upward trend in prices seems inconsistent with price patterns in many durable-good markets. For example, Pashigian and Bowen (1991), in their study of apparel pricing, find it difficult to explain the use of discounts at the beginning of the season, but also note that prices are on average lower at the end of the season. Taking the beginning of the season as date zero, my baseline model seems to suggest an explanation for the first observation but not the second. In my model, discounting is attractive early in the season precisely because it only implies a rent for buyers who arrive early, but the same logic suggests prices should tend to rise with time. The theory does become consistent with a price path that eventually declines, however, if one supposes that the value the buyer obtains from the product is smaller at later dates. This assumption seems to make sense in fashion markets (which, broadly defined, include markets for consumer electronics) since consumers anticipate the value of their purchase will be eroded by the new fashion they expect to come on the market. By allowing the possibility of declining values, I am able to rationalize both the use of initial and periodic discounts as well as prices that eventually decline.

The finding in the baseline model that sales eventually cease completely turns out to depend on the assumption that buyers remain in the market forever. More realistically, one would expect buyers may leave the market. This could reflect that they have no further use for the good, but it might also reflect

an inability to continue waiting. In either case, I show that price cycles may persist indefinitely, with the pattern becoming stationary over time. This provides a notion of “steady state” pricing which may be appropriate for goods that have been offered for a long time.

Another important empirical regularity that has created a puzzle for theory is countercyclical markups. Markups of prices over costs have often been observed to be lowest during booms, and, at the micro level, at times when shoppers are particularly active (e.g., weekends and Christmas). There are a number of competing theories, although none seem entirely satisfactory (see Chevalier et al. (2003)). The prediction of my model is that periods during which the rate of buyer arrival is high are more likely to be associated with lower prices. During these times, the concern for limiting the rents of buyers who arrive earlier is less important relative to the objective of obtaining profit from buyers arriving at the time in question. For the parameter values of interest, the latter concern requires lowering the price.

The rest of the paper is set out as follows. Next, I review the relevant literature. Section 2 introduces the model, Section 3 examines the optimal path of posted prices, and Section 4 examines the fully-optimal mechanism and compares it to the optimal posted prices. Section 5 concludes. The Appendix collect proofs for the various results and is split into two parts: Appendix A contains proofs of results on posted prices, and Appendix B contains those for the fully-optimal mechanism.

1.1 Related literature

The economics literature has frequently enquired whether a monopolist selling a durable good to buyers who arrive over time should occasionally discount its products. In what might be described as the “standard framework” (McAfee and te Velde (2006)), the answer is that it should not.⁵⁶ The mechanism the seller should commit to is a constant price equal to the monopoly level in the static problem. The standard framework includes at least the following assumptions: (i) once having entered the market, buyers’ values do not change, (ii) the distribution of buyer values does not depend on the time of arrival in the market, (iii) there is no rivalry for the good (i.e., the seller has at least as many units available as there are buyers), (iv) the seller’s commitment power is absolute, and (v) there is no competition among sellers.

This paper is the first one to examine the role of stochastically changing values in the above framework

⁵McAfee and te Velde argue that this framework might apply to airline prices. But they remark: “The extent of price changes found in actual airline pricing is mysterious because a monopolist with commitment ability, in a standard framework, doesn’t want to engage in it at all!”

⁶See Riley and Zeckhauser (1983) and Conlisk et al. (1984) for early discussions of the optimality of a constant price when the seller can commit and buyers arrive over time.

(hence relaxing (i)). In order to highlight the implications of this assumption, I deliberately maintain the other assumptions (ii)-(v). Other authors, however, have examined their role. For example, Board (2008) characterizes price fluctuations when the values of different cohorts of buyers are distributed differently (relaxing (ii)). As explained above, I focus instead on an environment which is stationary: a buyer's value at any date is independent both of his arrival date and of the time since his arrival.

A growing literature in operations research and economics (examples include Wang (1993), Gallego and Van Ryzin (1994), Das Varma and Vettas (2001), Gallien (2006), Board (2008), Gershkov and Moldovanu (2009), Board and Skrzypacz (2010), Pai and Vohra (2011) and Said (2011)) considers the problem of selling to buyers who arrive over time under the assumption of scarcity (relaxing (iii)). One possibility (see, e.g., Gallego and Van Ryzin and Gershkov and Moldovanu), given both scarcity and a deadline for the seller, is that prices fall as the deadline approaches but jump whenever a unit is purchased. Although the resulting price path may be non-monotone, prices always decline monotonically (or remain constant) conditional on the inventory level in the specifications examined in these papers. Because price fluctuations in my model do not depend on changes in the inventory level, the two theories should be easily distinguished empirically.

Conlisk et al. (1984) and Sobel (1991) consider an environment where the seller cannot commit (hence relaxing (iv)) and argue that equilibrium often involves periodic sales.⁷ This paper, in contrast, seeks an understanding of the optimal choice of mechanism for firms which can establish credible commitments. Commitment seems possible, for instance, when firms can establish reputations for their pricing practices.⁸

Many of the earlier theories of sales are based on price dispersion arising due to competition between firms (see Salop and Stiglitz (1977), Shilony (1977), Rosenthal (1980) and Varian (1980)). For the most part, these models are entirely static, and it is only informally argued that they might have implications for the patterns of prices over time.⁹ Maskin and Tirole (1988) study how “Edgeworth cycles”, or cycles of price wars, can arise in a duopoly model with partial commitment to future prices. In contrast to these theories, where competition between firms is essential, I find price fluctuations in the absence of competition.

The possibility of changing values has also been given some attention. Stokey (1979) shows that committing to a declining price path can be optimal when buyers' values change gradually and deterministically with time. Board and Skrzypacz (2010) show that prices decline over time when the good expires at a known

⁷Lack of commitment is also central in a number of other papers which find periodic discounting. Examples are Sobel (1984), Pesendorfer (2002), Villas-Boas (2004), Chevalier and Kashyap (2010) and Eichenbaum et al. (2011).

⁸See Waldman (2003) for a review of the treatment of commitment in the analysis of durable-goods markets.

⁹An exception is Rosenthal (1982).

date. Deb (2010) finds an increasing price path in a model where buyers' values change stochastically.¹⁰ Although these papers often interpret the low point of the price path as a "sale", none of them explain why prices should vary non-monotonically, even less why they should fluctuate repeatedly.¹¹

As explained above, the paper makes an important contribution to the literature on optimal dynamic mechanisms for agents with stochastically evolving types. This is to allow agents' arrival times both to be stochastic as well as their private information. All of the other papers — for instance, Baron and Besanko (1984), Besanko (1985), Courty and Li (2000), Battaglini (2005), Board (2007), Eso and Szentes (2007), Zhang (2009), Pavan et al. (2011) and Garrett and Pavan (2011) — take it as given that agents are available to contract at the time the principal commits to the dynamic mechanism. Equivalently for the purposes of analysis, it is an implicit assumption that the principal perfectly observes agents' arrival times. What these papers therefore miss is an account of how the contracts the principal offers can be expected to evolve with time.

I do rely on some of the methods developed in the earlier papers; in particular, the use of a continuous-time framework with Markov switching between two values is the same as in Zhang's work on dynamic taxation. The reason for choosing this special structure is both that it facilitates the study of the optimal price path and that it allows me to completely characterize the fully-optimal mechanism. The companion paper, Garrett (2011), discussed in more detail in Section 4, studies the fully-optimal mechanism in a richer discrete-time environment, but is unable to achieve a complete characterization. It tends to confirm the key findings that buyers with low initial values expect strictly positive rent and that buyers who participate later in the mechanism receive tougher contractual terms and expect to wait longer to receive the good. In follow-on work, Ely et al. (2011) apply some related ideas to the problem of overbooking in the airline industry, considering a restricted class of mechanisms.¹²

Finally, one might conjecture that some of the work that does not consider dynamic arrivals could accommodate the possibility directly by including "latent types" who have especially low values until, according to some stochastic process, the buyer "arrives", i.e. has a relatively high value at which he might conceivably

¹⁰Contrary to my paper and to the "standard framework" described above, the buyer in Deb's model is known to be in the market at the time the seller commits to a price path. Prices increase monotonically in his model, but this seems to rely on the assumption that either (i) the buyer's value changes only once, or (ii) the buyer's value continues to change but any purchase must be made on the date of a shock to his value.

¹¹A number of other papers find prices that rise or decline monotonically for other reasons. A few examples include Lazear (1986), Courty and Li (1999), Su (2007), Nocke and Peitz (2007), and Horner and Samuelson (2011).

¹²The class of mechanisms studied in that paper is motivated by the idea that a "ticket" should be thought of as an option to fly, which might be given up or retained on the date of travel.

purchase. This kind of approach is mentioned, for instance, by Bergemann and Valimaki (2010) for the problem of maximizing efficiency. However, the approach does not apply for profit maximization (assuming no restriction on mechanisms), where the time at which the buyer is *available to contract* is critical. In particular, the ability to contract with latent types at the time the mechanism is determined would allow the seller to extract more rent and would imply a different allocation rule. For example, if the buyer is always a latent type at date zero and has no private information at that time, then the profit-maximizing mechanism involves date-zero contracting, full efficiency and zero expected rent for the buyer.

2 Model

Players. There is a single seller and a single representative buyer. All of the results will apply also if there is a population of buyers, provided the problem remains separable across them, i.e. provided that all buyers are ex-ante identical and have independently distributed information.

Arrival. My analysis is chiefly concerned with the case where the buyer’s arrival time τ is distributed exponentially with parameter $\lambda > 0$ (I also consider the case where the seller knows the buyer’s arrival date in advance; i.e., where there is no uncertainty about the date of arrival). Note that this process is memoryless; i.e., the distribution of the time until the buyer’s arrival, conditional on his having not yet arrived, remains the same at each date. Before arrival to the market, any communication with the seller is taken to be impossible. This is simply because the buyer is unable to observe any offers or communication from the seller until after arrival.

Payoffs. The buyer and seller are risk neutral and have a common discount rate $r > 0$. The buyer has unit demand and the seller’s cost of production is normalized to zero.

There is no re-sale; i.e., the good is “sold outright” and, after purchase, the buyer keeps it forever. More generally, one might want to allow that the buyer can return the good to the seller. In Section 3, where I restrict attention to outright sale at posted prices, the seller cannot benefit from instead posting prices for rental of the good (this possibility is discussed further in Section 3 and has been discussed in the context of a related model by Biehl (2001)). In Section 4, where the seller can use any mechanism, there is also no possibility to benefit from temporary allocation of the good (see footnote 35). These observations, however, would not be expected to apply in more general environments; e.g., environments with rivalry for the good among buyers or where there are more than two distinct values for the good.

Abusing notation, the buyer's flow value of holding the good is $\omega_t \in \{\omega_L, \omega_H\}$, where $\omega_H > \omega_L > 0$. Thus the buyer's value at time t is either "low" (the flow consumption value equals ω_L) or "high" (the flow consumption value equals ω_H). The buyer's expected value from obtaining the good will be determined by the evolution of values described below.

Evolution of values. Upon arrival at any date τ , the buyer has a high value with probability $\gamma \in (0, 1)$ and a low value with probability $1 - \gamma$. Values change at a Poisson rate: the value switches from low to high at rate α_L and from high to low at rate α_H . Unless otherwise specified, these values are strictly positive. The initial distribution is then "stationary" in case γ is equal to $\gamma^S \equiv \frac{\alpha_L}{\alpha_L + \alpha_H}$. This means that the probability the buyer has a high value at any date does not depend on the time since arrival.

Denote by θ_t the buyer's date- t expected discounted value from owning the good forever conditional on his date- t flow value. Again abusing notation, if the buyer's flow value at t is high, this is equal to¹³

$$\theta_H = \mathbb{E} \left[\int_t^\infty e^{-r(s-t)} \tilde{\omega}_s ds \mid \tilde{\omega}_t = \omega_H \right].$$

If his flow value at t is low, it is equal to

$$\theta_L = \mathbb{E} \left[\int_t^\infty e^{-r(s-t)} \tilde{\omega}_s ds \mid \tilde{\omega}_t = \omega_L \right].$$

Note that

$$\begin{aligned} \theta_H &= \frac{(\alpha_L + r)\omega_H + \alpha_H\omega_L}{r(\alpha_L + \alpha_H + r)}, \text{ and} \\ \theta_L &= \frac{\alpha_L\omega_H + (\alpha_H + r)\omega_L}{r(\alpha_L + \alpha_H + r)}. \end{aligned}$$

Clearly, the entire analysis may be carried out by reference to the values θ_L and θ_H . However, it is often helpful to refer to the flow values ω_L and ω_H .

A sample path for the process from the buyer's arrival time τ is denoted $\theta^{[\tau, \infty)}$, where, for any $s \geq \tau$, $\theta^{[\tau, \infty)}(s) = \theta_s \in \{\theta_L, \theta_H\}$. A history of values over any interval $I \subset \mathbb{R}_+$ is recorded by letting, for any $s \in I$, $\theta^I(s) = \theta_s \in \{\theta_L, \theta_H\}$. Throughout, if the buyer arrives at date τ , then at any subsequent date $t \geq \tau$, the history of values $\theta^{[\tau, t]}$, as well as the arrival time, remains his private information.

Finally, one statistic which shall be used frequently in the analysis is the probability that the buyer's

¹³Throughout, I denote random variables using tildes.

value is high conditional on an earlier value. For any dates $s, t, s > t$,

$$\begin{aligned} \Pr\left(\tilde{\theta}_s = \theta_H | \tilde{\theta}_t = \theta_H\right) &= \frac{\alpha_L + \alpha_H e^{-(\alpha_L + \alpha_H)(s-t)}}{\alpha_L + \alpha_H}, \text{ and} \\ \Pr\left(\tilde{\theta}_s = \theta_H | \tilde{\theta}_t = \theta_L\right) &= \frac{\alpha_L (1 - e^{-(\alpha_L + \alpha_H)(s-t)})}{\alpha_L + \alpha_H}. \end{aligned}$$

For clarity, I will often leave expressions in terms of these conditional probabilities.

3 Posted prices

This section considers profit-maximizing posted prices. Because my main objective is to demonstrate the optimality of price fluctuations in an environment which is stationary, I focus on the case in which the distribution of initial values is the stationary one, i.e. $\gamma = \gamma^S$. However, the finding that optimal prices may fluctuate does not depend upon this restriction.

I require that the seller commits at date zero to a price path $p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ with the implication that, if the buyer purchases the good at date t , he pays price $p(t)$. The buyer perfectly observes the price path upon arrival to the market.¹⁴ The price the seller charges at any date cannot depend on communication by the buyer (e.g., communication that signals his arrival).¹⁵¹⁶ The inability to adjust prices in response to communication simply reflects the prevailing norm in retail markets. The restriction to deterministic price paths does not harm the seller.¹⁷

The 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 $\sigma(\theta_L, \cdot), \sigma(\theta_H, \cdot) : \mathbb{R}_+ \rightarrow \{0, 1\}$. The set Σ is taken to be the set of feasible strategies for the buyer. The restriction to Markov strategies is without loss of generality given that (i) the buyer's problem is Markov (i.e., dependent only on his current value for the good, and

¹⁴Perfect observability, however, is not essential; all that is required is that, upon arrival to the market, the buyer correctly anticipates that he can purchase the good from the seller and the corresponding price path p . For instance, it is enough that the seller sufficiently values that the buyer perceives the choice of full-commitment prices when he purchases the good. This may be the case in the context of a repeated game involving the seller and a population of buyers who can communicate their purchase prices among one another.

¹⁵Although I do not explore the issue of when posted prices might be the optimal institution, note that the mechanism described above is optimal in an environment in which the entire transaction (in particular, both communication and delivery of the good) must take place at a single instance selected by the buyer. This seems a reasonable approximation for inexpensive goods, where the cost of repeated interactions between the buyer and seller is large relative to the size of the information rents at stake.

¹⁶Although not immediate, it turns out that the seller can do no better by committing to the infinite price path at a date later than zero (say after the buyer has arrived and made some communication with the seller). This is a straightforward consequence of Result 1 in the next subsection.

¹⁷The 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 (see the expression for the seller's total profit (23) in the Appendix).

not on past values or time of arrival; this means that the value of the buyer’s problem can be attained by a Markov strategy), and (ii) the seller can never profit by asking the buyer to follow a strategy that is not Markov.¹⁸ For any arrival date τ and any $\theta^{[\tau, \infty)}$, the date of purchase for the buyer when using the strategy $\sigma \in \Sigma$ is $\mu_\sigma(\theta^{[\tau, \infty)}) \equiv \inf \{s \geq \tau : \sigma(\theta_s, s) = 1\}$. In case $\mu_\sigma(\theta^{[\tau, \infty)}) = +\infty$, the interpretation is that the buyer never purchases.

A “posted-price mechanism” $\mathcal{M}_P = \langle p, x \rangle$ is fully specified by appending to the price path p the prescription $x \in \Sigma$ of whether the buyer is to purchase for each value at each date (without loss of generality, the seller specifies x also at date zero). The interpretation is that the 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 , (b) he has not yet purchased the good, and (c) his value for the good at date t is θ_t , with $x(\theta_t, t) = 1$. If the buyer finds it optimal to follow this prescription given p , then the mechanism is said to be incentive compatible.

3.1 Benchmarks

Known arrival date. A natural starting point for the analysis is to consider an environment which is simpler to analyze, namely one where the seller knows the buyer’s arrival date in advance. The date of arrival can be normalized to zero (equivalently, one can think of the buyer as always being active and receptive to the seller’s offers, so that the buyer is active in the market at the time the seller has the opportunity to commit to a path of prices).

Profit-maximizing posted prices for buyers with known arrival dates and stochastically changing values are analyzed by Conlisk (1984), Biehl (2001) and Deb (2010). The analysis here is best thought of as extending Biehl’s (see Proposition 2 of his paper) from a setting with two periods to one with an infinite horizon (I also generalize his model by allowing $\alpha_L \neq \alpha_H$). The discussion is somewhat heuristic, but all results can be derived from arguments made formally in the next subsection.

The kinds of price paths the seller might choose in the environment with a known arrival date can be divided between situations in which the buyer purchases only if his value is high (Case 1) and those where the buyer purchases at some point also if his value is low (Case 2). In Case 1, the best the seller can do is to set a constant price θ_H , with the buyer purchasing either on date zero if his initial value is high, or on the first date that his value turns high. In Case 2, if the buyer finds it optimal to purchase with a low value on some date z , then it is easy to show that he must find it optimal to purchase at that date also if his value

¹⁸The restriction to deterministic strategies is also without loss of generality (see the previous footnote).

is high (see Lemma 1 below). As a consequence, no purchases occur in equilibrium after date z , and any price path that remains above θ_H after date z is consistent with profit maximization.

For Case 2, the highest price the seller can charge on date z consistent with purchase by the buyer when his value is low is θ_L . It is easy to see that the optimal path of prices between date zero and date z is the one which induces the buyer to purchase if his value is high, but ensures his expected rent is the same as if he were to purchase at date z . Thus, for all $t \in [0, z]$, let

$$p(t) = \theta_H - \Pr\left(\tilde{\theta}_z = \theta_H | \tilde{\theta}_t = \theta_H\right) e^{-r(z-t)} (\theta_H - \theta_L). \quad (1)$$

The profit the seller expects for a given date z is then

$$\begin{aligned} R(z) &= \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{\alpha_L + r} \right) \theta_H + (1 - \gamma^S) e^{-(\alpha_L + r)z} \left(\theta_L - \frac{\alpha_L \theta_H}{\alpha_L + r} \right) \\ &\quad - \gamma^S e^{-rz} (\theta_H - \theta_L). \end{aligned} \quad (2)$$

The first term is the expected surplus when the buyer obtains the good only if his value is high. The second term is the expected gain in surplus from instead having the buyer obtain it at date z if his value begins low and remains so until this date.¹⁹

The final term is the buyer's expected rent. This can be decomposed according to whether the initial value is low or high: conditional on a high initial value it is

$$\Pr\left(\tilde{\theta}_z = \theta_H | \tilde{\theta}_0 = \theta_H\right) e^{-rz} (\theta_H - \theta_L),$$

and conditional on a low initial value it is

$$\Pr\left(\tilde{\theta}_z = \theta_H | \tilde{\theta}_0 = \theta_L\right) e^{-rz} (\theta_H - \theta_L). \quad (3)$$

The former is decreasing faster than the rate of discounting (since $\Pr\left(\tilde{\theta}_z = \theta_H | \tilde{\theta}_0 = \theta_H\right)$ is decreasing in z), but the latter is either increasing, or it is decreasing slower than the rate of discounting (since $\Pr\left(\tilde{\theta}_z = \theta_H | \tilde{\theta}_0 = \theta_L\right)$ is increasing in z). For the following result, what it is enough to notice is that the buyer's ex-ante expected rents are exactly proportional to the discount factor e^{-rz} .

¹⁹This term can be understood better by noting that $\theta_L - \frac{\alpha_L \theta_H}{\alpha_L + r} = \frac{\omega_L}{\alpha_L + r}$, which is the present value of an annuity that pays ω_L until termination at rate α_L , when the discount rate is r . After discounting, this annuity coincides with the difference in surplus between the two purchasing policies given that the buyer's value begins low and remains so until date z .

Result 1 *Suppose the seller knows that the buyer arrives at date zero. If*

$$\theta_L > \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{\alpha_L + r} \right) \theta_H, \quad (4)$$

it is optimal to set $p(0) = \theta_L$ and $p(t) = \theta_H$ for all $t > 0$; the buyer purchases immediately irrespective of his value. Otherwise, it is optimal to set $p(t) = \theta_H$ for all t ; the buyer purchases if and only if his value is high (and he has not yet purchased a unit).

The interesting qualitative result is that inducing the buyer to purchase with a low value at a positive date, i.e. choosing $z > 0$ in Case 2, is never optimal. As noted above, increasing z reduces the buyer's expected rents by a factor of e^{-rz} , reflecting discounting. However, it reduces the gain in surplus from having the buyer purchase with a low value by a factor of $e^{-(\alpha_L+r)z}$, reflecting both discounting and the probability that the buyer with a low initial value purchases before date z .

A similar insight is available when buyer values are constant, as first analyzed by Stokey (1979) (note that, in case of constant values, the fully-optimal mechanism coincides with the posted price mechanism). In this case, letting γ^C be the probability of a high value, it is optimal to have the buyer purchase immediately irrespective of his value if $\theta_L > \gamma^C \theta_H$, and to have him purchase immediately only if his value is high otherwise, with the low-value buyer never purchasing. In particular, any date on which the buyer purchases must be date zero. The optimal date-zero prices are θ_L and θ_H respectively.

Uncertain arrival date. Now consider the possibility that the buyer's arrival date is uncertain (say, drawn from the exponential distribution defined in the model set-up, although the precise distribution is not important) and is the buyer's private information. If buyer values are constant, it is easy to see that the seller should simply choose a constant price θ_L for all of time in case $\theta_L > \gamma^C \theta_H$ and should choose a constant price θ_H for all of time otherwise. The buyer finds it optimal to make any purchase immediately and, in this sense, the solution replicates the outcome for the problem where the buyer's arrival date is known in advance.

When values change stochastically, things are different and more interesting. To see this, suppose that (4) holds; then, the above logic does not apply. Choosing a constant price of θ_L does not induce purchase by the buyer when his value is low. The reason is simply that the buyer expects zero rent from purchasing if his value is low. He can do better by delaying purchase until his value changes to θ_H . At such time, his expected value of holding the good forever, compared to the price paid, is positive. This means that the

option value of waiting to satisfy his unit demand is positive.

It turns out that the highest constant price the seller can charge to induce purchase when the buyer's value is low is equal to $\frac{\omega_L}{r}$, the discounted value of receiving flow payoff ω_L in perpetuity. To see this, suppose the constant price is p^C . Suppose the buyer's value is low at t' and remains so until $t'' > t'$ and compare the strategy of purchasing the good at date t' with that of purchasing instead at date t'' (say, when the buyer's value changes to high). The difference between the buyer's payoffs under the two strategies, as evaluated at date t' , is equal to

$$\left(\frac{\omega_L}{r} - p^C\right) \left(1 - e^{-r(t'' - t')}\right).$$

Therefore, if $p^C > \frac{\omega_L}{r}$, the buyer should clearly purchase only once his value turns high.

Given that setting a constant price $\frac{\omega_L}{r}$ fails to extract any of the surplus anticipated from the buyer's value becoming high, it is not too surprising that such a price path turns out never to be optimal. What one should predict is that the seller must find ways to limit the rents the buyer expects to earn, without giving up entirely on the objective of inducing purchase when the buyer's value is low. The rest of this section shows how the seller optimally responds to these tradeoffs given that she is restricted to using posted prices. Of course, if the seller is not restricted in her choice of mechanism, she can also find other ways to reduce the buyer's rent (this is true whether or not she knows the buyer's arrival date). Section 4 will show how she chooses the mechanism in this case.

3.2 Optimal prices for a given sales policy

Allowing for the possibility that the buyer's arrival date is uncertain, this subsection derives profit-maximizing prices for given dates at which the buyer is induced to purchase if his value is low (i.e., for a given "sales policy"). I start by describing more carefully the buyer's problem when confronted with a mechanism \mathcal{M}_P .

Buyer's problem. The expected payoff to the buyer from using the Markov strategy $\sigma \in \Sigma$ at any date t at which he has not yet purchased is given by

$$\begin{aligned} u_t^{\mathcal{M}_P}(\theta_L; \sigma) &= \mathbb{E} \left[e^{-r(\tilde{\mu}_\sigma - t)} \left(\tilde{\theta}_{\tilde{\mu}_\sigma} - p(\tilde{\mu}_\sigma) \right) \mid \tilde{\theta}_t = \theta_L \right] \text{ and} \\ u_t^{\mathcal{M}_P}(\theta_H; \sigma) &= \mathbb{E} \left[e^{-r(\tilde{\mu}_\sigma - t)} \left(\tilde{\theta}_{\tilde{\mu}_\sigma} - p(\tilde{\mu}_\sigma) \right) \mid \tilde{\theta}_t = \theta_H \right], \end{aligned}$$

where $\tilde{\mu}_\sigma$ is the stopping time determined by σ .

Incentive compatibility of the posted-price mechanism $\mathcal{M}_P = \langle p, x \rangle$ is then the requirement that, for

each t and θ_t ,

$$u_t^{\mathcal{M}_P}(\theta_t; x) = \nu_t^{\mathcal{M}_P}(\theta_t)$$

where

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

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

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

Therefore, the present value of total expected profit is

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

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

Sales policy. Although, in this section, I am considering only a restricted set of mechanisms, I adopt the approach common in the mechanism design literature of formulating the seller's objective in terms of the allocation rule rather than prices. For any incentive-compatible mechanism $\mathcal{M}_P = \langle p, x \rangle$, define the corresponding sales policy to be the set $A \subset \mathbb{R}_+$ of times t such that $x(\theta_L, t) = 1$. If A is the corresponding sales policy, then I say that the mechanism is "consistent with A " and denote it by $\mathcal{M}_{P,A} = \langle p_A, x_A \rangle$. The collection of feasible sales policies is simply the collection of left-closed subsets of the positive reals.²⁰

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 is²¹

$$m_A(t) \equiv \inf \{s \in A : s \geq t\}. \quad (5)$$

Thus, $m_A(t) = +\infty$ simply states that A is bounded below t (below I use the convention that, for any function $g : \mathbb{R} \rightarrow \mathbb{R}$, $g(+\infty) = \lim_{s \rightarrow +\infty} g(s)$).

²⁰By left-closed it is meant that the limit of any decreasing sequence of points in A is contained in A .

²¹Note that because A is left-closed, it is measurable. Hence m_A is a measurable function.

Value of buyer's problem and optimal prices given sales policy. The first step to finding conditionally optimal prices is to deduce for a given sales policy A certain conditions that the value functions $\left(\nu_t^{\mathcal{M}_{P,A}}(\cdot)\right)_{t \geq 0}$ must satisfy. In particular, I establish tight lower bounds for these functions, which will turn out to coincide with the rents the buyer expects to earn for the price path that maximizes profits given A .

The following lemma establishes a monotonicity property of the allocation rule associated with any mechanism $\mathcal{M}_{P,A}$, which will be helpful in determining the aforementioned lower bounds.

Lemma 1 *For any sales policy A , any mechanism $\mathcal{M}_{P,A}$, and any $t \in A$, (i) $x_A(\theta_H, t) = 1$, and hence (ii) $\nu_t^{\mathcal{M}_{P,A}}(\theta_H) = \nu_t^{\mathcal{M}_{P,A}}(\theta_L) + \theta_H - \theta_L$.*

The property is simply that, if the buyer is willing to purchase when his value is low, he must also be willing to purchase if his value is high. The additional rent he earns is the difference in the expected value from holding the good conditional on a high or low value at the time of purchase.

For any sales policy A , I now derive a lower bound on the buyer's expected rent in mechanisms consistent with A . To do so, I first restrict attention to the problem where the buyer can only employ a limited set of strategies (since this can only lower his achievable rent). In particular, I restrict attention to strategies where the buyer is unable to make purchases except at dates in the sales policy (i.e., to $\{\sigma \in \Sigma : \sigma(\theta_t, t) = 0 \text{ for all } t \notin A \text{ and all } \theta_t \in \{\theta_L, \theta_H\}\}$), and define $\nu_t^{\mathcal{M}_{P,A,rest.}}(\theta_t)$ to be the value of the buyer's problem given the restricted strategies for each t and each θ_t .

Consider any sales policy A . By Lemma 1, given a mechanism $\mathcal{M}_{P,A}$, the buyer finds it optimal to purchase at all dates in A irrespective of his value. Therefore, at any date t ,

$$\begin{aligned} \nu_t^{\mathcal{M}_{P,A,rest.}}(\theta_L) &= e^{-r(m_A(t)-t)} \left(\nu_{m_A(t)}^{\mathcal{M}_{P,A,rest.}}(\theta_L) + \Pr\left(\tilde{\theta}_{m_A(t)} = \theta_H | \tilde{\theta}_t = \theta_L\right) (\theta_H - \theta_L) \right), \text{ and} \\ \nu_t^{\mathcal{M}_{P,A,rest.}}(\theta_H) &= e^{-r(m_A(t)-t)} \left(\nu_{m_A(t)}^{\mathcal{M}_{P,A,rest.}}(\theta_L) + \Pr\left(\tilde{\theta}_{m_A(t)} = \theta_H | \tilde{\theta}_t = \theta_H\right) (\theta_H - \theta_L) \right). \end{aligned}$$

Evaluating the probabilities,

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

The buyer, if his value is low at date t , anticipates that he can obtain at least the rent from avoiding all purchases until such time as his value becomes high, and then following a (restricted) optimal strategy from

then on. For any $s \geq t$, this is given by $\nu_s^{\mathcal{M}_{P,A,rest.}}(\theta_H)$ if his value changes to high at date s . Therefore,

$$\begin{aligned} \nu_t^{\mathcal{M}_{P,A,rest.}}(\theta_L) &\geq \int_t^\infty \alpha_L e^{-(r+\alpha_L)(s-t)} \nu_s^{\mathcal{M}_{P,A,rest.}}(\theta_H) ds \\ &= \int_t^\infty \alpha_L e^{-(r+\alpha_L)(s-t)} \left(\nu_s^{\mathcal{M}_{P,A,rest.}}(\theta_L) + e^{-(r+\alpha_L+\alpha_H)(m_A(s)-s)} (\theta_H - \theta_L) \right) ds. \end{aligned} \quad (6)$$

A lower bound for the buyer's rents is obtained by supposing that the buyer is in fact willing to wait until his value changes to high before making future purchases, i.e. that (6) holds with equality. The unique solution to this equation consistent with the transversality condition that buyer rents remain bounded is²²

$$\nu_t^{LB,A}(\theta_L) = \int_t^\infty e^{-r(s-t)} \alpha_L e^{-(r+\alpha_L+\alpha_H)(m_A(s)-s)} (\theta_H - \theta_L) ds.$$

This equation, together with

$$\nu_t^{LB,A}(\theta_H) = \nu_t^{LB,A}(\theta_L) + e^{-(r+\alpha_L+\alpha_H)(m_A(t)-t)} (\theta_H - \theta_L),$$

defines a lower bound on the agent's payoffs in an incentive-compatible mechanism consistent with A , not only for the restricted strategies considered, but also for the unrestricted strategies Σ .

The next result gives a mechanism consistent with a sales policy A for which the lower bounds coincide with the value function of the agent. The idea is to choose this mechanism so that, subject to consistency with A , both (i) the lower bound on rents is attained, and (ii) surplus is maximized, i.e. the buyer purchases the good whenever his value is high. I then show in the Appendix that the proposed mechanism is incentive compatible. To simplify this proof, I restrict attention to sales policies A which are at most countable collections of intervals and points. In the following subsection, I show that restricting attention to these kinds of sales policies is without loss of optimality.

²²The solution is found by considering the linear first-order differential equation

$$r\nu_t^{LB,A}(\theta_L) = \frac{\partial \nu_t^{LB,A}(\theta_L)}{\partial t} + \alpha_L e^{-(r+\alpha_L+\alpha_H)(m_A(t)-t)} (\theta_H - \theta_L).$$

Lemma 2 *Let A be any sales policy that is an at most countable union of intervals and points. Define the price path p_A^* by*

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

if $t \in A$ and by

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

if $t \notin A$. Choose the purchase prescription x_A^ consistent with A and such that the buyer purchases immediately on all dates such that his value is high (i.e., $x_A^*(\theta_H, t) = 1$ for all t). Then, $\mathcal{M}_{P,A}^* = \langle p_A^*, x_A^* \rangle$ maximizes the seller's expected profit conditional on consistency with 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 the 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. Second, at any date $t \notin A$ at which the 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 again if there are no future dates ($m_A(t) = +\infty$).

Before moving to consider optimal sales policies, it is useful to consider more carefully how prices are affected by the choice of sales policy. Suppose firstly that $[t', t''] \subset A$, with $t' < t''$. Then, using

$$\nu_{t'}^{LB,A}(\theta_L) = e^{-r(t''-t')} \nu_{t''}^{LB,A}(\theta_L) + \left(1 - e^{-r(t''-t')}\right) \frac{\alpha_L (\theta_H - \theta_L)}{r}, \quad (7)$$

one can easily deduce that

$$p_A^*(t') - e^{-r(t''-t')} p_A^*(t'') = \left(1 - e^{-r(t''-t')}\right) \frac{\omega_L}{r}. \quad (8)$$

This simply states that the difference in the value of payments at t' from purchasing at t' rather than t'' equals the value of holding the good between t' and t'' when the buyer's value remains low. This reflects the same principle described in the previous subsection for constant prices.

Second, suppose that $t', t'' \in A$ with $t' < t''$, but that $(t', t'') \cap A = \emptyset$. In this case

$$\nu_{t'}^{LB,A}(\theta_L) = e^{-r(t''-t')} \left(\nu_{t''}^{LB,A}(\theta_L) + \Pr\left(\tilde{\theta}_{t''} = \theta_H | \tilde{\theta}_{t'} = \theta_L\right) (\theta_H - \theta_L) \right) \quad (9)$$

i.e., if the buyer's value is low at t' , his rent must equal that from waiting to purchase at t'' , when there is a positive probability that his value is high.

Now compare the two scenarios, assuming that the dates t' and t'' , as well as the rent at date t'' (as determined by $\nu_{t''}^{LB,A}(\theta_L)$), are the same in each. It is easy to see that the buyer's rent at t' (as given by $\nu_{t'}^{LB,A}$) is strictly larger if the dates (t', t'') are included in A (as given by (7)) than if they are excluded (as for (9)). Thus the price at t' in the first scenario is also lower. The reason for this result is that the option value of waiting to purchase at date t' is greater if prices are known to be lower in the future, i.e. if there are more future dates in the sales policy A .

Finally note that, for optimal prices, a given change in the rents expected at any date in the sales policy (e.g., consider a change in the sales policy after that date) affects the buyer's rents at all earlier dates in present value (i.e., date zero) terms by the same amount. This observation plays a crucial role in determining the key features of the optimal sales policy, which is characterized in the next subsection. A similar observation will apply also when I turn to discuss the fully-optimal mechanism in the next section.

3.3 Optimal sales policy

Characterization of optimal sales policy. This subsection considers profit-maximizing sales policies when the buyer's arrival date is uncertain and determined according to the exponential distribution defined in the model set-up. I next give the main characterization result. To do so, I restrict attention to parameters such that the inequality (4) holds, since for the remaining parameters it is optimal simply to choose a constant price equal to θ_H . To simplify the statement of the result, and without loss of optimality, I restrict attention to sales policies that are closed sets.

Proposition 1 *Suppose that (4) holds.²³ An optimal sales policy exists. Let*

$$T = \begin{cases} 0 & \text{if } \omega_L \leq \gamma^S \omega_H \\ \frac{1}{\lambda} \log \left(1 + \frac{\lambda(\omega_L - \gamma^S \omega_H)}{\alpha_L(\omega_H - \omega_L)} \right) & \text{if } \omega_L > \gamma^S \omega_H \end{cases}. \quad (10)$$

Any optimal policy A^ satisfies*

- (i) *there exists $K \neq \emptyset$ such that $A^* = [0, T] \cup \{t_k : k \in K\}$, where $t_k > T$ for all $k \in K$, and*
- (ii) *for any $\eta > 0$, there exist at most finitely many points in A^* above $T + \eta$,²⁴ hence A^* is bounded.*

This result, together with the observations in Subsection 3.1, imply the following corollary for the environment with uncertain arrival.

Corollary 1 *If values for the good do not change, then optimal prices are constant over time. If values change stochastically (i.e., given $\alpha_L, \alpha_H > 0$), and if (4) holds, then optimal prices fluctuate.*

The proof of Proposition 1 proceeds by determining the seller's expected total profit for a given sales policy $PROF^{\mathcal{M}P,A}$ (see (23) in the Appendix) and then analyzing a dynamic program. This centers on a particular recursive statement of the problem, the details of which are delayed until the end of this subsection. Next, I provide an example (calculated by implementing the recursive approach numerically) to highlight further the key features of the optimal price path and then provide intuition for why these features arise.

Example 1 *Suppose that $\omega_L = 1\frac{1}{4}$, that $\omega_H = 2$, $\alpha_L = \alpha_H = \frac{1}{4}$, that $\lambda = r = \frac{1}{10}$ and that $\gamma = \gamma^S = \frac{1}{2}$. It follows that $\theta_L = 15\frac{5}{8}$ and $\theta_H = 16\frac{7}{8}$. I approximate A by considering only subsets of $\{n/100 : n \in \mathbb{N}\}$. The optimal sales policy and price path is depicted in Figure 1.*

Several features are apparent from considering Figure 1. As stated in Part (ii) of Proposition 1, the optimal sales policy is bounded, which means that prices are eventually constant and equal to θ_H . Before this time, prices fluctuate as they fall up to isolated points in the sales policy and jump thereafter. These "sales" events are anticipated by the proposition, for it states the optimality of a discrete sales policy above $T + \eta$ for any $\eta > 0$ (here, $T \approx 1.25$). Further, the sales policy includes an initial interval $[0, T]$ and then the

²³The result of Proposition 1 holds true also if $\gamma \neq \gamma^S$ provided that $\alpha_L, \alpha_H > 0$. Whilst (4) remains sufficient, a weaker condition often suffices for other values of γ . What is needed is simply that there be a positive probability that the buyer purchases with a low value under the optimal price path. None of the arguments given in the proof turn out to depend on the value of γ . An earlier working paper version provides more detailed discussion.

²⁴It can be shown, in addition, that if $\omega_L \geq \gamma^S \omega_H$, T must be an accumulation point of $A^* \setminus [0, T]$.

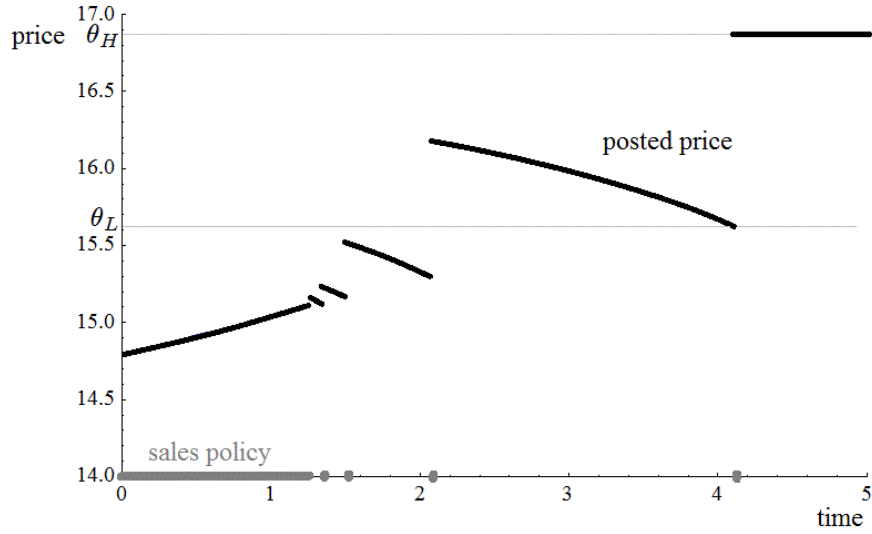


Figure 1: Optimal price path and sales policy for Example 1

subsequent dates are spaced further apart with time.²⁵ This, together with the boundedness of the optimal policy, means that prices trend upwards over time. I now provide intuition for these key features.

Boundedness of sales policy. That the optimal sales policy A^* is bounded follows for the reason discussed at the end of the previous subsection. Including additional dates in the sales policy after some date t implies higher buyer rents and lower prices at all dates before date t back to date zero. Obviously the expected effect on the buyer’s rent is therefore larger the later the date t . In addition, the probability of the buyer arriving on the market is smaller the later the date under consideration, so the efficiency gains from including additional dates in the sales policy are less. Both these effects work in the same direction. Together with the intuition discussed below, they also explain why the distance between dates in the sales policy grows with time. In Subsection 3.5, however, I explain that the sales policy may be unbounded and that the pattern of prices may converge to a stationary one in case the buyer leaves the market stochastically at a sufficiently fast rate.

Discreteness of sales policy after T (“price fluctuations”). Now consider the finding of isolated points in the optimal sales policy. In particular, it seems important to understand why it is optimal to induce the buyer to purchase when his value is low after a delay when the buyer’s arrival date is uncertain and is his private information, but not (as stated by Result 1) when the seller knows the arrival date in

²⁵The finding that sales become uniformly less frequent with time is common to all the numerical examples I have considered. I have not yet been able to verify that this result always holds.

advance.

A significant component of the intuition can be understood by recalling Stokey’s (1979) intuition about the optimality of intertemporal price discrimination. The following example illustrates.

Example 2 *Suppose that $\alpha_L = 0$, but that $\gamma \in (0, 1)$, and suppose as in Subsection 3.1 that the buyer is known to arrive at date zero. Again, other than the empty policy, one can restrict attention to sales policies that contain only one point, say z . The corresponding optimal prices are such that $p_{\{z\}}^*(0) = \theta_H - e^{-(r+\alpha_H)z}(\theta_H - \theta_L)$ and $p_{\{z\}}^*(z) = \theta_L$; prices at other dates are irrelevant provided they are sufficiently high. The seller’s expected profit is equal to*

$$\gamma\theta_H + e^{-rz}((1 - \gamma)\theta_L - \gamma e^{-\alpha_H z}(\theta_H - \theta_L)).$$

If α_H is set equal to zero, this is linear in the discount factor e^{-rz} , and so one recovers the result for constant values mentioned in Subsection 3.1: the optimal sales policy is either empty or $\{0\}$ and so any purchases are immediate. However, if $\alpha_H > 0$ and γ is sufficiently large, then a sales policy $\{z\}$ is optimal for $z > 0$.

The reason for the optimality of $z > 0$ in Example 2 is that, if the buyer’s value is high, he expects it to decline. This makes the strategy of purchasing at date $z > 0$ unattractive. In effect, the buyer is “impatient” to obtain the good if his value is high, and so the seller can charge a higher price at date zero. Stokey (1979) identifies the same effect in a model where values change deterministically (here, it is instead *expected* changes in the value that make the buyer effectively impatient when his value is high).

The above effect is clearly not the only force at play. If one extends Example 2 to allow for an uncertain arrival time, the optimal price may be constant at $\theta_L = \frac{\omega_L}{r}$ with the corresponding sales policy $A = \mathbb{R}_+$ (this is true whenever the parameters are such that the optimal sales policy in the example is $\{0\}$). The reason is that, at any date t in the sales policy, the choice of subsequent dates in the policy has no effect on rents before date t . For the subsequent dates to affect the buyer’s rents when arriving before t requires that he anticipate the possibility of purchasing with a high value on such dates when his value is low, i.e. that $\alpha_L > 0$.

Instead, Proposition 1 assumes that both α_L and α_H are positive and states that, provided the buyer purchases with a low value with positive probability (i.e., provided (4) holds), the sales policy *always* includes isolated points, possibly after an initial interval $[0, T]$.²⁶ The reason relates to the weight that must be placed

²⁶As mentioned in footnote 23, this conclusion holds provided $\alpha_L, \alpha_H > 0$ regardless of the value of γ .

on limiting the rent the buyer earns in case of arrival at earlier dates. At dates in $[0, T]$, this weight is low enough that the efficiency advantage of inducing immediate purchase with a low value dominates. At dates above $\max A^*$, the weight is high enough that it dominates the efficiency advantage of inducing any purchase with a low value. Between T and $\max A^*$, the seller balances the trade-offs by inducing purchase with a low value only occasionally, i.e. by “holding sales”. This allows the seller to provide the good to the buyer if he is waiting with a low value, whilst (by exploiting his “impatience”) still charging a relatively high price if he arrives with a high value (or if his value switches to high) at times in between.²⁷ This of course also lowers the option value of delaying purchase if the buyer arrives at an earlier time, so it reduces the rents the buyer earns at earlier arrival dates as well.

Another way to understand the discreteness of the sales policy after T is by considering mechanisms that involve both rental and outright sale. In case (4) holds, it follows easily from Result 1 that the sales policy $A = [0, \bar{t}]$ achieves a higher expected profit than $A = \emptyset$, provided that \bar{t} is sufficiently small. The question of why a simple interval $[0, \bar{t}]$, for optimally chosen \bar{t} (see Lemma 9 in the Appendix), does not maximize profit can then be understood from the fact that it raises the same profit as the following mechanism which involves rental. First, the good is rented at rate ω_L from date zero up to \bar{t} . Second, the mechanism reverts to the original mechanism at \bar{t} , i.e. the buyer is given the opportunity to purchase outright at \bar{t} at a price θ_L before the price changes to θ_H for outright purchase for the rest of time. That the two mechanisms are equivalent is immediate from (8). Suppose that $\omega_L \leq \gamma^S \omega_H$. Then, if the seller must rent the good rather than sell it between date zero and \bar{t} , a better policy would be to set a rental price ω_H , inducing the buyer to rent only if his value is high (again reverting to the original mechanism at \bar{t}). This mechanism is, however, strictly dominated by using the posted price mechanism with outright sale with a sales policy equal to $\{\bar{t}\}$ and optimally chosen prices. (The reason that selling dominates rental in this case is simply that outright selling is more efficient, as the buyer continues to hold the good even if his value turns low, while his expected rents, assuming optimal prices, are the same.) The same intuition can be extended to the case where $\omega_L > \gamma^S \omega_H$ for dates later than T by recognizing that a commitment to rent at a price of ω_L over an interval of positive length (rather than, say, to sell at a price which induces the buyer to purchase only if his value is high) implies that earlier prices must be lower.

²⁷The effect of “impatience” is crucial. In case $\alpha_H = 0$, price fluctuations are never optimal for any level of γ . If γ is sufficiently small, the optimal sales policy is an interval of the form $[0, \bar{t}]$ with $\bar{t} > 0$, which means that prices gradually rise up to \bar{t} and then jump to θ_H .

Recursive formulation. I now discuss the recursive formulation that is used to prove Proposition 1 and which allows optimal sales policies to be derived numerically using value function iteration. The idea is to consider a sequence of sub-problems such that, for each date t that is included in the sales policy, the problem is to choose the subsequent date in the sales policy. Each sub-problem accounts for the effect of the future selection of dates not only on the revenue the seller expects after date t , but also on the rents earned by the buyer when arriving before t . The value of each sub-problem $W^*(t)$ satisfies the functional equation

$$W^*(t) = \sup_{t' > t} \left\{ \begin{array}{l} \phi(t' - t) - (e^{\lambda t} - 1) \Pr(\tilde{\theta}_{t'} = \theta_H | \tilde{\theta}_t = \theta_L) e^{-r(t' - t)} (\theta_H - \theta_L) \\ + e^{-(r+\lambda)(t' - t)} W^*(t') \end{array} \right\} \quad (11)$$

on \mathbb{R}_+ , where

$$\phi(z) = \int_0^z \lambda e^{-(\lambda+r)s} R(z-s) ds \quad (12)$$

and where $R(\cdot)$ is given by (2).

As explained in the Appendix, the functional equation is helpful for identifying the points in an optimal sales policy following date T . If A^* is a solution, then for dates $s'' > s' > T$ to be consecutive members of A^* it must be that s'' attains the value $W^*(s')$ for the sub-problem at s' .²⁸ Note that the inclusion of no further dates in the sales policy can be approximated by taking t' arbitrarily large in (11) (the value from including no future dates is $\frac{\lambda}{r+\lambda} \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{\alpha_L + r} \right) \theta_H$).

Now consider the optimization problem. The first term, $\phi(t' - t)$, is the revenue the seller expects to earn over the interval $[0, t' - t]$ for the hypothetical sales policy which includes only the point $\{t' - t\}$. The third term, $e^{-(r+\lambda)(t' - t)} W^*(t')$, is the value of the seller's problem of choosing additional dates in the policy after t' , discounted to reflect the rate of time preference as well as the reduced probability that the buyer arrives after date t' .

Finally, consider the second term. This term accounts for the effect of including date t' in the sales policy on the buyer's rent when arriving before date t ; see (9). It is weighted by $e^{\lambda t} - 1$, which, after taking into account the fact that the value of the seller's problem is discounted also by the probability of arrival subsequent to t (i.e., $e^{-\lambda t}$) corresponds to the probability of an earlier arrival (i.e., $1 - e^{-\lambda t}$). It is this term which captures the non-stationarity of the trade-offs described above, with the consequence that the

²⁸That the functional equation is valid on $[0, T]$ reflects the continuity of the problem: the profit obtained from any sales policy which includes an interval can be approximated arbitrarily closely by a discrete sales policy. This is immediate from considering (23) in the Appendix.

distance between sales increases over time and that eventually there are no further sales.

3.4 Diminishing value for the good

As mentioned in the Introduction, the finding that prices tend to rise over time is not borne out in many markets where the evolution of fashion and technology play an important role. One possible explanation is that, as fashions and technology are replaced, consumers derive less and less value from the older products. As this occurs, it seems reasonable to expect the price also to decline. The model can be extended to accommodate this possibility, and it is easy to see that the key tradeoffs in determining when and whether to reduce prices remain.

One particularly tractable possibility is to posit that the flow value derived from the good vanishes with the time since date zero at a constant rate $f > 0$. In this case, $\omega_H(t) = \omega_H(0)e^{-ft}$ and $\omega_L(t) = \omega_L(0)e^{-ft}$ for all $t \geq 0$, where $\omega_H(0) > \omega_L(0) > 0$. This implies that the expected values from obtaining the good in the high and low states are $\theta_H(t) = e^{-ft}\theta_H(0)$ and $\theta_L(t) = e^{-ft}\theta_L(0)$, where

$$\begin{aligned}\theta_H(0) &= \frac{(\alpha_L + r + f)\omega_H(0) + \alpha_H\omega_L(0)}{(r + f)(\alpha_L + \alpha_H + r + f)} \\ \theta_L(0) &= \frac{\alpha_L\omega_H(0) + (\alpha_H + r + f)\omega_L(0)}{(r + f)(\alpha_L + \alpha_H + r + f)}.\end{aligned}$$

Using the same arguments as in the previous subsection, it is straightforward to check that the optimal sales policy is identical for the problem where the discount rate is given by $r + f$. For any sales policy A , conditionally optimal prices are given by

$$p_A^*(t) = e^{-ft} \left(\theta_L(0) - \int_t^\infty e^{-(r+f)(s-t)} \alpha_L e^{-(r+f+\alpha_L+\alpha_H)(m_A(s)-s)} (\theta_H(0) - \theta_L(0)) ds \right)$$

if $t \in A$ and by

$$p_A^*(t) = e^{-ft} \left(\begin{array}{l} \theta_H(0) - \int_t^\infty e^{-(r+f)(s-t)} \alpha_L e^{-(r+f+\alpha_L+\alpha_H)(m_A(s)-s)} (\theta_H(0) - \theta_L(0)) ds \\ - e^{-(r+f+\alpha_L+\alpha_H)(m_A(t)-t)} (\theta_H(0) - \theta_L(0)) \end{array} \right)$$

if $t \notin A$. The following example illustrates one possibility for the pattern of prices.

Example 3 Suppose that $\omega_L = 1\frac{1}{4}$, $\omega_H = 2$, $\alpha_L = \alpha_H = \frac{1}{4}$, $\lambda = \frac{1}{10}$, $r = \frac{4}{50}$ and $f = \frac{1}{50}$ (hence $r + f = \frac{1}{10}$). It follows that $\theta_L(0) = 15\frac{5}{8}$ and $\theta_H(0) = 16\frac{7}{8}$. The optimal sales policy is the same as in Example 1. The optimal sales policy and price path is depicted in Figure 2.

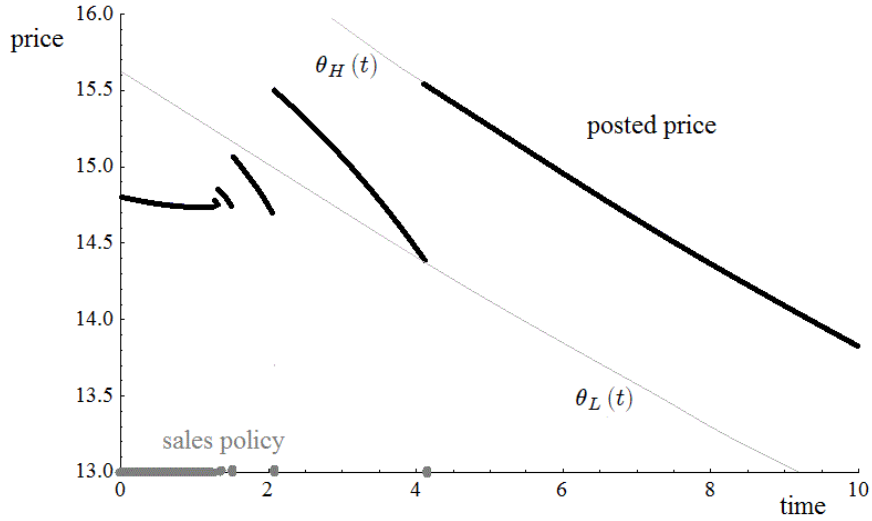


Figure 2: Optimal prices and sales policy for Example 3

In Example 3, after sufficient time, the buyer purchases only once his value turns high. At such a date t , the price is equal to the expected value from acquiring the good at date t , $\theta_H(t)$, which is declining. Thus, by combining the hypothesis of private uncertainty about future values with the hypothesis of an anticipated and predictable decline in values, one can account for discounting both early and late in the product life cycle, which, as mentioned in the Introduction, is consistent with price patterns found in fashion retail (see Pashigian and Bowen (1991)).²⁹

3.5 Buyer attrition from the market

Another important feature of the model considered so far is that the buyer anticipates having the opportunity to purchase the good into the infinite future. As discussed in Subsections 3.2-3.3, including additional dates in the sales policy after some date t in the policy therefore increases the rent expected at all dates before t by the same amount in date-zero terms. This is what drives the optimality of a bounded sales policy.

The above conclusion does not remain true, however, if the buyer expects to exit the market. Suppose that the buyer, having arrived to the market, exits it permanently at some rate $d > \lambda$. Then the ex-ante probability the buyer is in the market at any date t is $e^{-\lambda t} \frac{\lambda}{d-\lambda} (1 - e^{-(d-\lambda)t})$, which eventually becomes proportional to the ex-ante arrival rate $\lambda e^{-\lambda t}$.

²⁹In particular, Pashigian and Bowen document discounting of shirts both at the beginning and end of the relevant season, noting (p 1030) "The percentage of shirts sold on sale at the beginning of each season is unexpectedly and surprisingly high."

To understand the implications for the optimal price path, consider how one must modify the recursive formulation discussed in Subsection 3.3. For simplicity, suppose that the buyer continues to use the good after exiting the market, so his value from obtaining it is either θ_L or θ_H , as defined in the model set-up. The timing is such that the buyer simply learns that he can no longer purchase the good and cannot subsequently purchase (i.e., there is no “last chance” to purchase the good before exiting). If the buyer instead stops having any need for the good, then the values from obtaining the good when he is still in the market should be appropriately adjusted, but the rest of the analysis remains the same.

Consider the expected profit given arrival at date zero when the sales policy is $\{z\}$ for $z \geq 0$. Abusing notation, instead of the profit defined in (2), this is

$$R(z; d) = \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{\alpha_L + r + d} \right) \theta_H + (1 - \gamma^S) e^{-(\alpha_L + r + d)z} \left(\theta_L - \frac{\alpha_L \theta_H}{\alpha_L + r + d} \right) - \gamma^S e^{-(r+d)z} (\theta_H - \theta_L).$$

The value of the seller’s problem of choosing the next date in the sales policy is then given by

$$W^*(t) = \sup_{t' > t} \left\{ \begin{array}{l} \phi(t' - t; d) - \frac{\lambda}{d - \lambda} (1 - e^{-(d - \lambda)t}) \Pr(\tilde{\theta}_{t'} = \theta_H | \tilde{\theta}_t = \theta_L) e^{-(r+d)(t' - t)} (\theta_H - \theta_L) \\ + e^{-(r+\lambda)(t' - t)} W^*(t') \end{array} \right\} \quad (13)$$

where, again abusing notation,

$$\phi(z; d) = \int_0^z \lambda e^{-(\lambda+r)s} R(z - s; d) ds.$$

Assuming $d > \lambda$, the seller’s continuation problem of choosing the next date in the sales policy becomes stationary over time, with the value function converging to

$$W_{LR}^* = \sup_{z > 0} \left\{ \frac{\phi(z; d) - \frac{\lambda}{d - \lambda} \Pr(\tilde{\theta}_z = \theta_H | \tilde{\theta}_0 = \theta_L) e^{-(r+d)z} (\theta_H - \theta_L)}{1 - e^{-(r+\lambda)z}} \right\}. \quad (14)$$

To understand the long-run pattern of prices, one can focus on an auxiliary problem with a value function satisfying (14). For instance, as stated in the following result, if a discrete sales policy is uniquely optimal in the auxiliary problem, then any optimal policy in the problem of interest must become discrete after sufficient time, with the distance between sales dates the same as in the solution to the auxiliary problem.

Result 2 Suppose that W_{LR}^* in (14) is uniquely attained by $z^* > 0$, and that this exceeds the value attained at $z = +\infty$.³⁰ Suppose A^{opt} is an optimal sales policy in the problem of interest (see (13)). Then

- (i) A^{opt} is unbounded, and
- (ii) for any $\varepsilon > 0$, there exists a date $\hat{t} > 0$ such that (a) $A^{opt} \setminus [0, \hat{t}]$ is a discrete set, and (b) if s', s'' are consecutive dates in $A^{opt} \setminus [0, \hat{t}]$, then $||s'' - s'| - z^*| < \varepsilon$.

Using Result 2, one can easily find parameter values such that any optimal sales policy is discrete after sufficient time (to give an example, this is true if $d = 1/5$ and the other parameters are as given in Example 1; in this case the distance between dates in the sales policy converges to $z^* = 2.1$ with time). Thus, one possibility is that price fluctuations persist and become stationary in the long run.

3.6 Countercyclical markups

The theory developed above provides a possible explanation for countercyclical markups that seems to have gone unnoticed. The hypothesis of countercyclical markups was proposed originally by Pigou (1927) and Keynes (1939), who observed that prices tend to change little in response to increases in aggregate demand. Murphy et al. (1989) report countercyclical movements in the prices of finished goods relative to intermediate inputs in the post-war period, with the finding most pronounced for consumer durables. At the micro level, Pashigian and Bowen (1991) document the evolution of prices for men's shirts and Warner and Barsky (1995) document the prices of a wider range of durables and "semi-durables". They find evidence that prices tend to be lower during times that shoppers are more active, such as before Christmas and on weekends. Chevalier et al. (2003) find lower markups during high demand periods in supermarket data.

One way to model variation in the level of buyer activity is to allow variation in the arrival rate of buyers to the market (see the discussion of Bils' (1989) theory below). One might expect that dates with a faster arrival rate would be associated with lower prices because of a simple trade-off made by sellers when deciding their price paths. At any date, a seller trades off between two objectives: (i) extracting as much profit as possible from buyers arriving at (or just before) a given date, which may require setting a low price, and (ii) lowering the option value of waiting for earlier arrivers, which requires setting a high price at the date in question. The first objective becomes much more important if there is a rush of buyers to the market at (or just before) this date. Thus, one might expect to see lower prices at dates when the rate of arrival

³⁰That is, $W_{LR}^* > \frac{\lambda}{r+\lambda} \left(\gamma^S + (1 - \gamma^S) \frac{\alpha_L}{\alpha_L + r + d} \right) \theta_H$. In other words, the value from including an optimal choice of future dates in the auxiliary problem exceeds that from not including any future dates.

is high. However, this intuition is incorrect in a model where values for the good do not change because optimal prices are constant and set at a level that maximizes the expected profit for each possible arrival date (see the discussion of Stokey’s (1979) result in Subsection 3.1). It can be correct, though, if values for the good change stochastically.

One way to formalize the intuition is to allow a mass point in the distribution of the buyer’s arrival date. The following result is an immediate consequence of the previous analysis.

Result 3 *Suppose that (4) holds. Suppose that the buyer arrives at date τ_{mass} with probability ρ , and that with complementary probability he arrives according to the exponential distribution with parameter $\lambda > 0$ (all other assumptions are as set out in Section 2). For any $\varepsilon > 0$, there exists $\bar{\rho} \in [0, 1)$ such that, for all $\rho \in (\bar{\rho}, 1]$, any optimal sales policy includes a date $t \in [\tau_{mass}, \tau_{mass} + \varepsilon)$.*

The intuition for this result is indeed the one given above. If the probability the buyer arrives at a given mass point is taken close enough to one, then the objective of maximizing profit for the event of arrival at that date dominates. If maximizing profit conditional on arrival at that date requires inducing immediate purchase irrespective of the buyer’s value (i.e., if (4) holds), then it is optimal to induce purchase irrespective of value either at that date, or just after it, also when arrival is uncertain. If other nearby dates are not included in the sales policy, then prices must be especially low at the date of the mass point. It seems reasonable to expect that a similar implication would hold more generally also with less extreme variation in the arrival rate, at least often enough to show up when averaged across different products with different initial dates and evolution of buyer values.

There are a couple of other points to note. First, one might object that sellers often cannot anticipate spikes in buyer activity in advance (this seems more the case with business cycles than with the Christmas and weekend effects mentioned above). Full-commitment prices might, however, be conditioned on some publicly observable variable that is correlated with the rate of arrival. Secondly, countercyclical markups are also observed for goods usually considered non-durables. However, the theory does seem likely to extend to goods that buyers perceive to be substitutable across time; in particular, to goods with the property that purchase at one moment reduces the value obtained from purchasing again at a later date.

A number of theories for countercyclical markups have already been proposed in the literature. Bils’ (1989) theory has a similar flavor to mine in that it is based on the optimal response of the seller to changes in the rate of arrival. He suggests that, for experience goods, firms would like to lower prices in order to attract

new customers but want to raise them to extract rent from customers who have previously purchased and know that the product is effective or to their liking. Periods associated with the greatest influx of new buyers are ones where the trade-off between extracting rent from previous purchasers and enticing new customers to experiment is resolved in favor of the latter. Assuming the seller *cannot* commit, Bilal shows that prices are lowest during these periods. One difficulty with this theory is that it requires experimentation to play a key role in the purchase decision, which seems unlikely to be true for all goods (e.g., where the consumer only makes a one-off purchase, or where relevant information is available without the need to experiment).

Chevalier et al. (2003) review other theories (especially, Rotemberg and Saloner (1986) and Warner and Barsky (1995)) and suggest that they do not match their data set well. They propose an explanation based on loss-leadership in advertising. Because advertising is costly and consumer attention scarce, it is optimal to lower the prices of the most popular products. They therefore suggest that, when the consumption rate of a good is high for exogenous reasons (e.g., the consumption of tuna during Lent), the good will tend to be sold at a discount. Again, however, the applicability of this theory across a range of markets is an open question.

What drives countercyclical markups therefore seems unresolved in the literature. The contribution of this subsection is simply to show that countercyclical pricing patterns may be consistent with full-commitment pricing when buyer values change stochastically. The intuition driving this result is necessarily absent from models where the seller does not commit. A seller lacking commitment ability does not take into account the effect of prices at each date on the profit earned in case of earlier arrival. It is variation in the weight assigned to the the profit earned in case of earlier arrival which drives countercyclical markups in my theory.

4 The fully-optimal mechanism

I now consider the unrestricted optimal mechanism. As discussed in the Introduction, the seller will want to incentivize the buyer to participate in the mechanism immediately. At the same time, since the buyer is forward-looking and strategic, the time at which he wants to participate potentially depends on the seller's choice of mechanism over the infinite future. An apparent difficulty is therefore understanding how the seller can choose the mechanism to induce immediate participation without giving up excessive rent. Moreover, since the level of rent will depend on the equilibrium allocation of the good, it will be important to understand

the relationship between this allocation and the size of the information rents that the buyer receives.

Without loss of optimality, the seller fully commits to a deterministic mechanism at date zero.³¹ As described above, the buyer cannot participate in the mechanism until he arrives. However, upon arrival, he perfectly observes the mechanism and can engage in unrestricted communication with the seller.³²

Because the seller can fully commit, the revelation principle applies in this environment, and so it is without loss of optimality to restrict attention to incentive-compatible direct mechanisms \mathcal{M}_D in which the buyer participates in the mechanism at the time of arrival and then reports how his value evolves over time (by the revelation principle, the seller cannot do better than inducing immediate participation; it turns out that a policy of inducing immediate participation is strictly preferred). Direct mechanisms are defined formally in the Appendix, where I also note standard measurability restrictions.

What is important to understand for the purposes of my analysis is that, for any mechanism, there exists a measurable function $m : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that, in equilibrium, if the buyer arrives at date τ , he receives the good at $m(\tau) \geq \tau$ if his value is low upon arrival and remains low. I say that such a mechanism is “consistent with m ”.

The approach is to initially consider a “relaxed program”. For any function m , I find a mechanism that maximizes expected profit among mechanisms consistent with m taking account only of certain incentive constraints (and thus assuming that the buyer always participates and reports truthfully, provided these constraints are satisfied). The constraints of the relaxed program are:

1. If the buyer has a high value but has reported only to have a low value in the past, then he prefers to report truthfully from then on rather than continuing to report a low value forever.
2. Suppose that the buyer follows a truthful reporting strategy after participating. If the buyer’s initial value upon arrival is low, then he prefers to participate immediately at this time rather than wait and participate if and only if his value becomes high.

Having derived the mechanism that maximizes profit subject to consistency with m in the relaxed program, I find the optimal choice of m for this program, and then verify that, for this choice of m , all incentive

³¹As for posted prices, it is optimal to consider deterministic mechanisms due to the “bang-bang” nature of the solution to the profit-maximization problem — see (21) in the Appendix.

³²Some care is needed in extending the analysis to a population of agents. The novel concern with respect to the single-buyer environment is the possibility the buyer tries to contract more than once. In effect, the buyer may find it advantageous to claim to have two or more identities. In case the seller cannot identify buyers who have contracted previously, the implementation of the optimal mechanism may need to be chosen carefully to discourage the buyer obtaining multiple contracts.

constraints are indeed satisfied. In deriving the solution to the relaxed program, the following result, which follows from arguments in Battaglini (2005), is helpful.

Lemma 3 *Let \mathcal{M}_D be any mechanism satisfying Constraints 1 and 2. There exists an alternative mechanism \mathcal{M}'_D such that (a) the buyer is always indifferent between the alternative strategies in Constraint 1,³³ (b) Constraint 2 is satisfied, and (c) the profit for the seller is at least as high as for the original mechanism \mathcal{M}_D .*

For the purposes of solving the relaxed program, Lemma 3 allows me to restrict attention to mechanisms satisfying Part (a) of the lemma. For each date τ and each $\theta_\tau \in \{\theta_L, \theta_H\}$, let $\nu_\tau^{\mathcal{M}_D}(\theta_\tau; \emptyset)$ be the buyer's expected payoff from participating for the first time at date τ in the mechanism \mathcal{M}_D and then following an optimal reporting strategy (the reason for conditioning $\nu_\tau^{\mathcal{M}_D}$ on \emptyset is that the buyer has not yet reported any values to the seller at the time he first participates).³⁴ For any function m as defined above, the relevant mechanisms are the ones such that, assuming truth-telling is optimal for the buyer, at any date τ ,

$$\begin{aligned} & \nu_\tau^{\mathcal{M}_D}(\theta_H; \emptyset) - \nu_\tau^{\mathcal{M}_D}(\theta_L; \emptyset) \\ &= e^{-r(m(\tau)-\tau)}(\theta_H - \theta_L) \left(\Pr(\tilde{\theta}_{m(\tau)} = \theta_H | \tilde{\theta}_\tau = \theta_H) - \Pr(\tilde{\theta}_{m(\tau)} = \theta_H | \tilde{\theta}_\tau = \theta_L) \right) \\ &= e^{-(r+\alpha_L+\alpha_H)(m(\tau)-\tau)}(\theta_H - \theta_L). \end{aligned} \tag{15}$$

That is, the difference in the buyer's expected payoff conditional on high and low values at the participation date τ is simply the discounted expected difference in the values at date $m(\tau)$.

Known arrival time. The analysis when the seller knows the buyer's arrival time is essentially the same as in Battaglini (2005), and it turns out one can characterize the optimal mechanism directly by appropriately adjusting the “dynamic virtual surpluses” for his problem to the present one.³⁵ To provide a benchmark for my results when the buyer's arrival time is uncertain and his private information, I briefly show how the above relaxed program applies to the case where the buyer's arrival time is known.

Using the same normalization as in Subsection 3.1, suppose the buyer arrives at date zero. Assuming

³³That is, if the buyer has a high value but has reported only low values, he is *indifferent* between reporting truthfully from then on and instead continuing to report a low value forever.

³⁴In terms of the notation in the Appendix, $\nu_t^{\mathcal{M}_D}(\theta_t; h^{[\tau, t]})$ denotes the value of the buyer's problem at date $t > \tau$ when his value is θ_t and when he has participated at date τ and reported his values since that time to be $h^{[\tau, t]}$.

³⁵In particular, one can solve the problem by focussing on the flow values, and then noticing that the optimal mechanism has the property that, once the buyer receives the good, he holds it forever. Thus, whilst in richer environments there is a difficulty related to the non-time separability of the optimization problem (see Pavan, Segal and Toikka (2011) and Garrett (2011)), this difficulty is not present when the buyer has only two values.

truth-telling is optimal conditional on participation, Constraint 2 is then simply $\nu_0^{\mathcal{M}^D}(\theta_L; \emptyset) \geq 0$, since the seller optimally commits not to allow the buyer another chance to obtain the good if he does not participate at date zero. (Alternatively, the seller finds it optimal to induce purchase only when the buyer's value is high, i.e. to set $m(\tau) = +\infty$ for all $\tau > 0$.) The indirect mechanism in the following result ensures that this constraint holds with equality.

Result 4 *Consider the environment of the model set-up (with any value $\gamma \in (0, 1)$), but where the seller knows that the buyer's arrival date is equal to zero. The following mechanism is optimal conditional on the buyer receiving the good at $m(0)$ if his value remains low. The buyer obtains at date zero an option to purchase the good at prices*

$$p_0^\#(t; m(0)) = \theta_H - \Pr\left(\tilde{\theta}_{m(0)} = \theta_H | \tilde{\theta}_t = \theta_H\right) e^{-r(m(0)-t)} (\theta_H - \theta_L)$$

for dates $t \in [0, m(0)]$, and for this option he pays at date zero the fee

$$f_0^\#(m(0)) = \Pr\left(\tilde{\theta}_{m(0)} = \theta_H | \tilde{\theta}_0 = \theta_L\right) e^{-rm(0)} (\theta_H - \theta_L).$$

The buyer purchases the good at the first moment his value is high, or at $m(0)$ if he has not already purchased.

The optimal choice of $m(0)$ is

$$m^\#(0) = \frac{1}{\alpha_H} \ln \left(\frac{\gamma(\omega_H - \omega_L)}{(1 - \gamma)\omega_L} \right).$$

The price path $p_0^\#$ coincides with (1) derived in Section 3.1. It is chosen to keep the buyer indifferent between purchasing the good and waiting to purchase at date $m(0)$ (the expiration date of the option) when his value is high. It therefore achieves the lower bound $e^{-(r+\alpha_L+\alpha_H)m(0)}$ on the difference in expected payoffs from participation given high and low initial values (as given by (15)). The fee $f_0^\#(m(0))$ is set equal to the rent the buyer expects conditional on a low initial value from having the option to purchase at the prices $p_0^\#$ (see (3)). These two observations imply that the buyer's expected rents are as small as possible for a mechanism consistent with $m(0)$. Moreover, the buyer receives the good as soon as his value is high at dates before $m(0)$, so the mechanism is as efficient as possible. Hence, the proposed mechanism maximizes expected profit subject to Constraints 1 and 2.

To see that all of the incentive constraints are satisfied, not only Constraints 1 and 2, note that after the buyer has purchased the option, he simply faces the optimal price path for a sales policy $A = \{m(0)\}$

as studied in the previous section. That the buyer optimally purchases at the times indicated given these prices is shown in the proof of Lemma 2. That he finds it optimal to participate, i.e. to purchase the option, simply follows because he expects a non-negative payoff from doing so.

Note also that the optimal mechanism always induces purchase eventually if the buyer's value remains low, i.e. $m^\#(0) < +\infty$. That the buyer is eventually induced to take the efficient action irrespective of his realized values reflects the “principle of vanishing distortions” demonstrated by Battaglini (2005) for a model with two values (“types”), and earlier by Besanko (1985) for a model with a continuum of types. The reason for the result is that the rents that must be left to the buyer are small provided $m(0)$ is sufficiently large, because the buyer's initial information is a poor predictor of his value at dates far in the future.

This result can be contrasted with Result 1, where it is seen that, if the reverse of (4) holds, the buyer never receives the good in case his value remains low forever (the price is simply constant at θ_H). The reason for the different result is that, when the seller is restricted to using posted prices, she has no way to recover rent that the buyer expects to earn from the option to purchase at given prices. It therefore becomes more important to limit buyer rents by the choice of the price path itself.

Arrival time uncertain and private information. Now consider the setting where the buyer's arrival time is uncertain and is his private information. The novel difficulty in considering this environment relates to the endogenous participation constraint. Because arrival is uncertain, one objective of the seller is to raise as much profit as possible contingent on each arrival date (as opposed to the case above where she cares only about the profit obtained when the buyer arrives at date zero). One might therefore expect that, for each arrival date, the seller may want to allocate the good to the buyer eventually if his value remains low; i.e., that $m(\tau) < +\infty$ for each arrival date τ . The buyer must then earn additional rent if his value is high at the time he participates. Contrary to the case for a known arrival time, the buyer must therefore expect to earn positive rent also if his initial value is low. The reason is that he must be discouraged from delaying participation until such time as his value becomes high.

Assuming truth-telling is optimal conditional on participating, Constraint 2 now requires, for each τ , that

$$v_\tau^{\mathcal{M}^D}(\theta_L; \emptyset) \geq \int_\tau^\infty \alpha_L e^{-(r+\alpha_L)(s-\tau)} v_s^{\mathcal{M}^D}(\theta_H; \emptyset) ds,$$

and hence that

$$\nu_{\tau}^{\mathcal{M}D}(\theta_L; \emptyset) \geq \int_{\tau}^{\infty} \alpha_L e^{-(r+\alpha_L)(s-\tau)} \left(\nu_s^{\mathcal{M}D}(\theta_L; \emptyset) + e^{-(r+\alpha_L+\alpha_H)(m(s)-s)} (\theta_H - \theta_L) \right) ds. \quad (16)$$

A lower bound for the expected rents of the buyer conditional on a low initial value is therefore given by the function that solves (16) with equality, subject to the transversality condition that buyer rents remain bounded. This is given by

$$\nu_{\tau}^{LB}(\theta_L; \emptyset) = \int_{\tau}^{\infty} e^{-r(s-\tau)} \alpha_L e^{-(r+\alpha_L+\alpha_H)(m(s)-s)} (\theta_H - \theta_L) ds. \quad (17)$$

Let

$$\nu_{\tau}^{LB}(\theta_H; \emptyset) = \nu_{\tau}^{LB}(\theta_L; \emptyset) + e^{-(r+\alpha_L+\alpha_H)(m(\tau)-\tau)} (\theta_H - \theta_L). \quad (18)$$

Assuming immediate participation and truth-telling is optimal, the following indirect mechanism ensures the buyer expects rents equal to $\nu_{\tau}^{LB}(\theta_L; \emptyset)$ and $\nu_{\tau}^{LB}(\theta_H; \emptyset)$ for low and high initial values at date τ . At each date τ , the seller offers an option to buy the good at prices

$$p_{\tau}^{\#}(t; m(\tau)) = \theta_H - \Pr\left(\tilde{\theta}_{m(\tau)} = \theta_H | \tilde{\theta}_t = \theta_H\right) e^{-r(m(\tau)-t)} (\theta_H - \theta_L) \quad (19)$$

at dates $t \in [\tau, m(\tau)]$. For this option, the seller charges the date- τ fee

$$f_{\tau}^{\#}(m(\tau)) = \Pr\left(\tilde{\theta}_{m(\tau)} = \theta_H | \tilde{\theta}_{\tau} = \theta_L\right) e^{-r(m(\tau)-\tau)} (\theta_H - \theta_L) - \nu_{\tau}^{LB}(\theta_L; \emptyset). \quad (20)$$

The buyer is asked to purchase the option immediately upon arrival and then purchase the good immediately whenever his value is high, or on date $m(\tau)$ if he has not purchased by that date. This mechanism is the solution to the relaxed program conditional on the function m : it maximizes expected surplus among mechanisms consistent with m and minimizes expected rent among mechanisms consistent with m and satisfying Constraints 1 and 2.

Next, for a given function m , one can calculate the expected profit assuming the buyer finds the mechanism incentive compatible. This is given by

$$\int_0^{\infty} \lambda e^{-(r+\lambda)\tau} b(\tau, m(\tau)) d\tau, \quad (21)$$

where

$$\begin{aligned}
b(\tau, m(\tau)) &= \gamma\theta_H + (1 - \gamma) \left(\frac{\alpha_L}{r + \alpha_L} \theta_H + e^{-(r+\alpha_L)(m(\tau)-\tau)} \left(\theta_L - \frac{\alpha_L}{r + \alpha_L} \theta_H \right) \right) \\
&\quad - \frac{1 - e^{-\lambda\tau}}{\lambda e^{-\lambda\tau}} \alpha_L e^{-(r+\alpha_L+\alpha_H)(m(\tau)-\tau)} (\theta_H - \theta_L) - \gamma e^{-(r+\alpha_L+\alpha_H)(m(\tau)-\tau)} (\theta_H - \theta_L).
\end{aligned}$$

The function b is a “dynamic virtual surplus” for an environment with *both* evolving valuations and dynamic arrival. The first two terms are the expected surplus given arrival at date τ when the buyer purchases at date $m(\tau)$ if his value remains low. The fourth term reflects the additional rent earned in case the buyer’s value is high when he arrives at date τ . The third term is the most interesting one and is absent in case the seller knows the buyer’s arrival date. For each date τ , it captures the effect of the choice of $m(\tau)$ on the rents of the buyer in case of arrival before date τ . The ratio of the probability of previous arrival to the ex-ante arrival rate $\frac{1 - e^{-\lambda\tau}}{\lambda e^{-\lambda\tau}}$ ensures an appropriate balance between reducing rent in case of earlier arrival and maximizing profit conditional on arrival at date τ .

Maximizing (21) simply requires maximizing the dynamic virtual surplus at each arrival date τ . Having identified the optimal function m , I verify that all incentive constraints are satisfied, yielding the following result.

Proposition 2 *Consider the environment of the model set-up (with any value $\gamma \in (0, 1)$), where the buyer’s arrival is determined according to the exponential distribution defined there. For all $\tau \geq 0$, let*

$$m^\#(\tau) \equiv \max \left\{ \tau + \frac{1}{\alpha_H} \ln \left(\frac{(\omega_H - \omega_L) (\alpha_L (e^{\lambda\tau} - 1) + \lambda\gamma)}{\omega_L \lambda (1 - \gamma)} \right), \tau \right\}.$$

The following mechanism is incentive compatible and optimal. Upon arrival at date τ , the buyer pays the fee $f_\tau^\#(m^\#(\tau))$ (given by (20)) and in return receives the option to purchase the good at prices $p_t^\#(t; m^\#(\tau))$ (given by (19)) for each $t \in [\tau, m^\#(\tau)]$.³⁶ The buyer purchases the good the first time his value turns high or at date $m^\#(\tau)$, whichever comes first.

The reason why the mechanism in the proposition is incentive compatible is as follows. Having purchased the option, the buyer faces the same price path as considered in Result 4. He therefore finds it optimal to follow the prescribed purchase rule. That it is optimal for him to purchase an option immediately upon arrival is established using a standard verification argument.

³⁶The buyer is only ever allowed to acquire one option.

What is critical for the optimality of immediate participation in the above mechanism is that the function m is non-decreasing, i.e. that conditional on his value remaining low, the buyer receives the good (weakly) sooner if he participates earlier. This turns out to be precisely what is required to ensure that, when the buyer's initial value is high, he expects to earn more rent by participating immediately than by waiting to participate at a later date.

The method can be applied also for other distributions of arrival times, but it is important that the arrival rate of the buyer does not increase too quickly with time. For instance, if one considers a distribution of arrival times with a mass point $\tau_{mass} > 0$ (see Result 2), the function m that solves the relaxed program jumps downwards at τ_{mass} . The solution to the relaxed program is then not implementable, and some form of “ironing” is required to calculate the optimal mechanism.

Properties of the optimal mechanism. Now consider the main features of the optimal mechanism. That $m^\#(\tau) < +\infty$ for all τ reflects that the principle of vanishing distortions still applies. In addition, a key innovation with respect to Battaglini's work and the rest of the dynamic mechanism design literature is to shed light on the way optimal long-term contracts evolve with the date they are signed. Examining the formula for $m^\#$ in Proposition 2 yields the following corollary.

Corollary 2 *Consider the environment of the model set-up (with any value $\gamma \in (0, 1)$), where the buyer's arrival is determined according to the exponential distribution defined there. The length of time the buyer waits to purchase when he arrives at date τ and his value remains low, $m^\#(\tau) - \tau$, is increasing without bound.*

Corollary 2 implies that contracts signed later are less efficient. The reason for the result is that the seller optimally reduces the rent earned at later arrival dates in order to discourage delayed participation in case of early arrival (because this consideration is absent at date zero, the value of $m^\#(0)$ in Proposition 2 and Result 4 coincide).

This result seems likely to be robust also in other settings. For instance, if H is a continuously differentiable distribution of arrival times on \mathbb{R}_+ , and if h is the corresponding density, then the same arguments as for Proposition 2 imply that the result holds also if $\frac{H(t)}{h(t)}$ is increasing in t . The result would be tempered in case the buyer departs stochastically from the market as discussed in Subsection 3.5. For the same reason discussed there, provided the departure rate is sufficiently fast, the length of waiting when the buyer's value remains low increases with time, but is bounded above. The possibility of other processes for the evolution

of values for the good is discussed below.

The following example illustrates the evolution of optimal contracts. In order to represent the payments graphically, I suppose that the buyer pays the fee at the same time as he procures the good. Of course, this means that the fee must be adjusted appropriately to reflect discounting.

Example 4 (Example 1 continued) Consider the parameters of Example 1. Consider the fully-optimal mechanism, adjusted so that the fixed fee is paid at the time the buyer receives the good. Figure 3 gives the execution price and total payment for selected arrival times ($\tau = 0, 1, 2, 3$).




Figure 3: Execution prices and total payments for Example 4

An immediate observation from Figure 3 is that the fee for the option may initially be negative. That is, a subsidy may be required to induce participation in case the buyer arrives sufficiently early.³⁷ It is therefore important that the buyer can obtain only one option; i.e., he must not be allowed to obtain the subsidy multiple times, foregoing his right to purchase on each. Note also that the buyer's total payment may be less than his expected value from obtaining the good, even if his value is low at the time of acquisition. This parallels the observation for the optimal path of posted prices, where prices are often set below θ_L at sufficiently early dates.

³⁷The use of subsidies to induce timely participation has been studied by Gershkov and Moldovanu (2010) in an environment where buyers have constant values, but where buyer arrival may provide information to the seller about the likely timing of future arrivals. However, their focus is on efficient mechanisms and to the extent they consider profit maximization, they identify environments where such subsidies are not required (in particular, where it suffices to consider "winner-pays mechanisms").

Second, the length of time between the sale of the option and the expiry date when the price drops to θ_L is greater the later the option is sold. This simply reflects the result in Corollary 2. A key implication is that the additional rent available in case the buyer's initial value is high is less at later contracting dates (equivalently, the initial posted price is higher). This means that the amount of rent that must be left to the buyer to induce his participation when his value is low, as opposed to waiting for his value to become high, diminishes with time. This explains why the fixed fee is negative only initially.³⁸

Other stochastic processes. Finally, an important question is the extent to which the principles developed in this section hold when the buyer's value evolves according to other stochastic processes. Although the simple continuous-time framework studied here has the advantage of tractability, and, for example, allows me to obtain the elegant closed-form expression for the expiry dates in Proposition 2, the fully-optimal mechanism can also be characterized for a richer class of first-order Markov processes. A companion paper, Garrett (2011), achieves this in a discrete-time setting with a continuum of values for the good. It demonstrates that the approach of relaxing the participation constraints for buyers whose initial values are above the minimum is also helpful in this setting.

A solution to this relaxed program involves the buyer being indifferent between participating upon arrival and waiting to participate in the next period when his initial value is equal to the minimum. Whether the solution to the relaxed program is incentive compatible is more difficult to check. A sufficient condition is that it has the following property. Supposing the buyer always reports values truthfully, for every possible path of values, he receives the good at least as soon by participating immediately as by delaying his participation. This is in fact the same condition required in the present framework with only two values for the good. However, for the broader class of stochastic processes, whether the solution to the relaxed program has this property is sensitive to the process in question.

The companion paper argues that, as established by Corollary 2 for the present environment, contracts signed at later dates tend to be more distorted; i.e., buyers who arrive later tend to wait longer to receive the good. On the other hand, the option contracts no longer take the simple form considered here. Instead, the optimal mechanism often requires the buyer to constantly update and adjust his option to purchase as new information about his value arrives. This requires ongoing communication with the seller rather than

³⁸Note that the fixed fee does not necessarily increase monotonically with the contracting date. As the contracting date τ tends to infinity, the waiting length for a buyer with a low value, $m^\#(\tau) - \tau$, tends to infinity as well. Hence the amount of rent expected by the buyer from holding an option purchased at date τ tends to zero. As a result, while in the example the fee becomes positive after sufficient time, it eventually converges to zero.

a one-off decision to purchase.

4.1 Comparison to posted prices

I now consider the key differences between optimal posted prices and the fully-optimal mechanism. One important observation is that there is a form of revenue equivalence. For allocation rules that can be implemented by posted prices, the seller cannot increase profit by using a more sophisticated mechanism. For any sales policy A , this is immediate from noticing that the conditionally-optimal posted price mechanism $\mathcal{M}_{P,A}^*$ (see Lemma 2) is payoff equivalent to the solution to the relaxed program as given by (19) and (20) provided m is taken to equal m_A . One way to understand this is that the value to the buyer of obtaining an option to purchase that is personalized to his date of arrival lies only in the advantage it conveys to him in being able to obtain the good sooner at lower prices. In case the price paid for the good does not depend on the buyer's time of arrival, a personalized option has no value.

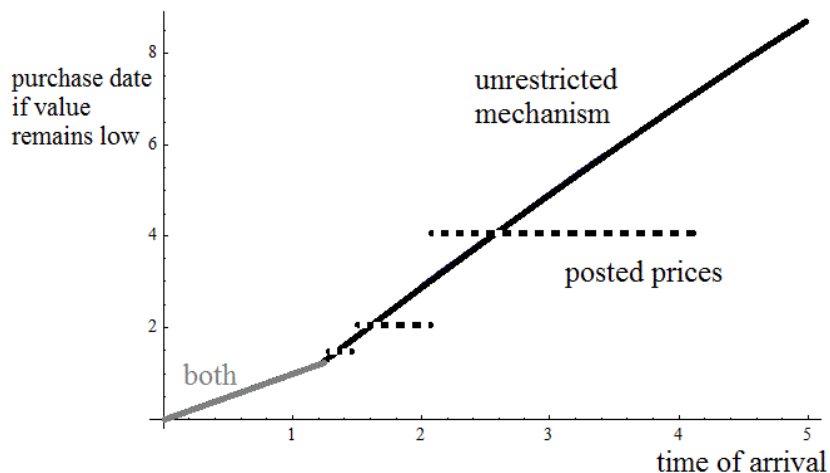


Figure 4: Time of purchase if values remain low: posted prices and fully-optimal mechanism

A key advantage of the option contracts proposed above is that they allow the seller to achieve finer intertemporal price discrimination by conditioning how long the buyer waits if his value remains low on the time of arrival. This is illustrated in Figure 4 for the posted price and unrestricted mechanisms of Examples 1 and 4 (parameter values are given in Example 1). The figure gives, for each arrival date, the date the buyer receives the good if his value remains low (for an arrival time τ , this is $m_{A^*}(\tau)$ for posted prices and $m^\#(\tau)$ for the fully-optimal mechanism). If the buyer arrives sufficiently early (in particular, before date

T as defined in Proposition 1; here, circa 1.25), purchase is immediate under both mechanisms. Later, the buyer waits to purchase on the next “sale” date under the posted price mechanism, and waits until a time particular to his date of arrival under the unrestricted mechanism. At dates sufficiently far in the future (circa date 4.1), the buyer waits indefinitely under the posted price mechanism but waits a finite length under the unrestricted mechanism.³⁹

These observations suggest a different way to understand the theory of sales with posted prices in Section 3. For the fully-optimal mechanism, the buyer receives the good eventually if his value remains low. The date he receives it depends on the date of arrival. With posted prices, the allocation of the good must be *pooled* across arrival dates: for two different arrival dates immediately preceding a given sales date, the buyer purchases at the same time if his value remains low. In this sense, when the seller is restricted to offering posted prices, fluctuations in prices reflect a lack of instruments to screen arrival times.

Waiting lists. Comparisons like the one above potentially offer a new lens through which to understand retail practices that diverge from posted prices. Although somewhat less sophisticated than the option contracts I propose, waiting lists and other delays in service appear to be used in practice to achieve similar objectives.⁴⁰ For example, retailers of luxury products sometimes require customers to join a waiting list with the possibility of obtaining the same or a similar good sooner at a higher price.⁴¹

To be more specific, consider the fashion designer Hermès, who, until recently, would place customers on a lengthy waiting list for its famous Birkin bags. Both the opportunity to join the waiting list, as well as to skip it altogether, appear to have been associated with fees implicit in the requirement of buying other Hermès items (see Tonello (2008) and Heit (2009, October 14)). This scheme, in addition to other possible advantages (e.g., creating an impression of scarcity), seems to have allowed intertemporal price discrimination to be tailored finely to the customer’s date of arrival.⁴² The implicit fee for joining the waiting list is also consistent with an attempt to extract rent expected by customers from the option to purchase. There do seem to be important differences, however, between the implicit Hermès mechanism and the fully-optimal mechanism described above. It is not clear how a customer was to be treated in case changing their mind

³⁹In the example, the posted price mechanism turns out to be less efficient than the fully-optimal mechanism. Whether this is true in general remains to be shown.

⁴⁰Related, Donaldson and Eaton (1981) suggest that one role of queuing is price discrimination based on a relationship between preferences for the good and preferences for waiting in queues. In my model, however, the buyer suffers no direct cost of waiting.

⁴¹Similarly, corrupt bureaucrats also often extract bribes in return for early service (the so-called “Myrdal effect”, Myrdal (1968)).

⁴²Another explanation for the use of waiting lists by Hermès is that the firm experienced higher than forecast demand. Whether that is a good explanation seems to depend on the extent to which Hermès could manipulate the length of its waiting lists by simply increasing its rate of production.

and wanting the bag earlier than originally forecast (it also seems there was considerable uncertainty about the waiting time for obtaining a bag).

5 Conclusion

This paper studied profit-maximizing mechanisms in an environment where buyers arrive over time and have values for the good which change stochastically. When the seller is restricted to offering posted prices, profit-maximizing prices fluctuate over time. Prices fall gradually up to dates when purchases are made by buyers whose values are low, which might be interpreted as “sales”, and jump thereafter. Prices also trend upwards over time.

When the seller can use any mechanism, she optimally sells options to purchase the good at prices that depend on the time the option is purchased as well as how long the option has been held. If a buyer’s value remains low, he executes the option after waiting a length which is greater the later he arrives. This uncovers a “principle of increasing distortions” in dynamic contracting with uncertain and unobservable arrival times which is new to the dynamic mechanism design literature.

The interest in posted price mechanisms was motivated by the fact that posted prices are the prevailing institution in many markets. Understanding the unrestricted optimal mechanism may cast some light on the cost to sellers of using this simple institution. It may also suggest why sellers sometimes deviate from this institution, for instance by requiring buyers to join waiting lists to obtain the good. The relative performance of various simple mechanisms in environments with changing values and stochastic arrival is a possible area for further research.

Finally, questions related to the ones studied here arise in a number of other principal-agent settings. Garrett and Pavan (2010) consider the problem of a firm hiring a sequence of managers to run the business. How should these contracts be designed if potential managers are not necessarily available to contract when the firm seeks a new one? Battaglini and Coate (2008), Zhang (2009), Farhi and Werning (2010) and Golosov et al. (2010) study optimal dynamic taxation when individuals’ abilities evolve stochastically and exhibit persistence. How should the tax code treat the incomes of people who decide to immigrate to the country in question? The principles and methods in this paper appear relevant for addressing these kinds of questions.

References

- [1] Baron, David and David Besanko (1984), ‘Regulation and Information in a Continuing Relationship’, *Information Economics and Policy*, 1, 447-470.
- [2] Battaglini, Marco (2005), ‘Long-Term Contracting with Markovian Consumers’, *American Economic Review*, 95, 637-658.
- [3] Battaglini, Marco and Stephen Coate (2008), ‘Pareto efficient income taxation with stochastic abilities’, *Journal of Public Economics*, 92, 844-868.
- [4] Bergemann, Dirk and Maher Said (2010), ‘Dynamic Auctions: A Survey’, Cowles Foundation Discussion Paper No. 1757R.
- [5] Bergemann, Dirk and Juuso Valimaki (2010), ‘The Dynamic Pivot Mechanism’, *Econometrica*, 78, 771-789.
- [6] Besanko, David (1985), ‘Multi-period Contracts between Principal and Agent with Adverse Selection’, *Economics Letters*, 17, 33-37.
- [7] Biehl, Andrew R. (2001), ‘Durable-Goods Monopoly with Stochastic Values’, *Rand Journal of Economics*, 32, 565-577.
- [8] Bils, Mark (1989), ‘Pricing in a Customer Market’, *Quarterly Journal of Economics*, 104, 699-718.
- [9] Bils, Mark and Peter J. Klenow (2004), ‘Some Evidence on the Importance of Sticky Prices’, *Journal of Political Economy*, 112, 947-985.
- [10] Board, Simon (2007), ‘Selling options’, *Journal of Economic Theory*, 136, 324-340.
- [11] Board, Simon (2008), ‘Durable-Goods Monopoly with Varying Demand’, *Review of Economic Studies*, 75, 391-413.
- [12] Board, Simon and Andrzej Skrzypacz (2010), ‘Optimal Dynamic Auctions for Durable Goods: Posted Prices and Fire-sales’, mimeo UCLA and Stanford University.
- [13] Chevalier, Judith A., Anil K. Kashyap and Peter E. Rossi (2003), ‘Why Don’t Prices Rise during Periods of Peak Demand? Evidence from Scanner Data’, *American Economic Review*, 93, 15-37.
- [14] Chevalier, Judith A. and Anil K. Kashyap (2010), ‘Best Prices’, mimeo, Yale SOM and Chicago Booth School of Business.
- [15] Conlisk, John (1984), ‘A Peculiar Example of Temporal Price Discrimination’, *Economics Letters*, 15, 121-126.
- [16] Conlisk, John, Eitan Gerstner and Joel Sobel (1984), ‘Cyclic Pricing by a Durable Goods Monopolist’, *Quarterly Journal of Economics*, 99, 489-505.
- [17] Courty, Pascal and Li Hao (1999), ‘Timing of Seasonal Sales’, *Journal of Business*, 72, 545-572.

- [18] Courty, Pascal and Li Hao (2000), ‘Sequential Screening’, *Review of Economic Studies*, 67, 697-717.
- [19] Das Varma, Gopal and Nikolaos Vettas (2001), ‘Optimal dynamic pricing with inventories’, *Economics Letters*, 72, 335-340.
- [20] Deb, Rahul (2010), ‘Intertemporal Price Discrimination with Stochastic Values’, mimeo Yale University.
- [21] Denardo, Eric V. (1967), ‘Contraction Mappings in the Theory Underlying Dynamic Programming’, *SIAM Review*, 9, 165-177.
- [22] Donaldson, David and B. Curtis Eaton, ‘Patience, More Than Its Own Reward: A Note on Price Discrimination’, *Canadian Journal of Economics*, 14, 93-105.
- [23] Eichenbaum, Martin, Nir Jaimovich and Sergio Rebelo (2011), ‘Reference Prices, Costs and Nominal Rigidities’, *American Economic Review*, 101, 234-262.
- [24] Ely, Jeff, Daniel Garrett and Toomas Hinnosaar (2011), ‘Overbooking’, mimeo Northwestern University.
- [25] Eso, Peter and Balazs Szentes (2007), ‘Optimal Information Disclosure in Auctions and the Handicap Auction’, *Review of Economic Studies*, 74, 705-731.
- [26] Farhi, Emmanuel and Ivan Werning (2010), ‘Insurance and Taxation over the Life Cycle’, mimeo MIT.
- [27] Gallien, Jeremie (2006), ‘Dynamic Mechanism Design for Online Commerce’, *Operations Research*, 54, 291-310.
- [28] Gallego, Guillermo and Garrett van Ryzin (1994), ‘Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons’, *Management Science*, 40, 999-1020.
- [29] Garrett, Daniel (2011), ‘Revenue Management: Evolving Valuations’, mimeo Northwestern University.
- [30] Garrett, Daniel and Alessandro Pavan (2010), ‘Managerial Turnover in a Changing World’, mimeo Northwestern University.
- [31] Garrett, Daniel and Alessandro Pavan (2011), ‘Dynamic Managerial Compensation: On the Optimality of Seniority-based Schemes’, mimeo Northwestern University.
- [32] Gershkov, Alex and Benny Moldovanu (2009), ‘Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach’, *American Economic Journal: Microeconomics*, 1, 168-198.
- [33] Gershkov, Alex and Benny Moldovanu (2010), ‘Demand Uncertainty and Dynamic Allocation with Strategic Agents’, mimeo University of Bonn.
- [34] Golosov, Mikhail, Maxim Troshkin and Aleh Tsyvinski (2010), ‘Optimal Dynamic Taxes’, mimeo Yale University.
- [35] Heit, Kimberley (2009, October 14), ‘A guide to the Hermes Birkin handbag’, *Helium*, <http://www.helium.com/items/1615933-a-guide-to-the-hermes-birkin-handbag>.

- [36] Horner, Johannes and Larry Samuelson (2011), ‘Managing Strategic Buyers’, *Journal of Political Economy*, 119, 379-425.
- [37] Kadet, Anne (2004), ‘The Price is Right’, *SmartMoney*, 1 December.
- [38] Keynes, John Maynard (1939), ‘Relative Movements of Real Wages and Output’, *Economic Journal*, 44, 34-51.
- [39] Lazear, Edward P. (1986), ‘Retail Pricing and Clearance Sales’, *American Economic Review*, 76, 14-32
- [40] Maskin, Eric and Jean Tirole (1988), ‘A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles’, *Econometrica*, 56, 571-599.
- [41] McAfee, R. Preston and Vera te Velde (2006), ‘Dynamic pricing in the airline industry’, in T.J. Hendershott (Ed.), *Handbook on Economics and Information Systems*, Elsevier, Amsterdam.
- [42] Milgrom, Paul and Ilya Segal (2002), ‘Envelope Theorems for Arbitrary Choice Sets’, *Econometrica*, 70, 583-601.
- [43] Murphy, Kevin M., Andrei Shleifer and Robert W. Vishny (1989), ‘Building Blocks of Market Clearing Business Cycle Models’, *NBER Macroeconomics Annual*, 4, 247-302.
- [44] Myrdal, Gunnar (1968). *Asian Drama: An Inquiry into the Poverty of Nations*. Pantheon, New York.
- [45] Nocke, Volker and Martin Peitz (2007), ‘A Theory of Clearance Sales’, *Economic Journal*, 117, 964-990.
- [46] Pai, Mallesh M. and Rakesh Vohra (2011), ‘Optimal Dynamic Auctions and Simple Index Rules’, mimeo Northwestern University.
- [47] Pashigian, B. Peter and Brian Bowen (1991), ‘Why Are Products Sold on Sale?: Explanations of Pricing Regularities’, *Quarterly Journal of Economics*, 106, 1015-1038.
- [48] Pavan, Alessandro, Ilya Segal and Juuso Toikka (2011), ‘Dynamic Mechanism Design: Incentive Compatibility, Profit Maximization and Information Disclosure’, mimeo Northwestern University and Stanford University.
- [49] Pesendorfer, Martin (2002), ‘Retail Sales: A Study of Pricing Behavior in Supermarkets’, *Journal of Business*, 75, 33-66.
- [50] Pigou, Arthur (1927), *Industrial Fluctuations*, Macmillan, London.
- [51] Riley, John and Richard Zeckhauser (1983), ‘Optimal Selling Strategies: When to Haggle, When to Hold Firm’, *Quarterly Journal of Economics*, 98, 267-289.
- [52] Rosenthal, Robert W. (1980), ‘A Model in which an Increase in the Number of Sellers Leads to a Higher Price’, *Econometrica*, 48, 1575-1579.
- [53] Rosenthal, Robert W. (1982), ‘A Dynamic Oligopoly Game with Lags in Demand: More on Monotonicity of Price in the Number of Sellers’, *International Economic Review*, 23, 353-360.

- [54] Rotemberg, Julio J. and Garth Saloner (1986), 'A Supergame-Theoretic Model of Price Wars during Booms', *American Economic Review*, 76, 390-407.
- [55] Said, Maher (2011), 'Auctions with Dynamic Populations: Efficiency and Revenue Maximization', mimeo Washington University in St. Louis.
- [56] Salop, Steven (1977), 'The Noisy Monopolist: Imperfect Information, Price Dispersion and Price Discrimination', *Review of Economic Studies*, 44, 393-406.
- [57] Salop, Steven and Joseph Stiglitz (1977), 'Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion', *Review of Economic Studies*, 44, 493-510.
- [58] Shilony, Yuval (1977), 'Mixed Pricing in Oligopoly', *Journal of Economic Theory*, 14, 373-388.
- [59] Sobel, Joel (1984), 'The Timing of Sales', *Review of Economic Studies*, 1984, 353-368.
- [60] Sobel, Joel (1991), 'Durable Goods Monopoly with Entry of New Consumers', *Econometrica*, 1991, 1455-1485.
- [61] Stokey, Nancy L. (1979), 'Intertemporal Price Discrimination', *Quarterly Journal of Economics*, 93, 355-371.
- [62] Su, Xuanming (2007), 'Intertemporal Pricing with Strategic Customer Behavior', *Management Science*, 53, 726-741.
- [63] Sutter, John D. (2010, November 29), 'Apple cuts iPad price to \$458', *CNN*, <http://www.cnn.com/2010/TECH/gaming.gadgets/11/26/apple.black.friday.sale/index.html>.
- [64] Tonello, Michael (2008), *Bringing Home the Birkin: My Life in Hot Pursuit of the World's Most Coveted Handbag*, HarperCollins, New York.
- [65] Varian, Hal R. (1980), 'A Model of Sales', *American Economic Review*, 70, 651-659.
- [66] Villas-Boas J. Miguel (2004), 'Price Cycles in Markets with Consumer Recognition', *RAND Journal of Economics*, 35, 486-501.
- [67] Waldman, Michael (2003), 'Durable Goods Theory for Real World Markets', *Journal of Economic Perspectives*, 17, 131-154.
- [68] Wang, Ruqu (1993), 'Auctions versus Posted-Price Selling', *American Economic Review*, 83, 838-851.
- [69] 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.
- [70] Zhang, Yuzhe (2009), 'Dynamic contracting with persistent shocks', *Journal of Economic Theory*, 144, 635-675.

Appendix A: Proofs of results on posted prices

This appendix provides proofs of the results in Section 3. Result 1 and Example 2 follow from arguments in the main text, and so proofs are omitted.

Proof of Lemma 1. Let A be any sales policy and let $t \in A$. Incentive compatibility of $\mathcal{M}_{P,A}$ and consistency of $\mathcal{M}_{P,A}$ with A imply that, for any stopping rule $\sigma \in \Sigma$,

$$\begin{aligned} \mathbb{E} \left[e^{-r(\tilde{\mu}_{x_A} - t)} \left(\tilde{\theta}_{\tilde{\mu}_{x_A}} - p_A(\tilde{\mu}_{x_A}) \right) \mid \tilde{\theta}_t = \theta_H \right] &\geq \theta_H - p_A(t) \\ &= u_t^{\mathcal{M}_{P,A}}(\theta_L; x_A) + \theta_H - \theta_L \\ &\geq \mathbb{E} \left[e^{-r(\tilde{\mu}_\sigma - t)} \left(\tilde{\theta}_{\tilde{\mu}_\sigma} - p_A(\tilde{\mu}_\sigma) \right) \mid \tilde{\theta}_t = \theta_L \right] + \theta_H - \theta_L. \end{aligned} \quad (22)$$

Suppose that $(\tilde{\psi}_s)_{s \geq t}$ is determined independently and identically to $(\tilde{\theta}_s)_{s \geq t}$. Then (22) and independence of the two stochastic processes imply that

$$\mathbb{E} \left[e^{-r(\mu_{x_A}(\tilde{\theta}^{[t, \infty)}) - t)} \left(\tilde{\theta}_{\mu_{x_A}(\tilde{\theta}^{[t, \infty)})} - \tilde{\psi}_{\mu_{x_A}(\tilde{\theta}^{[t, \infty)})} \right) \mid \tilde{\theta}_t = \theta_H, \tilde{\psi}_t = \theta_L \right] \geq \theta_H - \theta_L.$$

This is possible only if $x_A(\theta_H, t) = 1$. This establishes Part (i); Part (ii) is an immediate consequence. ■

Proof of Lemma 2. As explained in the text, if the buyer finds the mechanism $\mathcal{M}_{P,A}^* = \langle p_A^*, x_A^* \rangle$ incentive compatible, then this mechanism must maximize expected profit subject to consistency with A . That is, it is enough to verify x_A^* is indeed an optimal stopping rule for the buyer given p_A^* , so that the value of the buyer's problem is given by $\nu_t^{LB,A}$.

For any stopping rule $\sigma \in \Sigma$, any t , and any $\theta_t \in \{\theta_L, \theta_H\}$,

$$\begin{aligned} &\mathbb{E}_t \left[e^{-r(\tilde{\mu}_\sigma - t)} \left(\tilde{\theta}_{\tilde{\mu}_\sigma} - p_A^*(\tilde{\mu}_\sigma) \right) \mid \tilde{\theta}_t = \theta_t \right] \\ &\leq \mathbb{E}_t \left[e^{-r(\tilde{\mu}_\sigma - t)} \nu_{\tilde{\mu}_\sigma}^{LB,A} \left(\tilde{\theta}_{\tilde{\mu}_\sigma} \right) \mid \tilde{\theta}_t = \theta_t \right] \\ &\leq \nu_t^{LB,A}(\theta_t) \\ &\quad + \mathbb{E} \left[\int_t^{\tilde{\mu}_\sigma} e^{-r(s-t)} \left(\begin{aligned} &1 \left(\tilde{\theta}_s = \theta_L \right) \left(\begin{aligned} &-r\nu_s^{LB,A}(\theta_L) + \frac{\partial \nu_s^{LB,A}(\theta_L)}{\partial s} \\ &+\alpha_L \left(\nu_s^{LB,A}(\theta_H) - \nu_s^{LB,A}(\theta_L) \right) \end{aligned} \right) \\ &+ 1 \left(\tilde{\theta}_s = \theta_H \right) \left(\begin{aligned} &-r\nu_s^{LB,A}(\theta_H) + \frac{\partial \nu_s^{LB,A}(\theta_H)}{\partial s} \\ &+\alpha_H \left(\nu_s^{LB,A}(\theta_L) - \nu_s^{LB,A}(\theta_H) \right) \end{aligned} \right) \end{aligned} \right) ds \mid \tilde{\theta}_t = \theta_t \right] \\ &\leq \nu_t^{LB,A}(\theta_t) \end{aligned}$$

where $1(\cdot)$ is the indicator function. The first inequality follows by choice of the function $p_A^*(\cdot)$. The second inequality follows by applying Dynkin's formula after using that (by the assumption that A is a countable union of intervals and points) $\nu_s^{LB,A}(\theta_H)$ contains only countably many downward discontinuities and is continuously differentiable elsewhere, and that $\nu_s^{LB,A}(\theta_L)$ is continuously differentiable except at the

discontinuities of $\nu_s^{LB,A}(\theta_H)$. The third follows because

$$-r\nu_t^{LB,A}(\theta_L) + \frac{\partial \nu_t^{LB,A}(\theta_L)}{\partial t} + \alpha_L \left(\nu_t^{LB,A}(\theta_H) - \nu_t^{LB,A}(\theta_L) \right) = 0$$

and

$$-r\nu_t^{LB,A}(\theta_H) + \frac{\partial \nu_t^{LB,A}(\theta_H)}{\partial t} + \alpha_H \left(\nu_t^{LB,A}(\theta_L) - \nu_t^{LB,A}(\theta_H) \right) \leq 0$$

wherever the derivatives exist. When $\sigma = x_A^*$, all inequalities hold with equality (the second because the stopping is either immediate or is before any discontinuity in $\nu_t^{LB,A}$). Thus $\nu_t^{LB,A}$ is the value function associated with the buyer's problem and x_A^* is an optimal strategy for the buyer. ■

Proof of Proposition 1.

A first step to understanding the problem of choosing a profit-maximizing sales policy is to derive the expected profit for the seller for a given sales policy A . Here, I merely require that A be left-closed and derive expected revenue assuming that all such policies are implementable by the prices specified in Lemma 2. It turns out that the optimal policy always satisfies the mild restriction of Lemma 2 that it be a countable collection of intervals and points. Because none of the arguments rely on the assumption $\gamma = \gamma^S$, I will simply require that $\gamma \in (0, 1)$ as in the model set-up.

Using integration by parts, the buyer's ex-ante expected rent for a mechanism $\mathcal{M}_{P,A}^* = \langle p_A^*, x_A^* \rangle$ as defined in Lemma 2 is equal to

$$\begin{aligned} RENT^{\mathcal{M}_{P,A}^*} &= \gamma \int_0^\infty \lambda e^{-(r+\lambda)\tau} \nu_\tau^{LB,A}(\theta_H) d\tau + (1-\gamma) \int_0^\infty \lambda e^{-(r+\lambda)\tau} \nu_\tau^{LB,A}(\theta_L) d\tau \\ &= (\theta_H - \theta_L) \int_0^\infty e^{-(r+\lambda)\tau} (\alpha_L (e^{\lambda\tau} - 1) + \gamma\lambda) e^{-(r+\alpha_L+\alpha_H)(m_A(\tau)-\tau)} d\tau. \end{aligned}$$

The total expected surplus from the mechanism $\mathcal{M}_{P,A}^*$ is

$$SURP^{\mathcal{M}_{P,A}^*} = \int_0^\infty \lambda e^{-(r+\lambda)\tau} \left(\gamma\theta_H + (1-\gamma) \left(\frac{\alpha_L}{r+\alpha_L} (1 - e^{-(r+\alpha_L)(m_A(\tau)-\tau)}) \theta_H + e^{-(r+\alpha_L)(m_A(\tau)-\tau)} \theta_L \right) \right) d\tau.$$

Therefore, the seller's expected profit is equal to

$$\begin{aligned} PROF^{\mathcal{M}_{P,A}^*} &= SURP^{\mathcal{M}_{P,A}^*} - RENT^{\mathcal{M}_{P,A}^*} \\ &= \int_0^\infty \lambda e^{-(r+\lambda)\tau} b(\tau, m_A(\tau)) d\tau \end{aligned} \quad (23)$$

where

$$\begin{aligned} b(\tau, m) &= \gamma\theta_H + (1-\gamma) \left(\frac{\alpha_L}{r+\alpha_L} \theta_H + e^{-(r+\alpha_L)(m-\tau)} \left(\theta_L - \frac{\alpha_L}{r+\alpha_L} \theta_H \right) \right) \\ &\quad - \frac{1 - e^{-\lambda\tau}}{\lambda e^{-\lambda\tau}} \alpha_L e^{-(r+\alpha_L+\alpha_H)(m-\tau)} (\theta_H - \theta_L) - \gamma e^{-(r+\alpha_L+\alpha_H)(m-\tau)} (\theta_H - \theta_L). \end{aligned} \quad (24)$$

If $m_A(\tau)$ may be chosen without constraint, $PROF^{\mathcal{M}_{P,A}^*}$ can be maximized by simply maximizing pointwise under the integral (see the solution to the problem for the fully-optimal mechanism). However, $m_A(\tau)$ must be given by (5) for some sales policy A . For instance, if A is taken to be discrete, then $m_A(\cdot)$ is a step function and it is easy to see that pointwise maximization is not possible; i.e., the problem is not separable across the dates τ . However, the problem of selecting the policy subsequent to any date τ with $m_A(\tau) = \tau$, i.e. after any date $\tau \in A$, is separable. The following observation therefore allows one to focus on the problem of choosing the optimal sales policy after date T (the “continuation problem”).

Lemma 4 *It is always optimal for the sales policy to contain dates $[0, T]$, where T is defined by (10).*

Proof. Note that

$$\frac{\partial b(\tau, m)}{\partial m} = e^{-(r+\alpha_L)(m-\tau)} \left(\begin{array}{l} -(1-\gamma)\omega_L + \gamma e^{-\alpha_H(m-\tau)}(\omega_H - \omega_L) \\ + \frac{\alpha_L(e^{\lambda\tau}-1)}{\lambda} e^{-\alpha_H(m-\tau)}(\omega_H - \omega_L) \end{array} \right).$$

For all $\tau \leq T$ and $m \geq \tau$, $\frac{\partial b(\tau, m)}{\partial m} \leq 0$. Hence, for all $\tau \leq T$, $b(\tau, m)$ is maximized by choosing $m = \tau$. The separability of the continuation problem at dates in the sales policy described above then implies the result. ■

Now consider the continuation problem. The solution to this problem is a set of dates above date T , which I refer to as the “continuation policy” (to avoid possible confusion with the sales policies A defined on \mathbb{R}_+ , continuation policies are denoted by $B \subset [T, \infty)$).

The problem is most easily understood with a recursive formulation. To show that a solution exists, it is helpful to consider a sequence of constrained problems where the sales policy is required to be discrete, with the distance between dates in the sales policy bounded above zero (Denardo (1967) uses this approach to solve a similar problem). Thus define, for any $\varepsilon > 0$ and all $t \geq T$, the functional equation

$$W(t) = \sup_{t' \geq t+\varepsilon} \left\{ \begin{array}{l} \phi(t'-t) - (e^{\lambda t} - 1) \Pr(\tilde{\theta}_{t'} = \theta_H | \tilde{\theta}_t = \theta_L) e^{-r(t'-t)} (\theta_H - \theta_L) \\ + e^{-(r+\lambda)(t'-t)} W(t'), \\ \frac{\lambda}{r+\lambda} \left(\gamma + \frac{\alpha_L(1-\gamma)}{\alpha_L+r} \right) \theta_H \end{array} \right\}, \quad (25)$$

where $\phi(z)$ is given by (12). Here, I include explicitly the option to hold no further sales, which has value $\frac{\lambda}{r+\lambda} \left(\gamma + \frac{\alpha_L(1-\gamma)}{\alpha_L+r} \right) \theta_H$ (thus, any date t in the sales policy must be the last date if the value $\frac{\lambda}{r+\lambda} \left(\gamma + \frac{\alpha_L(1-\gamma)}{\alpha_L+r} \right) \theta_H$ is uniquely attained by this option). Otherwise, the interpretation is the same as for the problem described in the main text (where the distance between dates is not constrained) except that any subsequent date in the sales policy t' must be at least ε above t .

The proof of existence of an optimal policy is as follows. Lemma 5 establishes that, for any $\eta > 0$, the number of points above $T + \eta$ is bounded uniformly across ε . Lemma 6 establishes the existence of a value $t^{\max} > T$ such that any continuation policy which solves the ε -constrained problem, irrespective of the value ε , contains no dates above t^{\max} . Using Lemma 6 and standard arguments in dynamic programming,

Lemma 7 establishes the existence of a solution to the ε -constrained problem. Lemma 8 establishes that the solution to the ε -constrained problem converges to the solution to the problem of interest.

That any (closed)⁴³ optimal policy satisfies Parts (i) and (ii) of the proposition follows from applying the same arguments as in Lemmas 4, 5, and 6 (for Lemmas 5 and 6, these arguments must now be applied for the unconstrained problem). Part (i) requires also to argue that there is at least one point above T in any optimal policy, which is the result in Lemma 9.

Lemma 5 *For every $\eta > 0$, there is $\zeta_\eta > 0$ such that, for any $\varepsilon > 0$, any optimal ε -constrained continuation policy B_ε^* , if t', t'', t''' are dates in B_ε^* with $t''' > t'' > t' \geq T + \eta$, then $t''' - t' > \zeta_\eta$. Hence, for any $\eta > 0$, there exists $N_\eta \in \mathbb{N}$ such that, for any $\varepsilon > 0$, the number of points in any ε -constrained continuation policy B_ε^* above $T + \eta$ is no greater than N_η .*

Proof. Suppose with a view to contradiction that there is a value $\eta > 0$ for which there exists no value $\zeta_\eta > 0$ satisfying the property in the lemma. Note that there exists $\delta_\eta > 0$ such that, for all $\tau \geq T + \eta$, and all $m \in [\tau, \tau + \delta_\eta)$

$$\frac{\partial b(\tau, m)}{\partial m} > 0,$$

where b is defined by (24). By the above assumption, there exists an optimal ε -constrained continuation policy B_ε^* containing dates t', t'', t''' in B_ε^* with $t''' > t'' > t' \geq T + \eta$ such that $t''' - t' < \delta_\eta$. However, since b is increasing in m over the relevant values, $b(\tau, t''') > b(\tau, t'')$ for all $\tau \in (t', t'']$. It follows immediately that omitting t'' from the policy B_ε^* strictly increases profit, a contradiction. ■

Lemma 6 *There exists $t^{\max} > T$ such the solution to any ε -constrained problem is bounded below t^{\max} .*

Proof. Suppose with a view to contradiction that there is a solution to an ε -constrained problem, for some $\varepsilon > 0$, which is unbounded. It therefore contains an increasing sequence $(t'_n)_{n=1}^\infty$ which is unbounded. For each n , let $t''_n > t'_n$ be the next element in the solution.

Denote the value function of the ε -constrained problem by W_ε^* , which must be a solution to (25). Note that W_ε^* must be non-increasing and never take values less than $\frac{\lambda}{r+\lambda} \left(\gamma + \frac{\alpha_L(1-\gamma)}{\alpha_L+r} \right) \theta_H$.

First, there must exist $\delta > 0$ and $N \in \mathbb{N}$ such that $n > N$ implies $t''_n - t'_n \geq \delta$. To see this, note that $\pi'(z)$ is uniformly bounded over z in a sufficiently small neighborhood of zero. Also

$$\frac{\partial}{\partial z} \left(e^{\lambda t'_n} - 1 \right) \Pr \left(\tilde{\theta}_{t'_n+z} = \theta_H | \tilde{\theta}_{t'_n} = \theta_L \right) e^{-rz} (\theta_H - \theta_L)$$

grows without bound with n , uniformly over z in a sufficiently small neighborhood of zero. These observations, together with the fact that $W_\varepsilon^*(\cdot)$ is non-increasing, imply the result.

Next, there exists \bar{z} sufficiently large that $t''_n - t'_n \leq \bar{z}$ irrespective of n . To see this, notice that, for

⁴³The reason for considering closed policies is simply to ease exposition; the only other possibility is that the date T is omitted from an optimal policy.

$z > 0$,

$$\begin{aligned}
\phi(z) &= \gamma \left(\frac{\lambda \theta_H}{\lambda+r} (1 - e^{-(\lambda+r)z}) - \frac{\alpha_L(\theta_H - \theta_L)}{\alpha_L + \alpha_H} e^{-rz} (1 - e^{-\lambda z}) \right) \\
&\quad - \frac{\lambda \alpha_H (\theta_H - \theta_L)}{\alpha_L + \alpha_H} e^{-(r+\lambda)z} \frac{1 - e^{-(\alpha_L + \alpha_H - \lambda)z}}{\alpha_L + \alpha_H - \lambda} \\
&+ (1 - \gamma) \left(\theta_L e^{-rz} (1 - e^{-\lambda z}) + \frac{\lambda \alpha_L \theta_H}{(r+\lambda)(\alpha_L+r)} (1 - e^{-(r+\lambda)z}) \right) \\
&\quad + \lambda \left(\theta_L - \frac{\alpha_L \theta_H}{\alpha_L+r} \right) e^{-(\lambda+r)z} \frac{1 - e^{-(\alpha_L - \lambda)z}}{\alpha_L - \lambda} \\
&\quad - \left(\frac{\alpha_L(\theta_H - \theta_L)}{\alpha_L + \alpha_H} + \theta_L \right) e^{-rz} (1 - e^{-\lambda z}) \\
&\quad + \frac{\lambda \alpha_L (\theta_H - \theta_L)}{\alpha_L + \alpha_H} e^{-(\lambda+r)z} \frac{(1 - e^{-(\alpha_L + \alpha_H - \lambda)z})}{\alpha_L + \alpha_H - \lambda} \\
&= \frac{\lambda}{r+\lambda} \left(\gamma + (1 - \gamma) \frac{\alpha_L}{\alpha_L+r} \right) \theta_H - \frac{\alpha_L (\theta_H - \theta_L)}{\alpha_L + \alpha_H} e^{-rz} + o(e^{-rz}),
\end{aligned}$$

where the first equality holds whenever $\lambda = \alpha_L, \alpha_L + \alpha_H$ only after taking the appropriate limits. The result then follows from noting that W_ε^* is non-increasing and discounted by $e^{-(r+z)} \in o(e^{-rz})$ in the maximization problem of (25), implying that, for \bar{z} large enough, $t_n'' - t_n' > \bar{z}$ is inconsistent with optimality. This follows from noting that a value of $\frac{\lambda}{r+\lambda} \left(\gamma + (1 - \gamma) \frac{\alpha_L}{\alpha_L+r} \right) \theta_H$ can be obtained by including no further dates after t_n' (\bar{z} is chosen so that including $t_n'' > t_n' + \bar{z}$ achieves a strictly lower value).

Finally, note that $-(e^{\lambda t_n'} - 1) \Pr(\tilde{\theta}_{t_n'+z} = \theta_H | \tilde{\theta}_{t_n'} = \theta_L) e^{-rz} (\theta_H - \theta_L)$ becomes arbitrarily small with n uniformly across $z \in [\delta, \bar{z}]$, whilst $\phi(z)$ is uniformly bounded on $[\delta, \bar{z}]$. This contradicts the observation that W_ε^* remains above $\frac{\lambda}{r+\lambda} \left(\gamma + (1 - \gamma) \frac{\alpha_L}{\alpha_L+r} \right) \theta_H$. ■

Lemma 7 *For any $\varepsilon > 0$, there exists an optimal ε -constrained continuation policy $B_\varepsilon^* \subset [T, \infty)$.*

Proof. Consider the functional operator defined by the right-hand side of (25) on the space of bounded continuous functions $W : [T, \infty) \rightarrow \mathbb{R}_+$ with sup norm. This operator satisfies the ‘‘monotonicity’’ and ‘‘discounting’’ conditions of Blackwell’s result and is hence a contraction mapping. By the contraction mapping theorem, there exists a unique fixed point of (25); this is the value function $W_\varepsilon^*(\cdot)$ considered in the previous lemma. Importantly, since the space of functions considered is complete, this shows that $W_\varepsilon^*(\cdot)$ is continuous. Moreover, since one can restrict attention to points below t^{\max} , as defined in Lemma 6, solutions to the relevant maximization problems exist by the extreme value theorem. Denote by B_ε^* the union of the members of the sequence of solutions starting at T . A standard verification argument shows that B_ε^* is the solution to the ε -constrained continuation problem. ■

Lemma 8 *A solution B^* to the unconstrained continuation problem exists.*

Proof. Let $\varepsilon_n = \frac{1}{n}$ for all $n \in \mathbb{N}$. Let $(B_{\varepsilon_n}^*)_{n=1}^\infty$ be a sequence of solutions to the ε_n -constrained problem. Using the axiom of choice, the property of Lemma 5 and the compactness of the interval $[T, t^{\max}]$ (with t^{\max} defined in Lemma 6), it is easy to see that there must exist a subsequence of $(B_{\varepsilon_n}^*)_{n=1}^\infty$ which is convergent in an appropriate metric to some set B^* . In particular, for each $k \in \mathbb{N}$, there exists a subsequence of $(B_{\varepsilon_n}^*)_{n=1}^\infty$

such that the l^{th} -to-last point in each set is convergent (in Euclidean distance) to the l^{th} -to-last point in B^* uniformly over $l \leq k$. An appropriate subsequence of $(B_{\varepsilon_n}^*)_{n=1}^\infty$ therefore converges to B^* taking the metric to be the uniform distance between l^{th} -to-last elements (taking the minimum element of the set in case no l^{th} -to-last element exists). Since the expected profit is continuous in this metric (see (23)), it must be that B^* is a solution to the unconstrained problem. ■

Lemma 9 *Suppose that (4) holds. Any (unconstrained) optimal continuation policy contains at least one date above T .*

Proof. Total expected revenue from a sales policy $[0, \bar{t}]$, $\bar{t} \geq 0$, is given by

$$\begin{aligned} & \left(1 - e^{-(\lambda+r)\bar{t}}\right) \frac{\lambda}{\lambda+r} \frac{\omega_L}{r} + e^{-(r+\lambda)\bar{t}} \frac{\lambda}{\lambda+r} \left(\gamma + \frac{\alpha_L}{\alpha_L+r} (1-\gamma)\right) \theta_H \\ & + e^{-r\bar{t}} \left(1 - e^{-\lambda\bar{t}}\right) \left(\theta_L - \frac{\omega_L}{r}\right). \end{aligned}$$

A first-order condition yields that, given (4) holds, the optimal value of \bar{t} is positive and equal to

$$\frac{1}{\lambda} \ln \left(1 + \frac{\lambda r \left((1-\gamma)\theta_L - \left(\frac{\alpha_L}{r} + \gamma\right)(\theta_H - \theta_L)\right)}{\alpha_L (\alpha_L + r) (\theta_H - \theta_L)}\right) > T.$$

(This comparison is straightforward after noticing that either $T = 0$ or

$$T = \frac{1}{\lambda} \log \left(1 + \frac{\lambda r \left(\theta_L (1-\gamma) - (\theta_H - \theta_L) \left(\frac{\alpha_L}{r} + \gamma + \frac{\alpha_H}{r} \gamma\right)\right)}{\alpha_L (\alpha_L + \alpha_H + r) (\theta_H - \theta_L)}\right).$$

Hence, including only dates $[0, T]$ in the sales policy is suboptimal. ■

■

Proof of Result 2. Step 1. First, I verify that W_{LR}^* cannot be approached by taking z arbitrarily close to zero in (14). The value of the auxiliary program from a sales policy A is

$$V_{LR}(A) = \int_0^\infty \lambda e^{-(r+\lambda)\tau} b_{LR}(\tau, m_A(\tau); d) d\tau,$$

where

$$\begin{aligned} b_{LR}(\tau, m; d) &= \gamma \theta_H + (1-\gamma) \left(\frac{\alpha_L}{d+r+\alpha_L} \left(1 - e^{-(r+d+\alpha_L)(m-\tau)}\right) \theta_H + e^{-(r+d+\alpha_L)(m-\tau)} \theta_L\right) \\ &\quad - (\theta_H - \theta_L) \left(\alpha_L \frac{\lambda}{d-\lambda} + \gamma\right) e^{-(r+d+\alpha_L+\alpha_H)(m-\tau)}. \end{aligned}$$

The claim follows from the observation that, for any τ , it must be that

$$\frac{\partial b_{LR}(\tau, m; d)}{\partial m} > 0$$

for m in some sufficiently small neighborhood of (but greater than) τ . If this observation is not true, it is easy to show that $\frac{\partial b_{LR}(\tau, m; d)}{\partial m} < 0$ for all $m > \tau$, and this contradicts the optimality of a discrete solution to the auxiliary problem.

Step 2. For any $z > 0$, let

$$X(z) = W_{LR}^* - \frac{\phi(z; d) - \Pr\left(\tilde{\theta}_z = \theta_H | \tilde{\theta}_0 = \theta_L\right) e^{-(r+d)z} \frac{\lambda}{d-\lambda} (\theta_H - \theta_L)}{1 - e^{-(r+\lambda)z}}.$$

For any date t and any $z > 0$, let

$$Y(t, z) = \frac{W^*(t) - e^{-(r+\lambda)z} W^*(t+z)}{1 - e^{-(r+\lambda)z}} - \frac{\phi(z; d) - \Pr\left(\tilde{\theta}_{t+z} = \theta_H | \tilde{\theta}_t = \theta_L\right) e^{-(r+d)z} \frac{\lambda}{d-\lambda} (1 - e^{-(d-\lambda)t}) (\theta_H - \theta_L)}{1 - e^{-(r+\lambda)z}}.$$

I now establish that $X(z) \rightarrow Y(t, z)$ as $t \rightarrow \infty$ uniformly over $z > 0$. The key observation is that

$$\frac{W^*(t) - e^{-(r+\lambda)z} W^*(t+z)}{1 - e^{-(r+\lambda)z}} \rightarrow W_{LR}^* \quad (26)$$

as $t \rightarrow \infty$ uniformly over $z > 0$. To see this, note that, by the envelope theorem (see Milgrom and Segal, 2002), $W^*(\cdot)$ is absolutely continuous, and

$$e^{-(r+\lambda)z} W^*(t+z) = W^*(t) + \int_0^z e^{-(r+\lambda)s} (W^{*'}(t+s) - (r+\lambda) W^*(t+s)) ds, \quad (27)$$

where $W^{*'}(t+s)$ satisfies, for almost all $s > 0$,

$$|W^{*'}(t+s)| \leq \lambda e^{-(d-\lambda)t} (\theta_H - \theta_L).$$

From (27),

$$\frac{W^*(t) - e^{-(r+\lambda)z} W^*(t+z)}{1 - e^{-(r+\lambda)z}} = \frac{(r+\lambda) \int_0^z e^{-(r+\lambda)s} W^*(t+s) ds - \int_0^z e^{-(r+\lambda)s} W^{*'}(t+s) ds}{1 - e^{-(r+\lambda)z}}$$

The observation that $W^*(t) \rightarrow W_{LR}^*$, and that $W^{*'}(t) \rightarrow 0$, then implies (26).

Step 3. It follows from Part 1 that $X(z)$ is bounded above zero for z in a sufficiently small neighborhood of zero. Since z^* uniquely attains W_{LR}^* in (14), and since this value strictly exceeds that approached as $z \rightarrow \infty$ in (14), $X(z)$ is also bounded above zero for all sufficiently large z . Using the continuity of X , the following must therefore be true. For any $\varepsilon > 0$, there is η_ε such that $X(z) - X(z^*) = X(z) > \eta_\varepsilon$ for all $z > 0$ with $|z - z^*| \geq \varepsilon$.

Step 4. Now, suppose with a view to contradiction that the result does not hold. Note that $W^*(t)$ can be approximated arbitrarily closely by choosing a discrete unbounded sales policy. Using this, it follows that there must exist $\varepsilon > 0$ such that, for all \tilde{t} , there is $t > \tilde{t}$ with the following property. For all $\delta > 0$, there exists $z_{t,\delta} > 0$ such that $|z_{t,\delta} - z^*| \geq \varepsilon$ and $Y(t, z_{t,\delta}) < \delta$.

Take \check{t} sufficiently large that $|X(z) - Y(t, z)| < \frac{\eta_\varepsilon}{3}$ for all $t > \check{t}$ and all $z > 0$ (such a choice is possible by Step 2). Let $t > \check{t}$ be chosen to have the aforementioned property and let $\delta = \frac{\eta_\varepsilon}{3}$, with a corresponding choice of $z_{t,\delta}$. Since $Y(t, z) \geq 0$ for all z , it follows that,

$$\begin{aligned} X(z_{t,\delta}) &= X(z_{t,\delta}) - Y(t, z_{t,\delta}) + Y(t, z_{t,\delta}) - Y(t, z^*) + Y(t, z^*) - X(t, z^*) \\ &< \eta_\varepsilon, \end{aligned}$$

contradicting the property of η_ε in Step 3. ■

Proof of Result 3. Consider any sales policy A and suppose the buyer arrives at date τ_{mass} . The present value of expected profit given this arrival date (for conditionally-optimal prices; see Lemma 2) is

$$PROF_{mass}^{\mathcal{M}_{P,A}^*} = e^{-r\tau_{mass}} R(m_A(\tau_{mass}) - \tau_{mass}) - \int_{m_A(\tau_{mass})}^{\infty} e^{-rs} \alpha_L e^{-(r+\alpha_L+\alpha_H)(m_A(s)-s)} (\theta_H - \theta_L) ds,$$

where R is given by (2). Since (4) holds, a sales policy A maximizes $PROF_{mass}^{\mathcal{M}_{P,A}^*}$ iff $\max\{A\} = \tau_{mass}$.

Let $\varepsilon > 0$. There exists $\eta_\varepsilon > 0$ such that, if A is a sales policy with $m_A(\tau_{mass}) \geq \tau_{mass} + \varepsilon$, then $PROF_{mass}^{\mathcal{M}_{P,\{\tau_{mass}\}}^*} - PROF_{mass}^{\mathcal{M}_{P,A}^*} \geq \eta_\varepsilon$.

For any sales policy A , the expected profit conditional on arrival according to the exponential distribution with parameter λ is $PROF^{\mathcal{M}_{P,A}^*}$ in (23). Let $\bar{\rho} = \frac{PROF_{mass}^{\mathcal{M}_{P,A^*}^*}}{\eta_\varepsilon + PROF_{mass}^{\mathcal{M}_{P,A^*}^*}}$ where A^* maximizes $PROF^{\mathcal{M}_{P,A}^*}$. Then, for any $\rho \geq \bar{\rho}$ and any A that includes no date in $[\tau_{mass}, \tau_{mass} + \varepsilon)$, expected profit for the distribution of arrival times in the result is

$$\begin{aligned} \rho PROF_{mass}^{\mathcal{M}_{P,A}^*} + (1 - \rho) PROF^{\mathcal{M}_{P,A}^*} &\leq \rho \left(PROF_{mass}^{\mathcal{M}_{P,\{\tau_{mass}\}}^*} - \eta_\varepsilon \right) + (1 - \rho) PROF^{\mathcal{M}_{P,A^*}^*} \\ &\leq \rho PROF_{mass}^{\mathcal{M}_{P,\{\tau_{mass}\}}^*} + PROF^{\mathcal{M}_{P,A^*}^*} - \bar{\rho} \left(PROF^{\mathcal{M}_{P,A^*}^*} + \eta_\varepsilon \right) \\ &= \rho PROF_{mass}^{\mathcal{M}_{P,\{\tau_{mass}\}}^*}, \end{aligned}$$

so that A achieves less profit than the policy $\{\tau_{mass}\}$. ■

Appendix B: Proofs of results on the fully-optimal mechanism

This Appendix provides proofs of the results on the fully-optimal mechanism. As explained in the main text, since the seller can fully commit, the revelation principle applies in this environment. One therefore needs only to consider incentive-compatible direct mechanisms, i.e. mechanisms in which the buyer finds it optimal to participate upon arrival and subsequently report his values truthfully.⁴⁴

Direct mechanisms. In order to define the direct mechanism, one needs to define the probability space for the buyer's information. Let $a(\cdot) : \mathbb{R}_+ \rightarrow \{0, 1\}$ denote a sample path of the arrival process, i.e. a right-continuous function such that $a(t) = 1$ if and only if he has arrived by date t . Let $\tau_a = \inf_{t \geq 0} \{a(t)\}$. For any $t \geq 0$, let $a^{[0,t]}$ denote the restriction of a to $[0, t]$. A sample path for the joint process is therefore $(a, \theta^{[\tau_a, \infty)})$ where $\theta^{[\tau_a, \infty)}(\cdot)$ is right continuous and has a finite number of jumps on any finite interval. Let Ω be the set of all sample paths and endow it with the Borel filtration $\mathcal{F} = \{\mathcal{F}_t\}_{t \geq 0}$, where each \mathcal{F}_t is the sigma-algebra defined on feasible histories (i.e., restrictions of some sample path) $(a^{[0,t]}, \theta^{[\tau_a, t]})$. Define by P the probability measure consistent with the Markov processes for buyer arrival and the transition of values defined in the model set-up. Then, (Ω, \mathcal{F}, P) is a probability space.

The buyer can choose to report for the first time to the seller at any date after the arrival date τ_a . Let $\hat{\tau} : \Omega \rightarrow \mathbb{R}_+$ be a stopping time such that $\hat{\tau}(a, \theta^{[\tau_a, \infty)}) \geq \tau_a$ for all $(a, \theta^{[\tau_a, \infty)}) \in \Omega$. The buyer reports a value at each $t \geq \hat{\tau}(a, \theta^{[\tau_a, \infty)})$. For any $t \geq \hat{\tau}(a, \theta^{[\tau_a, \infty)})$, let $\iota_t(a^{[0,t]}, \theta^{[\tau_a, t]}) \in \{\theta_L, \theta_H\}$ be a progressively measurable function which specifies the buyer's report at date t . The buyer is restricted to reports that coincide with realizations of the process. In case $\hat{\tau}(a, \theta^{[\tau_a, \infty)}) > t$ there is no report at or before t , and, assuming truthful reporting, the seller infers that $a(s) = 0$ for all $s \leq t$. For any interval I bounded below by $\hat{\tau}(a, \theta^{[\tau_a, \infty)})$, let $\iota^I(a^{[0,t]}, \theta^{[\tau_a, \infty)}) = (\iota_t(a^{[0,t]}, \theta^{[\tau_a, t]}))_{t \in I}$. Let h^I denote an arbitrary history of values over I . Histories may be denoted by concatenations: for example, $h^{[t, t']}h^{[t', \infty)}$ denotes a realization of reports from date t . A truthful reporting strategy is defined by $\hat{\tau}_{truth}(a, \theta^{[\tau_a, \infty)}) = \tau_a$ and $\iota_{truth, t}(a^{[0,t]}, \theta^{[\tau_a, t]}) = \theta_t$ for all $t \geq \tau_a$.

Without loss of generality, the direct mechanism \mathcal{M}_D can be taken to be a collection of stopping times $y = (y_\tau)_{\tau \geq 0}$ and a collection of functions $q = (q_{\tau, t}(\cdot))_{\tau \geq 0, t \geq \tau}$ such that, if the buyer's realized report is the sequence of values $\theta^{[\tau, \infty)}$ from time τ , then (i) $y_\tau(\theta^{[\tau, \infty)})$ denotes the time the buyer receives the good, and (ii) $q_{\tau, t}(\theta^{[\tau, t]}) \in \mathbb{R}$ denotes the date- t flow payment from the buyer to the seller. Both functions are required to be progressively measurable with respect to the filtration defined above (a reporting time of τ is interpreted as $a(t) = 0$ for $t < \tau$ and $a(t) = 1$ for $t \geq \tau$; however, here and in what follows, I omit dependence on the sample path a). Without any significant loss of generality, I restrict attention to payments q such that, for any τ , any $t > \tau$, and any $\theta^{[\tau, \infty)}$ with $y_\tau(\theta^{[\tau, \infty)}) < t$, $q_{\tau, t}(\theta^{[\tau, t]}) = 0$.

Seller's payoff. The date- τ value of the seller's expected profit from an incentive-compatible direct

⁴⁴This is a well understood principle in dynamic mechanism design. It follows from the same replication argument as in static mechanism design. For any mechanism \mathcal{M} , one can consider the direct mechanism \mathcal{M}_D which prescribes, for any history of reports of an arrival date τ and history of values $\theta^{[\tau, t]}$, the same outcome as obtained in equilibrium under the original mechanism when the buyer's actual history corresponds to the reported one. Since it is optimal for the buyer to choose the given outcome in the original mechanism (and since the attainable outcomes under the direct mechanism is a subset of those under the original mechanism), it is optimal for him to report truthfully in the direct mechanism.

mechanism \mathcal{M}_D conditional on arrival at τ is

$$\pi^{\mathcal{M}_D}(\tau) = \mathbb{E} \left[\int_{\tau}^{\infty} e^{-r(s-\tau)} q_{\tau,s} \left(\tilde{\theta}^{[\tau,s]} \right) ds \right].$$

Therefore, the present value of total expected profit is

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

Buyer payoffs and incentive compatibility. The buyer's expected continuation payoff at date t , given that he contracted at date $\hat{\tau} \leq t$ and has not yet received the good at any earlier date, depends only on the contract he signed at date $\hat{\tau}$ (as given by $y_{\hat{\tau}}$ and $(q_{\hat{\tau},t})_{t \geq \hat{\tau}}$), his value θ_t at date t (and, since the process is Markov, not on earlier values), the history of reports $h^{[\hat{\tau},t]}$, and his continuation reporting strategy $\iota^{[t,\infty)}$. Omitting dependence of the buyer's reporting strategy on irrelevant past values, this is

$$\begin{aligned} & u_t^{\mathcal{M}_D} \left(\theta_t; h^{[\hat{\tau},t]}, \iota^{[t,\infty)} \right) \\ = & \mathbb{E} \left[e^{-r(y_{\hat{\tau}}(h^{[\hat{\tau},t]}, \iota^{[t,\infty)})(\tilde{\theta}^{[t,\infty)})) - t} \tilde{\theta}_{y_{\hat{\tau}}(h^{[\hat{\tau},t]}, \iota^{[t,\infty)})(\tilde{\theta}^{[t,\infty)})} - \int_t^{\infty} e^{-r(s-t)} q_{\hat{\tau},s} \left(h^{[\hat{\tau},t]} \iota^{[t,s]} \left(\tilde{\theta}^{[t,s]} \right) \right) ds \mid \tilde{\theta}_t = \theta_t \right]. \end{aligned}$$

The value of the reporting problem at date $t \geq \hat{\tau}$ after he has contracted with the seller and reported $h^{[\hat{\tau},t]}$, when his date- t value is θ_t (and he has not yet received the good), is

$$\nu_t^{\mathcal{M}_D} \left(\theta_t; h^{[\hat{\tau},t]} \right) = \sup_{\iota^{[t,\infty)}} u_t^{\mathcal{M}_D} \left(\theta_t; h^{[\hat{\tau},t]}, \iota^{[t,\infty)} \right).$$

The mechanism \mathcal{M}_D induces incentive-compatible reporting conditional on contracting at date $\hat{\tau}$ if and only if, for all dates $t \geq \hat{\tau}$ and feasible histories $h^{[\hat{\tau},t]}$ such that the buyer has not yet received the good at date t , and for each $\theta_t \in \{\theta_L, \theta_H\}$, $\nu_t^{\mathcal{M}_D} \left(\theta_t; h^{[\hat{\tau},t]} \right) = u_t^{\mathcal{M}_D} \left(\theta_t; h^{[\hat{\tau},t]}, \iota_{truth}^{[t,\infty)} \right)$.

Proofs of results.

Proof of Lemma 3. For all $\tau, t, t \geq \tau$, let $\theta_L^{[\tau,t]} = \theta^{[\tau,t]}$ such that $\theta^{[\tau,t]}(s) = \theta_L$ for all $s \in [\tau, t]$. Let $\iota_L^{[t,\infty)}$ be the strategy of reporting θ_L regardless of the buyer's values from t onwards.

Let $\mathcal{M}_D = \langle y, q \rangle$ be a mechanism satisfying Constraints 1 and 2. Suppose that \mathcal{M}_D is consistent with a function m , and define, for all τ , all $t \in [\tau, m(\tau)]$,

$$q'_{\tau,t} \left(\theta_L^{[\tau,t]} \theta_H \right) = q_{\tau,t} \left(\theta_L^{[\tau,t]} \theta_H \right) + \Delta_t \times \left(u_t^{\mathcal{M}_D} \left(\theta_H; \theta_L^{[\tau,t]}, \iota_{truth}^{[t,\infty)} \right) - u_t^{\mathcal{M}_D} \left(\theta_H; \theta_L^{[\tau,t]}, \iota_L^{[t,\infty)} \right) \right)$$

and let

$$q'_{\tau,\tau}(\theta_L) = q_{\tau,\tau}(\theta_L) - \Delta_{\tau} \times \left(\int_{\tau}^{m(\tau)} \alpha_L e^{-(\alpha_L+r)(t-\tau)} \left(u_t^{\mathcal{M}_D} \left(\theta_H; \theta_L^{[\tau,t]}, \iota_{truth}^{[t,\infty)} \right) - u_t^{\mathcal{M}_D} \left(\theta_H; \theta_L^{[\tau,t]}, \iota_L^{[t,\infty)} \right) \right) dt \right)$$

where Δ_s is the Dirac delta function indicating a strictly positive payment on date s . Let $q'_{\tau,t} = q_{\tau,t}$ for all other histories of reports.

Let $q' = (q'_{\tau,t}(\cdot))_{\tau \geq 0, t \geq \tau}$ and $y' = y$, and consider the mechanism $\mathcal{M}'_D = \langle y', q' \rangle$. By choice of q' , for

all τ , all $t \in [\tau, m(\tau)]$, $u_t^{\mathcal{M}_D}(\theta_H; \theta_L^{[\tau, t]}, \iota_{truth}^{[t, \infty)}) = u_t^{\mathcal{M}_D}(\theta_H; \theta_L^{[\tau, t]}, \iota_L^{[t, \infty)})$, which is the claim in part (a). Part (b) follows because, for all τ , $u_\tau^{\mathcal{M}'_D}(\theta_L; \emptyset, \iota_{truth}^{[\tau, \infty)}) = u_\tau^{\mathcal{M}_D}(\theta_L; \emptyset, \iota_{truth}^{[\tau, \infty)})$ and $u_\tau^{\mathcal{M}'_D}(\theta_H; \emptyset, \iota_{truth}^{[\tau, \infty)}) \leq u_\tau^{\mathcal{M}_D}(\theta_H; \emptyset, \iota_{truth}^{[\tau, \infty)})$. Hence, if the buyer arrives with a low value, he expects to earn the same rent by participating immediately as under the original mechanism, but he expects to earn weakly less by delaying participation. Part (c) follows because the buyer expects weakly less rents under the modified mechanism, assuming he tells the truth, but the allocation rule remains the same. ■

Proof of Result 4. All arguments are given in the text except for the derivation of the optimal choice of $m(0)$, which is chosen to maximize expected revenue (for the $m(0)$ -conditionally optimal mechanism) as given by

$$\begin{aligned} & \left(\gamma + (1 - \gamma) \frac{\alpha_L}{\alpha_L + r} \right) \theta_H + (1 - \gamma) e^{-(r + \alpha_L)m(0)} \left(\theta_L - \frac{\alpha_L \theta_H}{\alpha_L + r} \right) \\ & - \gamma e^{-(r + \alpha_L + \alpha_H)m(0)} (\theta_H - \theta_L). \end{aligned}$$

■

Proof of Proposition 2. As argued in the main text, the proposed mechanism with $m^\#$ chosen to maximize (21) solves the relaxed problem. The prescribed purchase rule is incentive compatible for the buyer once he has purchased an option, and his payoff from participating with a low or high value is therefore given by $\nu_\tau^{LB}(\theta_L; \emptyset)$ and $\nu_\tau^{LB}(\theta_H; \emptyset)$, as given by (17) and (18) in the main text. What is left to check is that the buyer is willing to purchase an option at the time of arrival.

Suppose the buyer arrives at date τ with value θ_τ and employs an arbitrary participation strategy $\hat{\tau}$. Using Dynkin's formula, and the continuous differentiability of $\nu_\tau^{LB}(\theta_L; \emptyset)$ and $\nu_\tau^{LB}(\theta_H; \emptyset)$ at all points except $\tau = T$, where T is defined by (10), his expected payoff is equal to

$$\begin{aligned} & \mathbb{E}_\tau \left[e^{-r(\hat{\tau}(\tilde{\theta}^{[\tau, \infty)}) - \tau)} \nu_{\hat{\tau}(\tilde{\theta}^{[\tau, \infty)})}^{LB}(\tilde{\theta}_{\hat{\tau}(\tilde{\theta}^{[\tau, \infty)})}; \emptyset) \mid \tilde{\theta}_\tau = \theta_\tau \right] \\ & = \nu_\tau^{LB}(\theta_\tau; \emptyset) + \mathbb{E}_\tau \left[\int_\tau^{\hat{\tau}(\tilde{\theta}^{[\tau, \infty)})} e^{-r(s - \tau)} \begin{pmatrix} 1(\tilde{\theta}_s = \theta_L) \begin{pmatrix} -r\nu_s^{LB}(\theta_L; \emptyset) + \frac{\partial \nu_s^{LB}(\theta_L; \emptyset)}{\partial s} \\ +\alpha_L(\nu_s^{LB}(\theta_H; \emptyset) - \nu_s^{LB}(\theta_L; \emptyset)) \end{pmatrix} \\ +1(\tilde{\theta}_s = \theta_H) \begin{pmatrix} -r\nu_s^{LB}(\theta_H; \emptyset) + \frac{\partial \nu_s^{LB}(\theta_H; \emptyset)}{\partial s} \\ +\alpha_H(\nu_s^{LB}(\theta_L; \emptyset) - \nu_s^{LB}(\theta_H; \emptyset)) \end{pmatrix} \end{pmatrix} ds \mid \tilde{\theta}_\tau = \theta_\tau \right]. \end{aligned}$$

The terms multiplied by the indicator function $1(\tilde{\theta}_s = \theta_L)$ are equal to zero for all $s \neq T$. Moreover, it is easily checked that, for each $s \neq T$,

$$\begin{aligned} & -r\nu_s^{LB}(\theta_H; \emptyset) + \frac{\partial \nu_s^{LB}(\theta_H; \emptyset)}{\partial s} + \alpha_H(\nu_s^{LB}(\theta_L; \emptyset) - \nu_s^{LB}(\theta_H; \emptyset)) \\ & = -(\alpha_L + \alpha_H + r) e^{-(r + \alpha_L + \alpha_H)(m(s) - s)} (\theta_H - \theta_L) m'(s), \end{aligned}$$

which is non-positive provided that $m'(\cdot) \geq 0$. Therefore, the payoff expected from contracting immediately, $\nu_\tau^{LB}(\theta_\tau; \emptyset)$, exceeds that from any alternative strategy involving delay. ■