A Dynamic Theory of Random Price Discounts

Francesc Dilmé* **Daniel Garrett†**

August, 2017
PRELIMINARY VERSION

Abstract

We study profit-maximizing pricing by a seller facing risk-averse buyers who arrive to the market over time. The seller can commit to any stochastic process determining the price at each date. The optimal stochastic price process is a constant "regular" price, with occasional price discounting that occurs at random times. Unlike much of the literature on intertemporal price discrimination, price discounts are unpredictable. Hence, they are sudden rather than gradual, as often observed in retail markets. We also show how the theory can account for other empirical regularities such as the tendency for discount dates to be regularly spaced (or "repulsive") and for sales to occur at the deadline for purchase (such as the end of a season).

JEL classification: D82

Keywords: dynamic pricing, sales, random mechanisms

* **University of Bonn.**

† **Toulouse School of Economics.**

1 Introduction

Durable goods prices at many retailers exhibit a distinct pattern that might seem difficult to square with much of the theory on dynamic pricing. Prices tend to remain constant at the highest level — often termed the "regular price" — apart from when they are occasionally discounted. Evidence for such patterns can be found in studies such as Warner and Barsky (1995), Nakamura and Steinsson (2008) and Février and Wilner (2016).[1]

A key reason these patterns seem difficult to reconcile with much of the theory is as follows. If the sellers in the theoretical models do choose to reduce their prices at certain dates, say to sell to buyers with lower values for the good, then the price discounts are typically *predictable*. Buyers in the market beforehand can anticipate (often perfectly) both the timing and size of the discount. As a consequence, strategic and forward-looking buyers become less willing to purchase at high prices as the date of a price discount approaches. In a range of models with flexible prices, this means that the seller gradually reduces prices as the date with the steepest discount approaches. Stokey (1979), Conlisk, Gerstner and Sobel (1984), Landsberger and Meilijson (1985), Sobel (1991), Board (2008), and Garrett (2016) are but a few examples.

Conversely, given that price discounts are usually discrete rather than gradual, purchases from strategic buyers with perfect foresight would dry up as a price discount approaches. Nonetheless, buyers *do make purchases* immediately before the price of a good is reduced. Perhaps the clearest example is provided by Février and Wilner's (2016) study of a French music retailer in the early 2000s, where sizeable price discounts are common but where purchases are found not to drop off as a discount approaches. A natural interpretation, and a view taken in much of the literature on dynamic demand estimation (see, for instance, Gowrisankaran and Rysman, 2012), is that buyers lack perfect foresight about prices.[2]

In this paper, we propose a novel theory for buyers' apparent failure to foresee the timing of price reductions, suggesting it is a consequence of optimal price setting by sellers that face risk-averse buyers. As in a number of existing papers, we posit forward-looking and strategic buyers with unit demand who arrive to the market over time. Unlike these earlier papers, buyers are averse to risk in the price paid. The seller commits to a stochastic price path, i.e. a stochastic process determining the price at each date.

We show that the seller often finds it optimal to commit to a constant high (or "regular") price punctuated by occasional price discounts whose timing is random from the perspective of buyers. Since high-value buyers cannot predict the timing of price discounts, they are willing to purchase at the high price, right up to the moment that a price discount occurs. These random discounts are a manifestation of the seller's optimal price discrimination policy, and they arise in spite of the

[1] Such patterns are also apparent for "semi-durables". Pesendorfer (2002) studies ketchup prices while Chevalier and Kashyap (2014) study prices of various goods such as peanut butter.

[2] Gowrisankaran and Rysman write simply (p 1183): "We believe that it is more realistic to assume that the consumer has only a limited ability to predict future model attributes."

absence of exogenous uncertainty in the environment (say about cost or demand conditions).

Buyer risk aversion generates two effects that work in tandem to yield optimal random discounts. The first provides the seller with a motive to discount prices. Unlike the well-known benchmark in Stokey (1979), where buyers are risk neutral and have values that are constant over time, intertemporal price discrimination can become optimal for buyers exhibiting diminishing marginal utility of price discounts.[3] Intuitively, selling to low-value buyers at all dates implies that high-value buyers also pay low prices. At these prices, their marginal utility for price reductions is relatively low. Intertemporally discriminating – by delaying the date at which the lowest price is available – allows the seller to reduce the payoff obtained by high-value buyers who purchase earlier. Since their marginal utility of money is small, the reduction in payoff corresponds to a relatively large increase in the price the seller can charge for early purchase.

The second effect provides a motive for randomization. An easy way to understand this is that committing to a deterministic price path, with prices that decline gradually rather than discretely, exposes buyers to pricing risk whose realization depends on the time of arrival to the market. Buyers who arrive not long before the steepest discount receive the lowest prices, while buyers arriving at other times pay higher prices. If the seller instead commits to a constant high price, broken only by occasional randomly timed discounts, then all buyers willing to purchase immediately pay the same price unless arriving at the time of a discount. In relation to the discussion above, that these high-value buyers are willing to purchase at a constant high price is made possible only because the price path is stochastic and the timing of price discounts is impossible to predict.

We begin (Section 3) by studying the model with two buyer values for the good. We start with the case with a single arrival date (Section 3.1), where there always exists an optimal deterministic price path, but also a range of optimal random paths. We then study (Section 3.2) the case where buyers arrive over time at a potentially time-varying rate. The seller's optimal policy is then always a constant price, possibly punctuated by random price discounts or "sales". When random sales are optimal, only high-value buyers purchase at the constant high price, while low-value buyers purchase at sales dates that arrive according to a homogeneous Poisson process. Sales induce purchases by the low-value buyers who have accumulated in the market, but last only an instant in our continuous-time setting with continuous arrivals. Hence, measure zero of high-value buyers pay the sales price and so these buyers are fully insured against pricing risk.

The price path for our two-value model appears to approximate quite well pricing patterns observed in data, where two focal prices (a "high" price and a discounted or sale price) often seem to persist over relatively long periods.[4] Typical price series from Février and Wilner's (2016) study

[3]Other variations relative to Stokey's benchmark include allowing buyer values to change (as studied in the same paper), buyers who are less patient than the seller (as in Landsberger and Meilijson, 1985), and recent work on buyers with multi-dimensional preferences over multiple products (see Rochet and Thanassoulis, 2017).

[4]Presumably for this reason, empirical work has sometimes simplified by considering two prices in parsimonious descriptions of buyers' purchasing problems. Chevalier and Kashyap (2014) suggest that two prices are most relevant for the consumer at any moment: the current list price and the "best price" available over a given time window. Hendel and Nevo's (2013) model of consumers with rational expectations divides prices into "sale" and "non-sale"

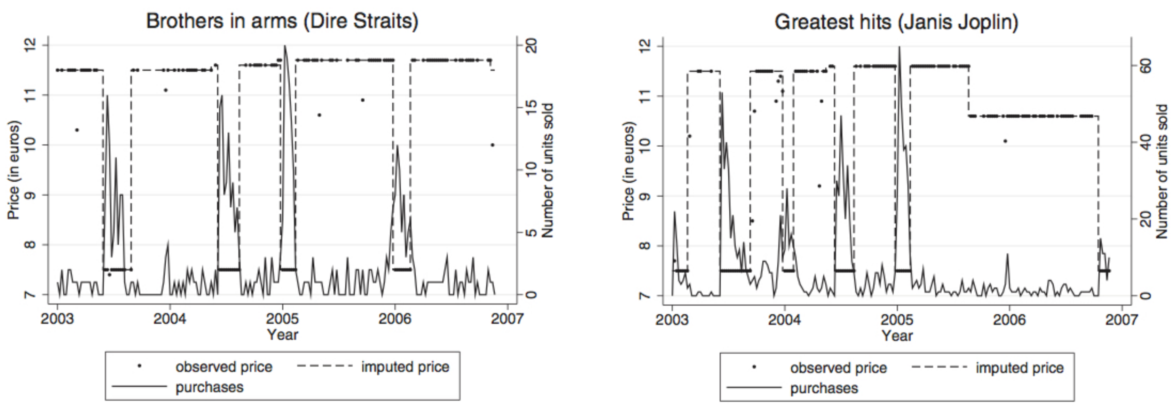

(a) (b)

Figure 1: Illustration of typical price and quantity patterns (Figure 1 from Février and Wilner (2016)).

(for two albums – 'Live At Olympia 1995' by Buckley, and 'Brothers in arms' by Dire Straits) are displayed in Figure 1.[5] Other price series with similar patterns can be found in Figure 1 of Garrett (2016) for prices of food processors, Figures 2, 7 and 8 (among others) in Warner and Barsky (1995) for various consumer durables, and Figure 1 in Pesendorfer (2002) for ketchup, which might be described as "semi-durable".

In terms of our main predictions, there are perhaps five that appear empirically important.

- First, consistent with many of the other theories of dynamic pricing and sales, the quantity purchased is much higher when prices are discounted. This "clearing out" of demand from low-value buyers is apparent in Février and Wilner's (2016) data, where about 35% of the retailer's revenue is generated from purchases at "discounted" prices, while such prices are charged on average only about 15% of the time.[6] Further evidence comes from their observation that "promotions are characterized by [an initial] peak followed by a strongly decreasing pattern of demand".

- Second, our theory predicts that the intensity of purchases at high or "regular" prices can be steady over time, since buyers fail to anticipate the timing of discounts in equilibrium. As noted above, steady demand at the regular price is also a feature of Février and Wilner's data.[7]

values.

[5] We should point out, however, that two-price patterns are not the only possibility that can arise. For instance, in Février and Wilner's data, it is also common for prices of individual albums to be set at more than two different levels at a given store over the horizon of their study.

[6] This pattern is also evident in other data such as in Pesendorfer (2002) and Chevalier and Kashyap (2014).

[7] In contrast, Chevalier and Goolsbee (2009), find evidence that buyers of college textbooks are able to anticipate well textbook revision cycles, with reduced purchases of new books when revision is (likely to be) imminent. This appears consistent with our theory in that it may be more difficult for a seller to adjust or randomize over time to revision, while information about the time since the last revision is immediately available to buyers (see also our fourth point below).

- Third, since buyers in our theory are forward-looking and strategic, high demanded quantities can be sustained at high prices only if price discounts occur infrequently or are not too deep. In line with this, Février and Wilner find that demand at high prices depends on the frequency and magnitude of price discounts (suggesting that consumers are indeed forward-looking, although they do not anticipate the timing of discounts).

- Fourth, while optimal price discounts occur at a homogeneous Poisson rate in our baseline model, things are different if (as is presumably common in practice) buyers fail to observe prices offered before their arrival to the market. In this case, as we show in Section 3.3, a range of processes for discounts is optimal. The key feature of discount dates, which can be viewed as a point process on the real line, is that they are "repulsive" in a sense we define formally below. Roughly, this means that sales dates exhibit "temporal homogeneity" or are "regularly spaced" rather than being "clustered" together. This is true in Février and Wilner's data, but also in some other studies such as Pesendorfer (2002).

- Fifth, price discounts last only a short time, long enough to fill the pent-up demand from low-value buyers. In our model, this demand can be satisfied instantaneously, while in reality this process would be expected to be more gradual (for instance, since buyers may take time to respond to low prices by visiting the store). In Février and Wilner's data, the mean duration of price discounting is five weeks.

It is worth pointing out that our theory can be extended to account for other temporal pricing patterns as well. Section 3.4 shows that, if the seller faces a deadline, then she often discounts the good with a relatively high probability at the deadline. This suggests a rationale for end-of-season sales based on intertemporal price discrimination (rather than, say, inventory concerns).

Given that randomization over prices plays a central role in our theory, it is worth considering what happens if the seller is instead *restricted* to a deterministic path. We characterize (Section 3.5) the seller's optimal choice of this price path and show that it often features cycles of gradually declining prices, familiar from the existing literature mentioned above (e.g., Conlisk, Gerstner and Sobel (1984), Sobel (1991) and Garrett (2016)). For the optimal deterministic price path, high-value buyers pay a different price depending on when they arrive. We show that, depending on the arrival date, they then pay either more or less than under the optimal stochastic price path. This provides a strong sense in which the optimal deterministic price path fails to insure high-value buyers against pricing risk.

We then turn (in Section 4) to study the optimal price path for more than two possible buyer values, again starting with the problem for a buyer who arrives on a fixed and known date (Section 4.1). We show that, under quite mild restrictions on preferences, there exists an optimal price process that is deterministic. We then consider buyers who arrive over time (Section 4.2) and show how the seller can achieve arbitrarily close to the profits per buyer from the case with a fixed arrival date (in the case of two values treated in Section 3 and discussed above, such profits can be attained exactly). This requires the seller to adopt a stochastic price path. Relative to the case with two buyer values, it also requires the seller to discriminate among the different buyer values by randomly dropping prices much as in a (single-bidder) Dutch auction with a random termination time. Unlike the gradual price declines seen elsewhere in the literature (as described above), episodes of price discounting occur at random times and are ephemeral.

Finally, Section 5 discusses several aspects of our theory. Focusing on the fixed-arrival benchmark of Section 4.1, we examine the relationship between optimal dynamic price paths and general direct revelation mechanisms. We show that whether the latter can attain higher profits depends on the nature of risk aversion. We also note the relationship to static revelation mechanisms, recalling (from work such as Salant, 1989) that delayed allocation in dynamic formats is a substitute for uncertain allocation of the good in static formats. Our findings are thus closely related to the literature on static mechanism design for risk-averse agents, such as Matthews (1983) and Maskin and Riley (1984). We also discuss briefly the literature on risk aversion in relation to the degree of risk aversion needed for our theory to be relevant in markets with relatively small pricing risks.

1.1 Review of pricing theories

We now provide some additional details on the relevant literature. First, note that several other papers develop theories in which buyers fail to anticipate and/or wait for price discounts. Pesendorfer (2002) and Chevalier and Kashyap (2014) simply assume that high-value buyers are myopic, so they have no reason to wait. As noted above, Février and Wilner (2016) argue that their data is inconsistent with consumer myopia; buyers expect price discounts but fail to forecast their timing. One possible rationalization of this can be understood from Öry (2017), who studies a seller who lacks the ability to commit to future prices. Buyers in her model are ignorant of their arrival date and cannot observe past prices. In equilibrium, prices are set at a high “regular” price, with discounts that are random from the perspective of consumers. In contrast, our theory endows the seller with full commitment power. We therefore emphasize that holding sales at random times (at least from the perspective of buyers) can be a fully-optimizing choice, including for sellers who find a way to resolve the commitment problem.

Related also are papers where buyers fail to perfectly anticipate futuree prices for other reasons, especially exogenous uncertainty in the environment. For instance, Hörner and Samuelson (2011), Dilmé and Li (2016), Board and Skrzypacz (2016) and Gershkov, Moldovanu and Strack (forthcoming) consider “revenue management” settings where the seller has finite inventory and demand is uncertain. Ortner (2017) studies a setting where the seller’s cost evolves stochastically over time. Unlike these settings, the exogenous environment of our model (demand, preferences and costs) is determinstic. Random pricing emerges as a feature of optimal price discrimination.

It is also worth mentioning that much of the early work on sales, such as Shilony (1977) and

Varian (1980), viewed variation in prices as reflecting randomizations by sellers competing for imperfectly informed buyers. While these frameworks were typically static (an exception is Fershtman and Fishman, 1992), the randomizations over prices are often interpreted as occurring over time. While the models of the early literature often predict smooth price distributions, Heidhues and Kőszegi (2014) suggest that "regular" or "focal" prices may emerge due to consumer loss aversion. They study a monopolist selling to buyers who are loss averse over both money and allocation of the good. Equilibrium prices are dispersed, but exhibit a mass point at the highest or "regular" price. The paper relates this to the same empirical observation we have emphasized above, that firms often set the same price for long periods but also occasionally discount. Unlike our theory for this behavior, Heidhues and Kőszegi's model is static. The role of price discounts is also different; in effect, random price discounts induce buyers to expect to buy the good creating an "attachment effect" that increases buyers' willingness to pay.

Finally, there has been interest in dynamic models with risk-averse buyers in the operations research literature, especially focusing on shortages and rationing as a deliberate price discrimination ploy. Liu and van Ryzin (2008) study a two period model with the possibility to ration in the second period, while Bansal and Maglaras (2009) study a discrete-time model with a finite horizon. These papers posit buyers arriving at a fixed initial date (as in Sections 3.1 and 4.1 of our paper) and do not explore the implications of dynamic arrivals. Among various important differences, note that these papers restrict attention to deterministic prices so that buyers necessarily face uncertainty not about prices but only whether they will be served.

2 Set-up

Environment. We begin by expositing the general model of the paper, specializing to two possible values only in the following section. The seller operates in continuous time that runs over dates $t \in [0, \infty)$. She faces infinitesimal buyers, each with demand for only one unit. Our main interest is in the case where buyers arrive to the market at a possibly time-varying finite rate $\gamma_t > 0$. Both the seller and buyers share a common discount rate $r > 0$ and we normalize demand by setting $\int_0^\infty \gamma_t e^{-rt} dt = 1$ so that the seller's total profits correspond to a per-buyer (weighted) average. Having arrived to the market, buyers who have not purchased exit at a constant Poisson rate $\rho \geq 0$.

Buyers' enjoyment of the good depends on their "type" $\theta_n \in \Theta$, where Θ is a finite set equal to $\{\theta_n | n=1, ..., N\} \subset \mathbb{R}_{++}$, with $\theta_1 > ... > \theta_N > 0$. Throughout, buyers' types θ_n are drawn randomly at the time of arrival to the market. We use $\beta_n > 0$ to denote the probability that the type is θ_n, with $\sum_{n=1}^N \beta_n = 1$.

Let $U = [0, \bar{u}]$ with $\bar{u} > \theta_1$. For each $p \in U$ and type θ_n, $v_n(p)$ denotes the instantaneous utility of a purchase by this type at price p. A buyer's intertemporal payoff is thus described by the Bernoulli utility $e^{-rt} v_n(p_t)$ if the good is purchased at price p_t on date t, while it is equal to zero if no purchase is ever made (or he exits the market). For all n, $v_n(\cdot)$ is a strictly decreasing, strictly concave, and twice continuously-differentiable function.[8] We normalize by setting $v_n(\theta_n) = 0$ for each n, i.e. a buyer's type corresponds to his value for the good. A natural example is $v_n(p) = \phi(\theta_n - p)$, with ϕ a strictly increasing, strictly concave, and twice-continuously differentiable function satisfying $\phi(0) = 0$. Another is $v_n(p) = \kappa(\theta_n) - \kappa(p)$, where κ is a strictly increasing, strictly convex, and twice-continuously differentiable function. Maskin and Riley (1984) describe further possibilities. Throughout, we will write $p_n(x)$ to denote the unique price charged to type θ_n that yields him an instantaneous utility $x \in [v_n(0), v_n(\bar{u})]$. By our normalization, $p_n(0) = \theta_n$ for all $\theta_n \in \Theta$.

The simplest interpretation of $v_n(p)$ is as a utility enjoyed by type θ_n when purchasing at price p. In this case, the good is best thought of as an "experience" that the buyer will only enjoy once, or a short-lived good like a fashion item that the buyer will only wear once.[9] However, v_n could also reflect discounted future flows. For instance, if $v_n(p) = \kappa(\theta_n) - \kappa(p)$, it is natural to view $\kappa(\theta_n)$ as the net present value of consumption utility over the (possibly infinite) future and $\kappa(p)$ as the lifetime disutility associated with reduced spending on other goods (assuming sufficient access to financial markets). Note, however, that our model does not, in general, easily extend to permit financial arrangements that take place *before* the purchase of the good (such as measures to hedge risk regarding the date or price of purchase).[10]

Finally, the seller is a risk-neutral profit maximizer who faces a constant cost $c \in [0, \theta_N)$ per unit and no capacity constraints.[11] The seller's payoff is determined by integrating over profits from each buyer, which are given by $e^{-rt}(p_t - c)$ in case of purchase at price p_t on date t.

Stochastic price processes. A stochastic price process is a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathcal{P})$ (with Ω a sample space, $\mathcal{F}$ a sigma algebra, $(\mathcal{F}_t)_{t\geq 0}$ a filtration of $\mathcal{F}$, and $\mathcal{P}$ a probability measure), and an adapted and progressively measurable process $P : [0, \infty) \times \Omega \to U$. Throughout we require $(\mathcal{F}_t)_{t\geq 0}$ to be the natural filtration generated by P. The price $P_t(\omega)$ is the amount charged in case of purchase at date t in outcome $\omega \in \Omega$. Where no ambiguity is created, dependence on ω will be suppressed. At any date t, a buyer in the market at that date knows the realization of prices $(p_s)_{s\in[0,t]}$ up to date t and correctly anticipates the continuation process by updating based on knowledge of both the random variable P and $(p_s)_{s\in[0,t]}$ (below, we will consider also the possibility that the buyer does not observe prices prior to arrival to the market).

Given that buyers may be indifferent between purchasing at a known price and waiting and

[8] While these restrictions will be enough for all results in Section 3, where there are only two types (i.e., $N = 2$), Section 4 introduces additional restrictions in the spirit of Maskin and Riley (1984).

[9] Fashion items are increasingly purchased to wear only a few times. See Tibbetts, Graham. 2008 "Primark effect' leads to throwaway fashion turning up in landfill.' *The Telegraph*, 25 November. http://www.telegraph.co.uk/news/uknews/3516158/Primark-effect-lead-to-throwaway-fashion-turning-up-in-landfill.html.

[10] We leave for future work the possibility that buyers privately save in anticipation of pricing risks.

[11] Note that none of our qualitative conclusions will hinge on either c or ρ being strictly positive, although it is reasonable to expect that both may be in many applications of interest. Our formulation is chosen to permit that buyers' desire to purchase at high prices, rather than wait for "sales", is related *both* to diminishing marginal utility of money and the risk of dropping out of the market. Hence, buyers' willingness to purchase at the high "regular" prices in Figure 1, for instance, would be consistent with moderate levels of risk aversion and discount rates provided that ρ is sufficiently large.

purchasing at a later date, we complete the definition of a stochastic price process by requiring that the seller specify the purchase date for all buyers. For any buyer of type θ_n arriving at date τ, this is a stopping time $\tilde{t}_n^\tau \in [\tau,\infty) \cup \{+\infty\}$, measurable with respect to the filtration generated by P, $(\mathcal{F}_t)_{t\geq 0}$.[12]

The stopping times $\tilde{t}_n^\tau$ must be incentive compatible for buyers in the sense that

$$\tilde{t}_n^\tau \in \arg\max_{\tilde{t}\in\mathcal{T}^\tau} \mathbb{E}\big[e^{-(r+\rho)(\tilde{t}-\tau)} v_n(P_{\tilde{t}})\big] \quad \text{for all } n \in N \text{ and all } \tau \geq 0,$$

where $\mathcal{T}^\tau$ is the set of stopping times taking values no less than τ. It is natural and without loss of optimality for the seller to restrict attention to stopping times such that, for all types θ_n, and for any τ', τ'' with $\tau' < \tau''$, $\tilde{t}_n^{\tau'} \geq \tau''$ implies $\tilde{t}_n^{\tau'} = \tilde{t}_n^{\tau''}$. That is, all buyers with the same type who are in the market and have not yet purchased do so at the same time. For the case of a fixed arrival date $\tau = 0$, as considered in Sections 3.1 and 4.1, we omit the superscript and write $\tilde{t}_n$.

3 Two Types

In this section, we restrict attention to two types and refer to type θ_1 as the "high type" and θ_2 as the "low type". We let $\beta = \beta_1 \in (0,1)$ be the probability of the high type θ_1.

3.1 Fixed arrival date

We begin by considering the case where measure one of buyers arrive at a fixed date, zero. The seller's expected discounted profits are

$$\beta\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}\left(P_{\tilde{t}_1} - c\right)\big] + (1-\beta)\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_2}(P_{\tilde{t}_2} - c)\big],$$

which she maximizes by choice of the price process P and incentive-compatible stopping times $\big(\tilde{t}_1, \tilde{t}_2\big)$. We begin with a few simple observations about optimal price paths.

Lemma 1 *Suppose there are two buyer types and that buyers are known to arrive at date zero. If P is an optimal price process, then $P_t \geq \theta_2$ for all t almost surely. Also, if it induces stopping times $\tilde{t}_1$ and $\tilde{t}_2$, then $\Pr(\tilde{t}_1 = 0) = 1$ and $\Pr(\tilde{t}_2 = 0) \in \{0,1\}$.*

The reasons for these claims are as follows. First, charging a price below θ_2 with positive probability gives a low-type buyer the opportunity to earn a rent. This may need to be passed in turn also to high-type buyers; say, to induce early purchase. This can be avoided by never lowering prices below θ_2. Second, akin to the usual "efficiency at the top" result in mechanism design, inducing high types to purchase immediately permits the seller to leave these buyers a fixed amount of rent (needed to deter mimicry of the purchase time for low types) while obtaining the

[12] A value equal to $+\infty$ corresponds to a decision not to purchase.

greatest possible profit. Third, if $\Pr(\tilde{t}_2 = 0) \in (0,1)$, the seller must reduce the price to θ_2 with positive probability at date zero, implying that high types face pricing risk that can be avoided if $\Pr(\tilde{t}_2 = 0) = 0$. In other words, profits are higher if the seller asks high types to purchase first (at a certain price, thus insuring these buyers against pricing risk) and delays purchase by low types to a later date with probability one.

Now note that, given high types can delay until low types purchase, they can obtain a payoff at least $\psi v_1(\theta_2)$, where $\psi = \mathbb{E}\left[e^{-(r+\rho)\tilde{t}_2}\right] \in [0,1]$ is the expected discount factor associated with type θ_2's purchase. Hence, the highest price the seller can charge high types at date zero is $p_1(\psi v_1(\theta_2))$, which is above θ_2 in case $\psi < 1$ and equal to θ_2 otherwise. Taking this as the date-zero price, the seller's expected profits can be written

$$\beta \left(p_1\left(\psi v_1(\theta_2)\right) - c\right) + (1-\beta)\psi\left(\theta_2 - c\right). \tag{1}$$

These expected profits are strictly concave in the expected discount factor $\psi = \mathbb{E}\left[e^{-(r+\rho)\tilde{t}_2}\right]$. Choosing ψ to maximize (1) over $[0,1]$ yields the following result.

Proposition 1 *Suppose there are two buyer types and that buyers arrive at a fixed date, zero. Then there exist $\underline{\beta}$ and $\bar{\beta}$, with $0 < \underline{\beta} < \bar{\beta} < 1$, such that, for any optimal price process, purchases can be described a.s. as follows:*

1. *If $\beta \leq \underline{\beta}$ then both types purchase at price θ_2 at date zero.*
2. *If $\beta \in \left(\underline{\beta}, \bar{\beta}\right)$ then the high type purchases at date zero at price $p_1(\psi^* v_1(\theta_2))$, while the low type purchases at a random time $\tilde{t}_2 > 0$ satisfying $\mathbb{E}[e^{-(r+\rho)\tilde{t}_2}] = \psi^*$, where $\psi^* \in (0,1)$ is the unique maximizer of (1).*
3. *If $\beta \geq \bar{\beta}$ then the high type purchases at price θ_1 while the low type never purchases.*

There is a large multiplicity of optimal price processes. In Cases 1 and 3, one possibility is a constant price. In Case 2, the optimal date-zero price is $p_1^* = p_1(\psi^* v_1(\theta_2))$. The price may then remain at this level except at the random time $\tilde{t}_2$, when it drops to θ_2 and the low type purchases. There is a wide range of possibilities for $\tilde{t}_2$. It may be degenerate at $\frac{-\log(\psi^*)}{r+\rho}$, or it could be exponentially distributed with parameter $\lambda^* = \frac{\psi^*(r+\rho)}{1-\psi^*}$.

We can conclude from Proposition 1 that intertemporal price discrimination occurs whenever $\beta \in \left(\underline{\beta}, \bar{\beta}\right)$. This contrasts with Stokey (1979) and Conlisk, Gerstner and Sobel (1984), who observed that, for risk-neutral buyers with constant values for the good, intertemporal price discrimination is not profitable.

One way to understand our result relative to the risk-neutral case is as follows. Suppose that $\beta = \frac{\theta_2}{\theta_1}$ and $c = 0$. In this case, the seller obtains the same profits whether setting a constant price equal to θ_1 or setting a constant price equal to θ_2 (and asking buyers to purchase at date zero whenever their reservation value exceeds the price). Indeed, these are the optimal profits when buyers are risk neutral. If, instead, buyers are risk averse, the seller can obtain higher profits setting a non-constant price. For example, the seller can charge some price p_0 at time 0, a high price $\bar{u}$ in $(0, \varepsilon)$ for some $\varepsilon > 0$, and charge a constant price afterwards. The price from ε onwards is determined randomly — it is θ_1 with some probability $\chi \in (0,1)$ and θ_2 otherwise. Choosing p_0 to make high types indifferent between purchasing at time 0 and waiting until time ε, the seller can ask high types to purchase at time 0 and low types to purchase at time ε whenever the price is θ_2. Hence, if ε is small, the seller obtains almost the same discounted revenue from the low types, and a strictly higher revenue from from high types since, given that they are risk averse, p_0 is bounded away from $\chi\theta_1 + (1-\chi)\theta_2$.

3.2 Dynamic Arrivals

Now we return to the setting of interest where buyers arrive at rate $\gamma_t > 0$ for each $t \geq 0$. The seller's problem is to choose a price process P and incentive-compatible stopping times $\left(\tilde{t}_1^\tau, \tilde{t}_2^\tau\right)$ for each arrival date τ (as described in the model set-up) to maximize expected profits

$$\int_0^\infty \gamma_\tau e^{-r\tau} \mathbb{E}\left[\beta e^{-(r+\rho)(\tilde{t}_1^\tau - \tau)}(P_{\tilde{t}_1^\tau} - c) + (1-\beta)e^{-(r+\rho)(\tilde{t}_2^\tau - \tau)}(P_{\tilde{t}_2^\tau} - c)\right] d\tau. \tag{2}$$

The next result characterizes optimal price processes as processes which feature a regular high price p_1^* and occasional random "sales", i.e. times when the price drops to θ_2.

Proposition 2 *Suppose that buyers arrive over time and that $\beta \in \left(\underline{\beta}, \bar{\beta}\right)$, the interval identified in Proposition 1. Let $\psi^* \in (0,1)$ be the unique maximizer of (1). Then, an optimal price process exists, with expected profits equal to those obtained in the fixed-arrival case of the previous subsection. Any optimal price process P, together with optimally specified incentive-compatible stopping times $\left(\tilde{t}_1^\tau, \tilde{t}_2^\tau\right)$, satisfies the following.*

1. *For almost all arrival dates τ, high types purchase immediately paying price p_1^* with probability one.*

2. *All low types arriving at $\tau > 0$ purchase at times $\tilde{t}_2^\tau$ that are exponentially distributed with parameter $\lambda^* = \frac{\psi^*(r+\rho)}{1-\psi^*}$, independently of past prices. The probability that a low type arriving at $\tau > 0$ purchases at a price other than θ_2 is zero.*

An optimal policy for the seller is thus a constant "regular" price p_1^*, punctuated by sales at price θ_2 that occur at a Poisson rate with parameter $\lambda^* = \frac{\psi^*(r+\rho)}{1-\psi^*}$. For any arrival date, the buyer then believes the date of the next sale to be exponentially distributed with parameter λ^*. Since the exponential distribution is memoryless, buyers' beliefs on the timing of the next sale are independent of time and the history of past prices. Expected profits per buyer are then equal to those when the price process could be conditioned on the buyer's arrival time as in Section 3.1. Moreover, such a policy is "essentially" the only optimal policy for the seller, in terms of purchase times and prices paid, in a sense made clear in the proposition.

The logic for the optimality of Poisson sales is remarkably similar to that in Conlisk, Gerstner and Sobel (1984). They pointed out that Stokey's (1979) observation about the sub-optimality of intertemporal price discrimination with a known arrival date extends straightforwardly to settings with dynamic arrivals. Indeed, for risk-neutral buyers, an optimal price process is a constant price, which induces any buyer to purchase upon arrival at the static monopoly price, replicating the outcome when the arrival date is fixed (as in Stokey).

The fact that sales arrive at a Poisson rate provides a formal sense in which buyers find them difficult to predict. As noted, the seller's past prices are uninformative about the timing of future sales. Related, buyers have as little information as possible about the time of the next sale in the sense of (Shannon) entropy, as the exponential distribution with parameter λ^* is the maximum entropy distribution among continuous distributions with mean $1/\lambda^*$ and support on the positive reals.

Comparative statics can provide some additional insights into the main forces in the model.

Corollary 1 *Suppose that buyers arrive over time ($\gamma_t > 0$ for all t), and that $\beta \in \left(\underline{\beta}, \bar{\beta}\right)$, the interval identified in Proposition 1. The Poisson rate for sales $\lambda^* = \frac{(r+\rho)\psi^*}{1-\psi^*}$ is (a) increasing in the discount rate r and exit rate ρ and (b) decreasing in the proportion of high types β.*

Part (a) of the result follows simply because the optimal value of the expected discount factor, ψ^*, is invariant to r and ρ. Part (b) follows because, with more high types, the seller favors reducing the rents left to these types, which requires fewer sales.

It is also worth noting that the frequency of sales, λ^*, can increase or decrease with the degree of risk aversion of the buyers. To see this, suppose that $v_n(p) = \kappa(\theta_n) - \kappa(p)$ for the CARA function $\kappa(x) = e^{\alpha x}$, for some $\alpha > 0$. Then

$$\psi^* = \frac{1}{1 - e^{-\alpha(\theta_1 - \theta_2)}} - \frac{\beta}{\alpha(1-\beta)(\theta_2 - c)}$$

whenever $\psi^* \in (0,1)$, which may be non-monotone with respect to α. To see this notice that, for a fixed $\alpha > 0$, a constant price equal to θ_2 is optimal for β sufficiently small. This corresponds to values ψ^* equal to one. However, taking α large, we necessarily have $\psi^* \in (0,1)$, and it approaches 1 as $\alpha \to +\infty$.

3.3 Unobserved past prices and patterns of discounts

As has sometimes been noted (see Öry (2017), but also Seiler (2013) and Février and Wilner (2016)), past prices are often not available to a buyer who newly arrives to the market. We now suppose that a buyer arriving to the market at any date τ has no information on prices before τ. We study the implications of this possibility for the set of profit-maximizing price processes P.

It is clear that the seller cannot raise profits above those attained by holding sales at a constant Poisson rate, as described in the previous subsection, since such a price process maximizes expected profits for every cohort. Nonetheless, various processes for the sales dates become optimal. To see this, suppose that $\beta \in (\underline{\beta}, \bar{\beta})$, the interval identified in Proposition 1, so that $\psi^* \in (0,1)$. Consider, for simplicity, price processes P where the price is set at p_1^* except at sale dates (when the price is θ_2), and suppose these sale dates are determined by a simple point process. Suppose, in addition, that the stopping times for high types specify immediate purchase, while for low types they specify purchase at sales. As usual, for an optimal price process, these stopping times are required to be incentive compatible (note that, here, the available stopping times for buyers are those adapted to the information generated through prices observed only since arrival to the market).

If P and stopping times $(\tilde{t}_1^\tau, \tilde{t}_2^\tau)$ are profit-maximizing (and hence the stopping times are incentive compatible) then, for almost all dates τ,[13]

$$\mathbb{E}\left[e^{-(r+\rho)(\tilde{t}_2^\tau - \tau)}\right] = \mathbb{E}\left[e^{-(r+\rho)(\tilde{t}_2^\tau - \tau)} \middle| \tilde{t}_2^\tau > \tau\right] = \psi^*. \tag{3}$$

For such τ, profits expected from low types are the same as in the case for a fixed arrival date in Section 3.1. Provided it is also incentive compatible for high types to purchase immediately upon arrival, then P is an optimal price process.

When (3) holds, a high type who arrives to the market at date τ and has no information about past prices obtains the same payoff purchasing immediately at arrival (paying price p_1^* almost surely) or, alternatively, waiting and purchasing at the next sale. A buyer who delays purchase, however, *is not restricted to purchasing at sales*, so the condition (3) is not sufficient to guarantee immediate purchase. We now provide a condition that determines the collection of optimal point processes for sales dates.

Proposition 3 *Suppose that buyers arrive over time, but observe prices only since arrival to the market. Consider a price process such that a "regular price" p_1^* is posted except at "sales" when the price is θ_2, and suppose that low types purchase at sales. For any τ such that (3) holds, it is incentive compatible for a high type arriving at τ to purchase immediately with probability one if and only if, for all $s > \tau$ such that $\tilde{t}_2^\tau > s$ with positive probability,*

$$\mathbb{E}\left[e^{-(r+\rho)(\tilde{t}_2^\tau - s)} \middle| \tilde{t}_2^\tau > s\right] \geq \psi^*. \tag{4}$$

One way to understand Condition (4) is that the absence of a sale since a buyer's arrival (i.e., $\tilde{t}_2^\tau > s$) is "good news" in the sense that the buyer then expects a sale relatively soon. In the language of point processes, if the condition is satisfied for all τ, one might say that the process for sale dates is "repulsive". Conversely, when the condition fails, the absence of a sale since arrival can be "bad news"(i.e., the buyer expects to wait a long time for a sale), a property associated

[13] Stopping times $\tilde{t}_n^\tau$ are now adapted to the information generated since date τ.

with "clustering". For the (time-homogeneous) Poisson process, the absence of a recent sale is not informative about the timing of the next, and such processes are usually described as exhibiting neither repulsion nor clustering.[14]

To help understand the importance of Condition (4), consider a point process that is determined as follows. With probability $\chi \in (0,1)$, sales occur at a Poisson rate, with the Poisson rate given by $\frac{\chi\psi^*(r+\rho)}{1-\chi\psi^*}$. With probability $1-\chi$, the good is never discounted. Then, while Condition (3) holds, Condition (4) fails for any τ and any $s>\tau$. Intuitively, a buyer arriving at date τ becomes pessimistic about the chances of the good ever being discounted upon delaying until $s>\tau$ without observing a sale. Deviating by delaying to date s (unless there is a sale) is profitable because the buyer (given the information available at date τ) considers it relatively likely that a sale occurs before s, but relatively unlikely that a sale occurs after s if none has occurred by this date. This example serves as a caution that, in empirical applications where some goods are occasionally discounted but not others, Conditions (3) and (4) must be evaluated in light of the information consumers are likely to have about which goods are subject to price discounts.

I order to compare to the empirical literature (see below) that focuses on hazard rate of inter-arrival distributions, consider stationary renewal processes such that inter-arrival times are distributed according to some distribution F. The distribution F is taken to be an absolutely continuous c.d.f. with support on $[0,\infty)$ and finite mean $\mu=\int_0^\infty x dF(x)$. As a consequence of the Renewal Theorem (Feller, 1968, Chapter XI), there exists a unique distribution G for the initial sales date $\tilde{t}_2^0$ that guarantees stationarity. By stationarity, we mean that, for any arrival date τ, the time to the next sale $\tilde{t}_2^\tau-\tau$ (as perceived by a buyer who has no information on prices before τ) is independent of τ. We have that (Feller, 1968, Chapter XI), for any τ, any $\xi\geq 0$,

$$\Pr\left(\tau\leq\tilde{t}_2^\tau\leq\tau+\xi|\tilde{t}_2^\tau\geq\tau\right)=G\left(\xi\right)=\frac{1}{\mu}\int_0^\xi\left(1-F\left(y\right)\right)dy.$$

We now consider properties of F such that the stationary renewal processes defined above are optimal. To do so, we consider standard properties of distribution functions (see, for instance, Hollander and Proschan, 1984).

Definition 1 *Fix a distribution F. We say that it satisfies*

1. *"decreasing mean remaining life" (DMRL) if $\int_t^\infty \frac{y-t}{1-F(t)}dF(y)$ is weakly decreasing in t.*
2. *"strictly increasing mean remaining life" (SIMRL) if $\int_t^\infty \frac{y-t}{1-F(t)}dF(y)$ is strictly increasing in t.*
3. *"new better than used in expectation" (NBUE) if $\int_t^\infty \frac{y-t}{1-F(t)}dF(y)\leq\int_0^\infty ydF(y)$ for all $t>0$.*

It will be important to note that DMLR implies NBUE. Note also that DMLR is implied if

[14] Note that various definitions of repulsion and clustering are considered in the literature on point processes, often depending on the application of interest.

F has an increasing hazard rate (see Hollander and Proschan, 1984). These notions are related to Proposition 3 as follows.

Corollary 2 *Consider a stationary renewal process as defined above, with interarrival times distributed according to F (absolutely continuous and with finite mean, as described above). If F satisfies DMLR, then Condition (4) holds, for all τ and all $s > \tau$, whenever Condition (3) does (i.e., whenever $\int_0^\infty e^{-y(r+\rho)} dG(y) = \psi^*$). If F satisfies SIMLR, then Condition (4) fails, for all τ and all $s > \tau$, whenever Condition (3) holds. If F fails to satisfy NBUE, then there exist parameters such that Condition (3) holds, yet for every τ there is some $s(\tau)$ such that Condition (4) fails.*

The result states that, for stationary renewal processes, DMLR is an adequate notion of repulsion to ensure optimality. SIMLR is always inconsistent with optimality. For distributions F that satisfy neither property, we offer an alternative "partial converse" to the first claim: If F fails NBUE, then the process for sales may fail to be "sufficiently repulsive", at least for some parameters of the problem. The corollary thus relates Condition (4) to standard concepts in survival analysis.

These observations hold particular interest for an empirical literature that studies the timing of price discounts across goods with varying degrees of durability (recall that some products such as ketchup appear to be "semi-durable" in that buyers can engage in some intertemporal substitution). These studies typically investigate how the probability of a price discount in a given period depends on the time since the most recent discount; i.e., they consider the hazard rate for discounting of a "regular" price. Findings are somewhat mixed: for instance, Berck et al. (2008) find some evidence that the hazard rate can decrease over time for frozen and refrigerated orange juice, while other studies find evidence of an increasing hazard rate — see Pesendorfer (2002) for ketchup, Février and Wilner (2016) for compact discs, and Lan, Lloyd and Morgan (2016) for supermarket food products. On the other hand, as emphasized by Nakamura and Steinsson (2008), unobserved heterogeneity in discounting policies across stores, products or time intervals can bias findings in the direction of decreasing hazard rates; Lan, Lloyd and Morgan argue that accounting for such heterogeneity is important for their finding of an increasing hazard rate. This observation is intimately related to our previous example, where the good is discounted at a Poisson rate with probability $\chi \in (0, 1)$ and not discounted otherwise. In that case, "clustering" resulted from buyers being unable to identify whether the good would be discounted or not.

In light of these possible applications, our analysis here makes several contributions. First, we have suggested a condition based on buyer incentive compatibility that could be readily tested in empirical settings. Recall that (for stationary renewal processes) this condition is implied by an increasing hazard rate for sales, and so we expect it to hold in those studies that have found an increasing hazard rate as mentioned above. However, Corollary 2 indicates that the weaker property of DMLR is sufficient (although it is also not necessary), and so our condition permits some violation of hazard rate monotonicity. Our observations also offer an interesting theoretical lens under which to consider discounting patterns. Notably, it is different from the other leading theory for "repulsive" discounting processes, which is based on sequential optimality of the seller (i.e., the idea that the seller waits for an accumulation of low types in the market since the last discounting date; see Pesendorfer, 2002). Second, it explains why price discounts may appear quite "predictable" to a researcher with data on price discounting, though not by buyers who lack this information. It also suggests that a seller with the opportunity to commit to an information policy providing information to consumers about the timing of previous discounts would elect not to reveal such information (in Pesendorfer's theory, which features myopic "high types", the seller would not lose from making such information available).[15] Indeed, it would appear that retailers rarely provide detailed scheduling information for price discounts.[16] Third, our theory permits a prediction that, as consumers become better able to access information about historical prices (or price predictions based on these prices), discounts should become less predictable; i.e., discounting should be at a time-homogeneous Poisson rate (as predicted by our theory when buyers have full information on past price realizations).

3.4 Deadlines

Of independent interest is the case where there is instead a known deadline. In particular, suppose that buyers arrive over an interval $[0, T]$, and $T > 0$ is the last date on which the good can be sold. As before, the seller and buyers discount time at the same rate r, and buyers arrive to the market at some finite rate γ_t for $t \in [0, T]$, with $\int_0^T \gamma_t e^{-rt} dt = 1$, and exit stochastically at rate ρ. For concreteness, buyers observe all previous prices, as assumed in the model set-up. Our results turn out to extend straightforwardly to this case, with the key difference that a positive probability mass is needed on a sale at the final date T to attain optimality.

Proposition 4 *Suppose that buyers arrive over a bounded interval* $[0, T]$*, and that* $\beta \in (\underline{\beta}, \bar{\beta})$*, the interval identified in Proposition 1. Then an optimal stochastic price process involves setting a constant price* p_1^* *except at either (i) random times in* $[0, T)$ *drawn according to a Poisson distribution with parameter* $\lambda^* = \frac{(r+\rho)\psi^*}{1-\psi^*}$*, or (ii) date* T *with probability* ψ^**. At such times, the price charged is* θ_2*.*

Proposition 4 shows that optimal intertemporal price discrimination calls for a mass point of the sales process at the deadline T.[17] To understand this, note that cohorts arriving close to the deadline have only a limited horizon over which to be offered a sale. They must therefore expect a sale to occur soon (by the deadline) or not at all. Our theory thus provides an alternative rationale for price reductions on the deadline relative to the one suggested in much of the revenue management literature

[15] In relation to the "Bayesian persuasion" literature that has received recent attention since Kamenica and Gentzkow (2011), the optimal information policy is a corner solution.

[16] The website https://isthereanydeal.com does provide historical information on price discounts to buyers of some computer games, but it notes that "we are not allowed to show you historical data for Amazon", perhaps suggesting Amazon would prefer the information not to be readily available.

[17] Its proof shows that any optimal process is essentially equivalent (in terms of payments and purchase dates) to the process described in the proposition.

that emphasizes the role of fixed but expiring inventory (see Hörner and Samuelson, 2011, Board and Skrzypacz, 2016, and Gershkov, Moldovanu and Strack, forthcoming, for recent examples). A key distinguishing prediction of our theory is that price reductions can occur with high probability at the deadline even absent uncertainty in the environment (especially regarding the level of demand).

3.5 Optimal deterministic price path

We now compare the optimal stochastic price process characterized above to the optimal choice of deterministic price path. We simplify by assuming that $\gamma_t = \gamma > 0$ for all t and restrict attention to parameters such that $\beta \in (\underline{\beta}, \bar{\beta})$.[18]

Without loss of optimality, prices are never below θ_2 and so low types purchase whenever the price θ_2 is charged; we again term such dates "sales". Let S denote the set of sales dates. If the time remaining before a sale is Δ, the highest price that can be charged to high types is

$$p_1^d(\Delta) \equiv p_1\big(e^{-(r+\rho)\Delta} v_1(\theta_2)\big).$$

Hence, if the seller chooses to hold sales at intervals of length z, expected discounted profits $\Pi(z)$ solve the following equation

$$\Pi(z) = \int_0^z \gamma e^{-r\tau} \left(\beta(p_1^d(z-\tau) - c) + (1-\beta)\, e^{-(r+\rho)(z-\tau)}\,(\theta_2 - c)\right) d\tau + e^{-rz}\Pi(z),$$

and $\Pi(0) = \frac{\gamma}{r}(\theta_2 - c)$.

The seller's problem thus has a clear recursive structure, since, for any sales date in S, prices offered after that date do not affect the decision problem of buyers arriving before it.[19] Hence, the seller's optimal profits can be stated recursively as

$$\sup_{z>0} \frac{\int_0^z \gamma e^{-r\tau}\left(\beta\left(p_1^d(z-\tau) - c\right) + (1-\beta)\, e^{-(r+\rho)(z-\tau)}\,(\theta_2 - c)\right) d\tau}{1 - e^{-rz}}. \tag{5}$$

We can characterize the solution to the seller's problem as follows.

Proposition 5 *Let $\beta \in (\underline{\beta}, \bar{\beta})$. Then the optimal deterministic price path is described by one of two possibilities:*

1. *There are regular sales, length $z^* > 0$ apart, with z^* uniquely determined; or*

2. *There are no sales at strictly positive dates;*[20] *an optimal deterministic price path is a constant price equal to θ_1.*

[18] Otherwise, there is an optimal stochastic price process that is deterministic, involving a constant price at either θ_1 or θ_2.

[19] Given that the seller never offers a price below θ_2, any buyer arriving before z purchases either at date z or earlier.

[20] Since measure zero of buyers arrive at date zero, the price at this date is indeterminate.

In the first case $e^{-(r+\rho)z^*} < \psi^*$, and hence high types are necessarily worse off than under the optimal stochastic price process for at least some arrival times.

Proposition 5 determines that, given $\beta \in (\underline{\beta}, \bar{\beta})$, the optimal deterministic price path is either a constant price θ_1, or it features cycling prices. Prices are highest at the beginning of a cycle and fall gradually to the price θ_2 at which low types purchase. High types arriving at the beginning of the cycle are worse off than under the optimal stochastic price process. Buyers arriving near the end of the cycle are of course better off. This provides a strong sense in which deterministic price paths fail to insure buyers against pricing risk associated with their arrival time to the market. One way to understand the result is that deterministic price paths reduce the seller's ability to tailor the expected discount factor ψ to buyers' arrival times. Given the restriction, the seller aims to do well "on average" across different arrival times, which requires that ψ can be smaller or greater than ψ^* depending on the arrival date.

4 Multiple Types

We now consider the case with more than two buyer types, imposing the following additional restrictions on buyer preferences.

Assumption A *The following additional assumptions are made on v_n:*

A1 Higher types are "more eager". For any $n = 1, ..., N-2$ and $p < \theta_{n+1}$, $\frac{-v_n'(p)}{v_n(p)} < \frac{-v_{n+1}'(p)}{v_{n+1}(p)}$.

A2 Higher types are less risk averse. For any $n = 1, ..., N-2$ and $p \in U$, $\frac{v_n''(p)}{v_n'(p)} \leq \frac{v_{n+1}''(p)}{v_{n+1}'(p)}$.

Note that Assumption A implies a joint restriction on the preferences of types $\{\theta_1, \ldots, \theta_{N-1}\}$, so there are no additional restrictions in the two-type case (our results below nest the case for two types). Relative to the types $\{\theta_1, \ldots, \theta_{N-1}\}$, the restrictions are close in spirit to the ones made by Maskin and Riley (1984), who consider continuous types. Assumption A1 is related to Maskin and Riley's Assumption B1, and ensures a higher type θ_n benefits proportionally less from price reductions; hence, higher types will be more "eager" to purchase early rather than delay and purchase at a reduced price. Assumption A2 is related to Maskin and Riley's Assumption B5, and requires risk aversion to be (weakly) decreasing in the type. The condition is perhaps natural in many settings, since buyers who gain more from purchasing (a higher θ_n) might be expected to be less concerned about price risks.

4.1 Fixed arrival date

We again begin with the case where measure one of buyers arrive at a fixed date, zero. A common starting point in the analysis of dynamic pricing models is to verify the so-called "skimming property"; i.e., that higher types purchase earlier. Interestingly, this property is not guaranteed across arbitrary stochastic price processes, even under the additional restrictions of Assumption A. The reason is that, when higher types are sufficiently less risk averse, these types may be willing to wait to purchase at risky future prices where lower types would prefer to pay a certain price rather than waiting. Conversely, absent such pricing risks, the skimming property does hold. We summarize these observations in the following result (which naturally extends to cases where buyers arrive over time).

Lemma 2 *(skimming property) Suppose that Assumption A is satisfied and that buyers arrive at a fixed date, zero. Consider any price process P such that the probability a price is offered below θ_N is zero, and let $(\tilde{t}_n)_{n=1}^N$ be some incentive-compatible stopping times.*[21] *Then:*

1. *The skimming property need not be satisfied in general, that is, there exist settings and price processes such that* $\Pr(\tilde{t}_n > \tilde{t}_{n+1}) > 0$ *for some* $n \in \{1, ..., N-1\}$.
2. *Suppose each type θ_n pays a sure price when purchasing, that is, there exists, for each θ_n, a price $p_n \in U$ such that* $\Pr\big(\big(P_{\tilde{t}_n} \neq p_n^*\big) \wedge \big(\tilde{t}_n < +\infty\big)\big) = 0$. *Then the skimming property holds:* $\Pr(\tilde{t}_n \leq \tilde{t}_{n+1}) = 1$ *for all* $n = 1, ..., N-1$.

The key assumption that guarantees the skimming property absent pricing risk is Assumption A1. Intuitively, it ensures that higher types value price discounts less than lower types, meaning they are willing to purchase earlier. Notice that the requirement that each type pays a sure price when purchasing is met when the seller uses a deterministic price path, but is also satisfied for price processes where buyers purchase at random times but at a sure price. As we state next, the absence of pricing risk turns out to be a property of optimal price processes under Assumption A.

Proposition 6 *Suppose that Assumption A is satisfied and that buyers arrive at a fixed date, zero. Then, an optimal price process exists. There are two unique (weakly) decreasing sequences,* $(p_n^*)_{n=1}^N \in \mathbb{R}_+^N$ *and* $(\psi_n^*)_{n=1}^N \in [0,1]^N$ *such that, for any optimal price process and for all n:*

1. *The purchasing time $\tilde{t}_n$ of a θ_n-buyer satisfies* $\mathbb{E}[e^{-(r+\rho)\tilde{t}_n}] = \psi_n^*$ *and* $\psi_1^* = 1$.
2. *If a type θ_n buys the good, he pays a certain price p_n^*; that is,* $\Pr\big(\big(P_{\tilde{t}_n} \neq p_n^*\big) \wedge \big(\tilde{t}_n < +\infty\big)\big) = 0$.
3. *Downward incentive constraints bind; i.e.,* $\psi_n^* v_n(p_n^*) = \psi_{n+1}^* v_n(p_{n+1}^*)$ *for all n, where* $\psi_{N+1}^* \equiv 0$.

We can summarize the properties of optimal price processes as follows. As stated by Part 2 of Proposition 6, buyers face no price risk. Hence, by Lemma 2, the skimming property applies, meaning that higher types purchase earlier. In fact, by Part 1 of the proposition, the highest type (and possibly others) purchase the good immediately and for sure. Finally, downward incentive constraints bind: each buyer is indifferent between following his prescribed strategy and following the strategy of the downward adjacent type.

[21] That prices below θ_N are offered with probability zero is a reasonably anticipated requirement for optimality, following the same logic as Lemma 1 (see the proof of Proposition 6).

It is worth pointing out here the important role played by Assumption A2. In particular, suppose that this assumption fails so that higher types are *more* risk averse than lower types. In this case (and when there are three or more types), the seller can often gain by exposing lower types to pricing risk, making it less attractive for high types (when offered the chance to purchase at a given price) to mimic them. Conversely, when higher types are less risk averse, exposing lower types to pricing risk does not relax the relevant incentive constraints (given Assumption A1, these are the constraints preventing mimicry of downward adjacent types). We discuss this intuition further in Section 5.1 (especially focusing on why no restriction was needed on the risk aversion of the lowest type θ_N).

An immediate corollary of Proposition 1 is that an optimal price process exists that is deterministic.

Corollary 3 *Suppose that Assumption A is satisfied and that buyers arrive at a fixed date, zero. Then there exists an optimal price process which is deterministic. Under any optimal deterministic price process, for each $n = 1, \ldots, N$, the type θ_n-buyer purchases at time $t_n^* \equiv -\log(\psi_n^*)/(r+\rho) \in \mathbb{R}_+$ at price p_n^* if $\psi_n^* > 0$ and does not purchase otherwise (i.e., $t_n^* = +\infty$).*

Other optimal price processes can then be determined. To give an example, fix $\mu \in (0, 1]$, and consider an optimal price process P with the following characteristics. We assume that θ_1 is offered at all times except, possibly, at (finite) times $t_n = \mu t_n^*$ for $n = 1, ..., N$, when the price p_n^* is offered. Types θ_n purchase the good at time t_n if they purchase the good at all. The probability that p_n^* is offered at t_n, for $n = 2, \ldots, N$, is

$$\Pr(P_{t_n} = p_n^* | P_{t_{n-1}} = \hat{p}) = \begin{cases} 0 & \text{if } \hat{p} \neq p_{n-1}^* \\ \left(\frac{\psi_n^*}{\psi_{n-1}^*}\right)^{1-\mu} & \text{if } \hat{p} = p_{n-1}^* \end{cases},$$

while we initialize by recalling that $\psi_1^* = 1$ and $P_0 = p_1^*$ with certainty. Thus, if $n \in \{2, \ldots, N\}$ is such that $p_n^* < p_{n-1}^*$, then the price p_n^* is offered only if p_{n-1}^* has been offered before. It is then optimal for type θ_{n-1} to purchase the good at time t_{n-1} at price p_{n-1}^* (if such a price is offered). Indeed, rejecting this price and waiting until t_n to purchase the good at price p_n^* (if such price is offered), gives her a payoff of

$$\left(\frac{\psi_n^*}{\psi_{n-1}^*}\right)^{1-\mu} e^{-(r+\rho)(t_n - t_{n-1})} v_{n-1}(p_n^*) = \frac{\psi_n^*}{\psi_{n-1}^*} v_{n-1}(p_n^*)$$

which, by Condition 3 of Proposition 6, is equal to the utility he obtains from accepting the price p_{n-1}^*, i.e. $v_{n-1}(p_{n-1}^*)$.

The parameter μ captures how the price process in the example substitutes between a random failure to allocate and delay in allocation, both of which are equivalent in terms of the players' expected payoffs (recall Salant, 1989). If $\mu = 1$ then we obtain an optimal deterministic price process. Instead, as $\mu \to 0$, expected delay in purchasing converges to zero. So, the outcome of an optimal price process with small delay is close to that of a Dutch auction with a random termination time (and no capacity constraints): after every price offer, there is a probability that the trade breaks down. More precisely, the price process is akin to the following discrete-time price process without discounting.[22] In the first period, the seller offers the good at price p_1^*. In the second period, either the seller "vanishes" (so she sells no further goods) with probability $1 - \frac{\psi_2^*}{\psi_1^*}$, or offers p_2^* with probability $\frac{\psi_2^*}{\psi_1^*}$. Iteratively, in each nth period, with $n \leq N$, if the seller is yet to vanish, then she either vanishes with probability $1 - \frac{\psi_n^*}{\psi_{n-1}^*}$, or remains and offers p_n^*.

4.2 Dynamic arrivals

We return to the environment with dynamic arrivals where buyers arrive at rate $\gamma_t > 0$ for each $t \geq 0$, with $\int_0^\infty \gamma_t e^{-rt} dt = 1$. Now, analogous to (2), the seller's profits are given by

$$\int_0^\infty \gamma_\tau e^{-r\tau} \mathbb{E}\left[\Pi_\tau\right] d\tau \quad \text{where} \quad \Pi_\tau \equiv \sum_{n=1}^{N} \beta_n e^{-(r+\rho)(\tilde{t}_n^\tau - \tau)} \left(P_{\tilde{t}_n^\tau} - c\right) . \tag{6}$$

When arrivals occur over time, ensuring that different types buy at different prices becomes difficult since, at any given moment, only one price is offered. Indeed, in order to ensure that each type of buyer pays a sure price, the seller must often drop prices systematically over time, giving higher types the opportunity to purchase at higher prices (but earlier and with higher probability). This means that, depending on the arrival time, a buyer may purchase paying a lower price than he would accept at other times or other realizations of the price process (see the details below). Nonetheless, as the next result shows, one can find price processes for which expected profits are arbitrarily close to those with a fixed arrival date.

Proposition 7 *Suppose that Assumption A is satisfied and that buyers arrive over time. For any $\varepsilon > 0$, there is a price process such that the seller's expected profits in (6) are at least $\Pi^* - \varepsilon$, where Π^* is the optimal profits obtained with a fixed arrival date, as in Proposition 6.*

We now illustrate how price processes can give expected profits arbitrarily close to those with a fixed arrival date. Assume, for simplicity, that there are 3 types and that $1 = \psi_1^* > \psi_2^* > \psi_3^* > 0$. In this case, $p_3^* = \theta_3$: the lowest type does not obtain any rent. A price process which approaches optimal profits is the following, described as a process with states $\{\sigma_1, \sigma_2\}$:

1. If the state is σ_1, the price offered is p_1^* and at a Poisson arrival rate λ_2^* such that $\psi_2^* = \frac{\lambda_2^*}{\lambda_2^*+r+\rho}$ the state changes to σ_2.

[22] The logic is closely related to Bose and Daripa (2009), who study a variation of a Dutch auction in a setting with competition among ambiguity-averse buyers.

2. If the state is σ_2 the price is p_2^*, and a price change arrives at rate Λ. Price changes are independent of past prices and either up with probability μ or down with probability $1-\mu$. If the price change is upwards, then the price changes to p_1^* and the state switches to σ_1. If it is downwards, then the price θ_3 is offered at the moment of the price change, and subsequently the state is σ_1.

The state σ_1 is a "regular state", where the usual "high" price is charged, while σ_2 is a "discount" state, where a low price is charged. The discount state ends with either a larger "final" price reduction with price θ_3 (with probability $1-\mu$), or by returning to the regular state. Notice that the expected payoff of the θ_2-buyer at state σ_1 if he plans on buying the first time where there is a discount is $\psi_2^* v_2(p_2^*)$. His expected payoff if instead waiting to obtain it at price θ_3 is the value x solving

$$x = \psi_2^* \left(\frac{(1-\mu)\Lambda}{\Lambda + r + \rho} v_2(\theta_3) + \frac{\mu\Lambda}{\Lambda + r + \rho} x \right).$$

Indifference between these options requires $x = \psi_2^* v_2(p_2^*)$ to solve the above equation, in which case the expected discounting of the time where the price is θ_3 is ψ_3^*, as prescribed in the optimal mechanism with known arrivals. Type θ_1 buyers purchase on arrival, θ_2-buyers purchase in state σ_2 or if the price is set at θ_3, and θ_3-buyers purchase when the price is θ_3.

Optimal profits are not achieved by this price process because there are θ_1-buyers who arrive when the state is σ_2, so they purchase at price p_2^* instead of p_1^*. Still, as Λ increases, the expected fraction of time where the state is σ_1 approaches one. Hence, the fraction of θ_1-buyers arriving at a time where the state is σ_2 becomes arbitrarily small, and profits approach those of the fixed-arrival case. This illustrates why optimal profits can be arbitrarily approximated but not achieved. Due to the restriction to charging all buyers purchasing on the same date the same price, the seller cannot discriminate between types θ_2 and θ_3 within a single instant.

5 Discussion

5.1 More general mechanisms

We now briefly consider the relationship between stochastic price processes and more general dynamic formats. Implicit in the description of the environment, we rule out temporary rentals of the good and we rule out payments at dates other than purchase dates. The most general class of mechanisms for our setting is revelation mechanisms in which buyers report their type at their arrival time, and the mechanism determines a (possibly random) date of purchase and a (possibly random) price to be paid on that date.

Consider now the fixed-arrivals setting with multiple types as in Section 4.1. We will argue that the additional generality of dynamic revelation mechanisms does not permit higher profits under the following strengthening of Assumption A2, which imposes a joint restriction on the preferences of all types $\{\theta_1, \ldots, \theta_N\}$.

Assumption A2* 1 *Higher types are less risk averse and the condition applies to all types* $\{\theta_1, \ldots, \theta_N\}$. *For any* $n = 1, ..., N-1$ *and* $p \in U$, $\frac{v_n''(p)}{v_n'(p)} \leq \frac{v_{n+1}''(p)}{v_{n+1}'(p)}$.

Note that, relative to Assumption A2, Assumption 12* now places a restriction on the risk aversion of the lowest type θ_N. Recall that a key reason for considering the more permissive Assumption A2 was that it imposes no restriction in the case of two types (while the two-type case seems to hold particular interest for applications).

Abusing somewhat notation, a (dynamic) revelation mechanism for our environment is a probability space $(\Omega, \mathcal{F}, \mathcal{P})$ and a collection of (measurable) mappings $\left(\tilde{t}_{\hat{n}}, P^{\hat{n}}\right) : \Omega \to \bar{\mathbb{R}}_+ \times U$, $\theta_{\hat{n}} \in \Theta$. A buyer who reports $\theta_{\hat{n}} \in \Theta$, in outcome ω, receives the good at time $\tilde{t}_{\hat{n}}(\omega)$, paying price $P^{\hat{n}}(\omega)$ (if $\tilde{t}_{\hat{n}}(\omega) = +\infty$, the buyer does not receive the good and the "price" $P^{\hat{n}}(\omega)$ does not enter payoffs). There is full commitment on both sides of the relationship; in particular, a buyer who participates in the mechanism and reports $\theta_{\hat{n}}$ is bound to purchase at the price and time (perhaps randomly) determined by the mechanism. The mechanism is required to be incentive compatible: for any $\theta_n, \theta_{\hat{n}} \in \Theta$,

$$\mathbb{E}\left[e^{-(r+\rho)\tilde{t}_n} v_n\left(P^n\right)\right] \geq \mathbb{E}\left[e^{-(r+\rho)\tilde{t}_{\hat{n}}} v_n(P^{\hat{n}})\right].$$

It must also be individually rational: $\mathbb{E}\left[e^{-(r+\rho)\tilde{t}_n} v_n\left(P^n\right)\right] \geq 0$ for all $\theta_n \in \Theta$.

Now, for any stochastic price process P and associated stopping times $\tilde{t}_n$, we can define a revelation mechanism that induces the same dates and price of purchase. In particular, for each report θ_n, let the random time $\tilde{t}_n$ be given by the aforementioned stopping time and let $P^n = P_{\tilde{t}_n}$ (with $P^n = \bar{u}$, say, when $\tilde{t}_n = +\infty$). Incentive compatibility of the mechanism $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$ (with probability space identical to that for P) is inherited from the incentive compatibility of stopping times for the process P. Hence, as one should expect, expected profits for optimal revelation mechanisms are no lower than for price processes (and this holds irrespective of assumptions on the shape of v_n). Whether strictly higher profits can be attained in revelation mechanisms depends on the nature of risk aversion, with Assumption 12* a sufficient condition for equivalence.

Proposition 8 *Consider the case with* $N \geq 2$ *types and a fixed arrival date as in Section 4.1.*

1. *If Assumptions A1 and 12* hold, then optimal expected profits are equal for price processes and revelation mechanisms.*

2. *If Assumption A1 and 12 hold, but 12* does not, then profits may be strictly higher for revelation mechanisms.*

To understand this result, first consider Part 2. For general revelation mechanisms, the individual-rationality requirement is ex-ante (as in much of the literature on mechanism design for risk-averse agents, e.g. Matthews (1983) and Maskin and Riley (1984)). Hence, the lowest type θ_N, in particular, can agree to a price that, with positive probability, is above θ_N. If this type is less risk averse than higher types, then asking the type to pay a random price that is above θ_N with positive probability (while holding the type's expected payoff constant) relaxes the incentive constraints of higher types. As we demonstrate in the proof (with an example based on one in Maskin and Riley), this can permit higher profits than for any (stochastic) price process, where the lowest type cannot be charged a price higher than θ_N with positive probability. In particular, under Assumptions A1 and A2, any optimal price process charges the lowest type a sure price θ_N whenever he purchases, since this is the price that maximizes profits for the type while making mimicry by higher types as unattractive as possible.

When Assumptions A1 and 12* hold, as in Part 1, replacing the random price for any type θ_n with an appropriately defined "certainty equivalent" which holds type θ_n's expected payoff constant (while specifying purchase at the same time) both increases profits for type θ_n (provided the price is higher than the per-unit cost c) and relaxes incentive constraints for all higher types. Due to Assumption 12* this is true for all types, not just those above θ_N (which explains why the strengthening of Assumption 12 is needed).

Note then that, when Assumptions A1 and 12* hold, the optimal revelation mechanism generates the same expected payoffs for both the seller and all types of buyer as the optimal deterministic price process, which induces purchase at dates t_n^* and prices p_n^*, as specified in Corollary 3. It is then easy to see (as explained, for example, by Salant, 1989) that a payoff-equivalent static revelation mechanism exists in which type θ_n buyers receive the good at date zero at price p_n^* with probability $e^{-(r+\rho)t_n^*}$, and do not receive the good (and pay nothing) otherwise. Hence, under these assumptions, any optimal stochastic price process is payoff equivalent to a static mechanism; i.e., one in which the good is traded at date zero or not at all.

Finally, note that, extending these results to dynamic arrivals is straightforward. For instance, the dynamic revelation mechanism can ask that buyers participate on their arrival dates. The revelation mechanism on that date is simply the optimal one for the above problem (taking the participation date as the fixed arrival date). Provided the seller keeps a record of buyers' identities, it can restrict buyers to participating in the mechanism only once. Under these conditions, buyers' private information on their arrival dates are not a source of rents, and expected profits are the same as for the case with a fixed arrival date. In particular, unlike what we find for a stochastic price process, revelation mechanisms permit optimal profits to be attained.

5.2 Patterns of price discounts

Another issue that is complicated by the presence of more than two types is the timing of price discounts with dynamic arrivals. One complication is that, as described in Section 4.2, optimal profits for stochastic price processes can be approached, but typically not attained. Predictions about the timing of price discounts (for near-optimal price processes) are therefore more cumbersome to state.

When buyers observe the history of past prices, as in Section 4.2, any near-optimal price process must satisfy stringent restrictions. For instance, we anticipate that episodes of price discounting should be ephemeral and arrive according to a process that is not too far (in an appropriate sense) from Poisson, with a constant rate (at least over bounded time horizons). Conversely, when buyers do not observe prices before their arrival to the market (recall Section 3.3 for two types), a wider range of processes for discounts can approach optimal profits. As for the case with two types, episodes of price discounting should not be clustered, since any buyer arriving to the market would have an incentive to delay purchase for a short time and then, if no discount arrives (so the buyer is pessimistic about the timing of future discounts), purchase at the undiscounted price.

5.3 Degree of risk aversion

A natural concern for our theory is whether the degree of risk aversion typically exhibited by consumers is large enough relative to the pricing risks faced in most retail markets. Clearly, for smaller items (such as compact discs), pricing risks are so small as to be practically irrelevant once we assume risk preferences are stable across a consumers' financial decisions (including decisions with larger stakes). On the other hand, risk preferences are widely believed *not* to be stable across contexts, and individuals often exhibit aversion to "small-scale" risks (see Rabin, 2000, for a discussion and see Sydnor, 2010, for evidence from insurance markets). Such aversion to small risks is often associated with loss aversion, which seems consistent with our theory although (unlike a large literature following Koszegi and Rabin, 2006) we have not explicitly modeled reference-dependent preferences.

A related point is that the key assumptions in our theory are (i) a particular shape of consumer preferences over prices and dates of purchase, and (ii) consumer expected utility maximization. Importantly, consumers who expect to have to wait a long time for discounted prices are willing to purchase at relatively high prices (the seller often finds it optimal to randomize prices to guarantee that all cohorts of consumers expect to wait a relatively long time for a price reduction). While we have related such preferences to risk aversion (which in turn permits us to relate our findings to a broader literature on mechanism design with risk-averse agents), it is possible the shape of these intertemporal preferences are determined by 'psychological' forces that are different to those usually associated with risk aversion in other environments (for instance, such preferences might stem from particular framing effects in (some) retail markets).

6 Conclusions

We have provided a theory of randomly timed price discounts which can account for buyers' apparent failure to wait and get the best deals in the market. Buyer risk aversion provides both a motive for intertemporal price discrimination – discounting the good to sell to low-value buyers only occasionally – and (when buyers arrive dynamically) for discounting the good at random times. The latter leads to a constant "regular" price, which insures high-value buyers against pricing risk associated with the timing of their arrival to the market. Optimal price processes also insure lower types against pricing risks.

With two types, price discounts take the form of instantaneous random sales. With more than two types, near-optimal price processes involve randomly timed episodes of discounts that are reminiscent of a Dutch auction with a random termination time. For many goods (such as those reviewed in the Introduction), episodes of many successively steeper discounts seem less pertinent than patterns of "high-low pricing" that consist of a high "regular" price and low "sale" price. This might simply reflect practical considerations such as menu costs. Indeed, a number of our insights appear robust even if price processes in practice (say due to menu costs or other practical considerations) depart from the predicted optimum. For instance, to the extent that buyers who purchase at high "regular" prices are risk averse and have time-independent beliefs about future discounts, the fact that such regular prices tend to be fairly stable over time is a form of insurance against pricing risk associated with the time of arrival to the market. That price discounts or sales often last a relatively short time also mitigates the pricing risk such buyers face.

Another finding we believe may be particularly robust is that price discounts should tend to exhibit "repulsion" rather than "clustering". The key idea is that buyers who are new to a market may struggle to access information on past pricing decisions (such as the timing of earlier discounts). While this means a range of processes for the dates of discounting can be optimal, these dates should not be clustered. Buyers who anticipate clustering of discounts would often prefer to delay purchase for a fixed amount of time, in order to "test" whether they have arrived at a time when the intensity of discounting is high.

It is worth reiterating that the rationale for price discounting in our model is related to other theories that posit full-commitment pricing. In particular, intertemporal price discrimination has often been associated with buyers that face higher discount rates than the seller (see Landsberger and Meilijson, 1985) or who have values for the good that change with time (either deterministically, as in Stokey, 1979, or randomly as in Garrett, 2016). Importantly, however, these theories do not suggest any motive for random pricing.[23] On the other hand, buyer uncertainty about the timing of future discounts seems likely to be a fairly robust empirical regularity across a range of markets, since few retailers provide a detailed account of their intentions to discount (both when such discounts are planned and how large they will be). Still, a range of factors may well be at play in determining both buyers' decisions and optimal pricing by sellers, so we view our theory as complementary to existing explanations for price discounts.

Finally, our work may prove relevant for future empirical studies. Such studies might have at least three objectives: First, to distinguish among various theories and suggest which are the most relevant in different settings; second, and related, to understand better buyers' expectations about future prices and discounting; and third, to fit parsimonious models of demand.

[23]For instance, Garrett argues that the optimal price path in his setting with risk-neutral consumers is deterministic.

References

[1] Bansal, Matulya and Costis Maglaras (2009), 'Dynamic pricing when customers strategically time their purchase: Asymptotic optimality of a two-price policy,' *Journal of Revenue and Pricing Management*, 8, 42-66.

[2] Berck, Peter, Jennifer Brown, Jeffrey M. Perloff and Sofia B. Villas-Boas (2008), 'Sales: Tests of theories on causality and timing,' *International Journal of Industrial Organization*, 26, 1257-1273.

[3] Board, Simon and Andrzej Skrzypacz (2016), 'Revenue Management with Forward-Looking Buyers,' *Journal of Political Economy*, 124, 1046-1087.

[4] Bose, Subir and Arup Daripa (2009), 'A dynamic mechanism and surplus extraction under ambiguity,' *Journal of Economic Theory*, 144, 2084-2114.

[5] Chevalier, Judith A. and Austan Goolsbee (2009), 'Are Durable Goods Consumers Forward-Looking? Evidence from College Textbooks.' *Quarterly Journal of Economics*, 124, 1853-1884.

[6] Chevalier, Judith A. and Anil K. Kashyap (2014), 'Best Prices: Price Discrimination and Consumer Substitution,' NBER Working Paper No. 20768.

[7] Conlisk, John, Eitan Gerstner and Joel Sobel (1984), 'Cyclic Pricing by a Durable Goods Monopolist,' *Quarterly Journal of Economics*, 99, 489-505.

[8] Dilmé, Francesc and Fei Li (2016), 'Revenue Management Without Commitment: Dynamic Pricing and Periodic Fire Sales,' mimeo University of Bonn and UNC at Chapel Hill.

[9] Feller, William (1968), *An introduction to probability theory and its applications* (Vol. 1, 3rd Ed.). New York: Wiley.

[10] Fershtman, Chaim and Arthur Fishman (1992), 'Price Cycles and Booms: Dynamic Search Equilibrium,' *American Economic Review*, 82, 1221-1233.

[11] Février, Philippe and Lionel Wilner (2016), 'Do consumers correctly expect price reductions? Testing dynamic behavior,' *International Journal of Industrial Organization*, 44, 25-40.

[12] Garrett, Daniel F. (2016), 'Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values,' *American Economic Review*, 106, 3275-3299.

[13] Gershkov, Alex, Benny Moldovanu and Philipp Strack (forthcoming), 'Revenue-Maximizing Mechanisms with Strategic Customers and Unknown, Markovian Demand ,' *Management Science*.

[14] Gowrisankaran, Gautam and Marc Rysman (2012), 'Dynamics of Consumer Demand for New Durable Goods,' *Journal of Political Economy*, 120, 1173-1219.

[15] Heidhues, Paul and Botond Kőszegi (2014), 'Regular prices and sales,' *Theoretical Economics*, 9, 217-251.

[16] Hendel, Igal and Aviv Nevo (2013), 'Intertemporal Price Discrimination in Storable Goods Markets,' *American Economic Review*, 103, 2722-2751.

[17] Hollander, Myles and Frank Proschan (1984), '27 Nonparametric concepts and methods in reliability,' in *Handbook of Statistics* (ed. Krishnaiah, P.R. and P.K. Sen), 4, 613-655.

[18] Hörner, Johannes and Larry Samuelson (2011), 'Managing Strategic Buyers,' *Journal of Political Economy*, 119, 379-425.

[19] Kamenica, Emir and Matthew Gentzkow (2011), 'Bayesian Persuasion,' *American Economic Review*, 101, 2590-2615.

[20] Kőszegi, Botond and Matthew Rabin (2006), 'A Model of Reference-Dependent Preferences,' *Quarterly Journal of Economics*, 121, 1133-1165.

[21] Landsberger, Michael and Isaac Meilijson (1985), 'Intertemporal price discrimination and sales strategy under incomplete information', *Rand Journal of Economics*, 15, 171-196.

[22] Lan, Hao, Tim Lloyd and Wyn Morgan (2016), 'Are Temporary Price Reductions Predictable? The Hazard Function of Sales,' mimeo U. of Bournemouth, U. of East Anglia and U. of Sheffield.

[23] Liu, Qian and Garrett J. van Ryzin (2008), 'Strategic Capacity Rationing to Induce Early Purchases,' *Management Science*, 54, 1115-1131.

[24] Maskin, Eric and John Riley (1984), 'Optimal Auctions with Risk Averse Buyers,' *Econometrica*, 52, 1473-1518.

[25] Matthews, Steven A. (1983), 'Selling to risk averse buyers with unobservable tastes,' *Journal of Economic Theory*, 30, 370-400.

[26] Nakamura, Emi and Jón Steinsson (2008), 'Five Facts about Prices: A Reevaluation of Menu Cost Models,' *Quarterly Journal of Economics*, 123, 1415-1464.

[27] Ortner, Juan (2017), 'Durable goods monopoly with stochastic costs,' *Theoretical Economics*, 12, 817-861.

[28] Ory, Aniko (2017), 'Consumers on a Leash: Advertised Sales and Intertemporal Price Discrimination,' mimeo Yale University.

[29] Pesendorfer, Martin (2002), 'Retail Sales: A Study of Pricing Behavior in Supermarkets', *Journal of Business*, 75, 33-66.

[30] Rabin, Matthew (2000), 'Diminishing marginal utility of wealth cannot explain risk aversion.' In *Choices, Values, and Frames* (Daniel Kahneman and Amos Tversky, eds.), 202–208, Cambridge University Press, Cambridge.

[31] Rochet, Jean-Charles and John E. Thanassoulis (2017), 'Intertemporal Price Discrimination with Multiple Products,' CEPR Discussion Paper No. DP12034.

[32] Salant, Stephen W. (1989), 'When is Inducing Self-Selection Suboptimal For a Monopolist?', *Quarterly Journal of Economics*, 104, 391-397.

[33] Seiler, Stephan (2013), 'The impact of search costs on consumer behavior: A dynamic approach,' *Quantitative Marking and Economics*, 11, 155-203.

[34] Shilony, Yuval (1977), 'Mixed pricing in oligopoly,' *Journal of Economic Theory*, 14, 373-388.

[35] Sobel, Joel (1991), 'Durable Goods Monopoly with Entry of New Consumers', *Econometrica*, 59, 1455-1485.

[36] Stokey, Nancy L. (1979), 'Intertemporal Price Discrimination,' *Quarterly Journal of Economics*, 93, 355-371.

[37] Sydnor, Justin (2010), '(Over)insuring Modest Risks,' *American Economic Journal: Applied Economics*, 2, 177-199.

[38] Varian, Hal R. (1980), 'A Model of Sales', *American Economic Review*, 70, 651-659.

[39] 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.

A Proofs

A.1 Proofs of Results in Section 3

Proof of Lemma 1

Proof. Fix an optimal price process P^* and optimal incentive-compatible stopping times $\tilde{t}_1$ and $\tilde{t}_2$. We begin by showing that $\Pr(\tilde{t}_1 = 0) = 1$ and that there is some value $\bar{p}_0 \in [\theta_2, \theta_1]$ such that the distribution of P_0^* is degenerate at $\bar{p}_0$.

Assume, for the sake of contradiction, that either $\Pr(\tilde{t}_1 = 0) < 1$ or there is no $\bar{p}_0$ such that $\Pr(P_0^* = \bar{p}_0) = 1$. Define $\check{p}_0$ as follows:[24]

$$\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}\big]v_1(\check{p}_0) = \mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}v_1(P^*_{\tilde{t}_1})\big] \ .$$

We first claim that $\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}\big]\,(\check{p}_0 - c) \geq \mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}\big(P^*_{\tilde{t}_1} - c\big)\big]$, with strict inequality whenever $P^*_{\tilde{t}_1}$ has a non-degenerate distribution on $\{\omega|\tilde{t}_1(\omega) < +\infty\}$. To see this, note that concavity of v_1 and Jensen's inequality implies that

$$v_1(\check{p}_0) = \frac{\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}v_1(P^*_{\tilde{t}_1})\big]}{\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}\big]} \leq v_1\left(\frac{\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}P^*_{\tilde{t}_1}\big]}{\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}\big]}\right),$$

with strict inequality whenever $P^*_{\tilde{t}_1}$ has a non-degenerate distribution on $\{\omega|\tilde{t}_1(\omega) < +\infty\}$.

Notice then that, since the price process P^* is assumed to be optimal, we have

$$\theta_2 - c \leq \beta\, \underbrace{\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}(P^*_{\tilde{t}_1} - c)\big]}_{\leq \mathbb{E}[e^{-(r+\rho)\tilde{t}_1}](\check{p}_0 - c) \leq \check{p}_0 - c} + (1-\beta)\, \underbrace{\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_2}(P^*_{\tilde{t}_2} - c)\big]}_{\leq \theta_2 - c} \ ,$$

and so $\check{p}_0 \geq \theta_2$. Now, fix some $\varepsilon > 0$, and consider the price process P' satisfying, for all $\omega \in \Omega$:

$$P'_t(\omega) = \begin{cases} \check{p}_0 & \text{if } t = 0, \\ \bar{u} & \text{if } t \in (0, \varepsilon), \\ P^*_{t-\varepsilon}(\omega) & \text{if } t \geq \varepsilon. \end{cases}$$

The seller can then specify that a high type purchases at date zero and a low type does so at time $\varepsilon + \tilde{t}_2$. Indeed, by purchasing immediately at date zero, a high type earns at least the payoff obtained under the original price path, while he obtains a lower payoff by not purchasing at date zero and following any continuation strategy. Using that $\check{p}_0 > c$ and that either $\Pr(\tilde{t}_1 = 0) < 1$ or $P^*_{\tilde{t}_1}$ has a

[24] Notice that, in any optimal price process, $\mathbb{E}[e^{-(r+\rho)\tilde{t}_1}] > 0$. This is the case because $\mathbb{E}[e^{-(r+\rho)\tilde{t}_1}] = 0$ implies $\mathbb{E}[e^{-(r+\rho)\tilde{t}_2}] = 0$ (since the θ_1-buyer has the option to imitate the θ_2-buyer), but the payoff of the seller is bounded below by $\theta_2 - c > 0$. Also, since the θ_1-buyer almost never accepts prices above θ_1 and there is either delay or a non-degenerate distribution of accepted prices, we have $\check{p}_0 < \theta_1$.

non-degenerate distribution on $\{\omega|\tilde{t}_1(\omega) < +\infty\}$, the seller's expected profits conditional on type θ_1 strictly increase. Such an increase is independent of ε, and it is equal to

$$\mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}\big]\check{p}_0 - \mathbb{E}\big[e^{-(r+\rho)\tilde{t}_1}P^*_{\tilde{t}_1}\big] > 0.$$

Hence, So, if $\varepsilon > 0$ is small enough, the seller's expected profits conditional on type θ_2 are close enough to that of the original process that the new price process generates a higher profit than P^*, a contradiction.

Finally, it is clear that $\bar{p}_0$ is weakly higher than θ_2, since the price paid by a low type is bounded above by θ_2 while the optimal profit per buyer is bounded below by θ_2.

The proof that $\Pr(\tilde{t}_2 = 0) \in \{0, 1\}$ follows from the above arguments. In particular, it follows because P^*_0 is almost surely a constant and the θ_2 buyer can only condition her decision at time 0 on the realization of P^*_0. We have $\Pr(\tilde{t}_2 = 0) = 0$ if he rejects the price $\bar{p}_0$ at time 0 and $\Pr(\tilde{t}_2 = 0) = 1$ otherwise.

Finally, consider why $P^*_t \geq \theta_2$ almost surely. If this is not true, then a low type expects a positive rent, purchasing with positive probability at a price below θ_2. The above arguments then imply that $\Pr(\tilde{t}_2 = 0) = 0$.

Let $\psi \equiv \mathbb{E}\left[e^{-(r+\rho)\tilde{t}_2}\right]$ and define $\bar{t}_2 \equiv -\log(\psi)/(r+\rho)$. Recall that P^*_0 is degenerate at $\bar{p}_0 \geq \theta_2$ and that $\Pr(\tilde{t}_1 = 0) = 1$. Specify the deterministic price process P' as

$$P'_t = \begin{cases} \bar{p}_0 & \text{if } t = 0, \\ \theta_2 & \text{if } t = \bar{t}_2, \\ \bar{u} & \text{otherwise.} \end{cases}$$

We claim that for this process, the high type can be asked to purchase at date zero and the low type at date $\bar{t}_2$. The latter follows because the low type's optimal payoff is zero under P'. The former follows because, if the high type delays purchase, it optimally purchases at $\bar{t}_2$. He then obtains a payoff

$$\psi v_1(\theta_2) \leq \mathbb{E}\left[e^{-(r+\rho)\tilde{t}_2} v_1\big(P^*_{\tilde{t}_2}\big)\right],$$

since $\Pr\big(P^*_{\tilde{t}_2} > \theta_2\big) = 0$. Hence, if the high type is not willing to purchase at date zero under P', the purchasing time $\tilde{t}_2$ generates a strictly higher payoff than date-zero purchase for this type under P^*. That is, $\tilde{t}_1$ satisfying $\Pr(\tilde{t}_1 = 0) = 1$ cannot be an incentive-compatible purchase time.

Finally, note that the expected profits for a high type buyer are the same under P^* and P'. Conditional on a low type, expected profits are $\psi(\theta_2 - c)$ under the deterministic price process P', but

$$\mathbb{E}\left[e^{-(r+\rho)\tilde{t}_2}\left(P^*_{\tilde{t}_2}-c\right)\right]=\mathbb{E}\left[e^{-(r+\rho)\tilde{t}_2}P^*_{\tilde{t}_2}\right]-\psi c$$
$$<\psi\left(\theta_2-c\right)$$

under P^*, where the inequality follows because $\Pr\left(P^*_{\tilde{t}_2}<\theta_2\right)>0$. ■

Proof of Proposition 1

Proof. Existence of an optimal expected discount factor ψ^* follows by continuity of (1) in ψ, while uniqueness follows from its strict concavity. The derivative of seller profits per buyer (1) with respect to ψ is given by

$$\beta v_1(\theta_2)p_1'\left(\psi v_1(\theta_2)\right)+(1-\beta)\,\theta_2, \tag{7}$$

which must be non-positive when $\psi=\psi^*=0$, non-negative when $\psi=\psi^*=1$, and equal to zero when $\psi=\psi^*\in(0,1)$. The result then follows taking $\underline{\beta}$ to satisfy

$$\underline{\beta}v_1(\theta_2)p_1'\left(v_1(\theta_2)\right)+\left(1-\underline{\beta}\right)\theta_2=0$$

and $\bar{\beta}$ to satisfy

$$\bar{\beta}v_1(\theta_2)p_1'\left(0\right)+\left(1-\bar{\beta}\right)\theta_2=0.$$

Notice that, by the concavity of v_1, we have $p_1'\left(v_1(\theta_2)\right)<p_1'(0)$, so $\underline{\beta}<\bar{\beta}$. It is then easy to see that $\psi^*\in(0,1)$ iff $\beta\in\left(\underline{\beta},\bar{\beta}\right)$. ■

Proof of Proposition 2

Proof. Existence of an optimal price process. Expected profits are defined in (2) to equal

$$\int_0^\infty\gamma_\tau e^{-r\tau}\mathbb{E}\left[\Pi_\tau\right]d\tau,$$

where (making dependence on the outcome ω explicit)

$$\Pi_\tau(\omega)=\beta e^{-(r+\rho)(\tilde{t}_1^\tau(\omega)-\tau)}(P_{\tilde{t}_1^\tau(\omega)}(\omega)-c)+(1-\beta)e^{-(r+\rho)(\tilde{t}_2^\tau(\omega)-\tau)}(P_{\tilde{t}_2^\tau(\omega)}(\omega)-c)$$

describes profit from date-τ arrivals for outcome ω. It is easy to see that, for all τ,

$$\mathbb{E}\left[\Pi_\tau\right]\leq\Pi^*=\beta\left(p_1\left(\psi^*v_1(\theta_2)\right)-c\right)+(1-\beta)\psi^*\left(\theta_2-c\right),$$

where Π^* is the profit obtained in case of a known arrival date, as in Proposition 1.

Now, consider a standard homogeneous Poisson process $\{N\left(t\right),t\geq0\}$ with parameter λ^*. Suppose then that p_1^* is offered at all instants except for the times of discontinuity of the Poisson process, where the price equals to θ_2. Incentive-compatible stopping times are then given, for all τ, by $\tilde{t}_1^{\tau} = \tau$ for type θ_1 and $\tilde{t}_2^{\tau} = \min\{t \in [\tau, \infty) : N(t) > \lim_{s \nearrow t} N(s)\}$ for type θ_2; i.e., high types purchase immediately while low types purchase at sales. It is then readily verified that $\mathbb{E}[\Pi_\tau] = \Pi^*$ for all τ.

Properties of any optimal price process. The necessary conditions for optimality (Parts 1 and 2 of the proposition) are a result of the following observation, together with the findings in Section 3.1. Fix a stochastic price process P (and associated stopping times). For any date t, and any positive probability event $A_t \in \mathcal{F}_t$, $\mathbb{E}[\Pi_t | A_t] \leq \Pi^*$.[25] Hence, the existence a positive probability event A_t with $\mathbb{E}[\Pi_t | A_t] < \Pi^*$ ensures $\mathbb{E}[\Pi_t] < \Pi^*$. It follows that, if there exists a positive measure set Z of times z with $\mathbb{E}[\Pi_z | A_z] < \Pi^*$ for a positive probability event $A_z \in \mathcal{F}_z$, then P cannot be optimal.

The first condition then follows by the same arguments as in Section 3.1, where we showed that attaining expected profits Π^* requires high types to purchase immediately and at the deterministic price p_1^*.

For the second condition, begin by considering the price paid by a low type arriving at $\tau > 0$. If there is a positive probability such a buyer purchases at a price greater than θ_2, then the stopping time $\tilde{t}_2^{\tau}$ is not incentive compatible. If there is a positive probability such a buyer pays a price strictly less than θ_2, then this is also true for a positive measure of arrival times earlier than τ. Again, this implies (using the arguments in Section 3.1 and the above claim) that P cannot be optimal.

Again from the arguments in Section 3.1, we can note that, for almost all $\tau > 0$, the conditional expectation $\mathbb{E}\left[e^{-(r+\rho)\left(\tilde{t}_2^{\tau} - \tau\right)} \middle| \mathcal{F}_\tau\right]$ must be degenerate at ψ^*. We now show that in fact this holds at *any* $\tau > 0$. Otherwise, there is a positive probability event $A_\tau \in \mathcal{F}_\tau$ with $\mathbb{E}\left[e^{-(r+\rho)\left(\tilde{t}_2^{\tau} - \tau\right)} | A_\tau\right] \neq \psi^*$ for some $\tau > 0$ and the claim follows from the following result.

Lemma 3 *Suppose that* $\mathbb{E}\left[e^{-(r+\rho)\left(\tilde{t}_2^{\tau} - \tau\right)} \middle| A_\tau\right] \neq \psi^*$ *for some* $\tau > 0$ *and positive probability event* $A_\tau \in \mathcal{F}_\tau$*. Then there there is a positive measure of dates* z *in a neighborhood of* τ *such that* $\mathbb{E}\left[e^{-(r+\rho)(\tilde{t}_2^{z} - z)} \middle| A_z\right] \neq \psi^*$ *for positive probability events* $A_z \in \mathcal{F}_z$*.*

Proof. Suppose that $\mathbb{E}\left[e^{-(r+\rho)\left(\tilde{t}_2^{\tau} - \tau\right)} \middle| A_\tau\right] < \psi^*$ for some $\tau > 0$ and positive probability event $A_\tau \in \mathcal{F}_\tau$. For any $z > \tau$, let $\check{A}_z \equiv \left\{\omega \in A_\tau : \tilde{t}_2^{\tau} < z\right\} \in \mathcal{F}_z$ so, whenever $\check{A}_z$ has positive probability, we have $\mathbb{E}\left[e^{-(r+\rho)(\tilde{t}_2^{\tau} - \tau)} \middle| \check{A}_z\right] \geq e^{-(r+\rho)(z-\tau)}$. For any $\eta \in \left(0, \psi^* - \mathbb{E}\left[e^{-(r+\rho)\left(\tilde{t}_2^{\tau} - \tau\right)} \middle| A_\tau\right]\right)$, we can find $\varepsilon_\eta > 0$ such that, for all $z \in (\tau, \tau + \varepsilon_\eta)$,

$$\mathbb{E}\left[e^{-(r+\rho)\left(\tilde{t}_2^{z} - z\right)} \middle| \check{A}_z\right] \leq 1 \leq e^{-(r+\rho)(z-\tau)} + \eta \leq \mathbb{E}\left[e^{-(r+\rho)\left(\tilde{t}_2^{\tau} - \tau\right)} \middle| \check{A}_z\right] + \eta$$

[25] If instead $\mathbb{E}[\Pi_t | A_t] > \Pi^*$ for some positive probability event A_t, then one can construct a price process for which expected profits exceed Π^* in the fixed-arrivals problem of Section 3.1. For example, this can be achieved by "shrinking" the time interval $[0, t)$ and increasing prices on that interval above θ_1 in such a way that the information contained in prices on the (shrinked) interval is still revealed to buyers, but buyers wait until after this interval to purchase.

whenever $\check{A}_z$ has positive probability, while

$$\mathbb{E}\big[e^{-(r+\rho)\left(\tilde{t}_2^z-z\right)}\big|A_\tau\backslash\check{A}_z\big] = e^{(r+\rho)(z-\tau)}\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|A_\tau\backslash\check{A}_z\big] \\ \leq \mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|A_\tau\backslash\check{A}_z\big] + \eta,$$

whenever $A_\tau\backslash\check{A}_z$ has positive probability, where we use that (by assumption) $\tilde{t}_2^z = \tilde{t}_2^\tau$ on $A_\tau\backslash\check{A}_z$. Hence, for all $z \in (\tau, \tau+\varepsilon_\eta)$,

$$\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^z-z)}\big|A_\tau\big] = \Pr(\check{A}_z|A_\tau)\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^z-z)}\big|\check{A}_z\big] + \big(1-\Pr(\check{A}_z|A_\tau)\big)\,\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^z-z)}\big|A_\tau\backslash\check{A}_z\big] \\ \leq \mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|A_\tau\big] + \eta \\ < \psi^*.$$

This establishes the result when $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|A_\tau\big] < \psi^*$.

Now suppose that instead that there is a $\tau > 0$ such that $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|A'_\tau\big] \geq \psi^*$ for all positive probability events $A'_\tau \in \mathcal{F}_\tau$, and $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|A_\tau\big] > \psi^*$ for some positive probability event $A_\tau \in \mathcal{F}_\tau$, so $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big] > \psi^*$. For any $z \in (0,\tau)$, we have that

$$\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^z-z)}\big] \geq \Pr(\hat{A}_z)e^{-(r+\rho)(\tau-z)} + \Pr(\Omega\backslash\hat{A}_z)\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-z)}\big|\Omega\backslash\hat{A}_z\big] \\ = e^{-(r+\rho)(\tau-z)}\big(1-\Pr(\Omega\backslash\hat{A}_z)\big(1-\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|\Omega\backslash\hat{A}_z\big]\big)\big)$$

where $\hat{A}_z \equiv \{\omega \in \Omega : \tilde{t}_2^z < \tau\}$. For any z is close enough to τ satisfying $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^z-z)}\big] \leq \psi^*$ it is necessay that the probability of $\Omega\backslash\hat{A}_z$ is positive so, from the previous equation,

$$\Pr(\Omega\backslash\hat{A}_z) \geq \frac{1-\psi^* e^{(r+\rho)(\tau-z)}}{1-\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|\Omega\backslash\hat{A}_z\big]}.$$

So, since $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|\Omega\backslash\hat{A}_z\big] \geq \psi^*$ by assumption, $\Pr(\Omega\backslash\hat{A}_z)$ is close to 1 if z is close to τ and, as a result, $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^z-z)}\big|\Omega\backslash\hat{A}_z\big]$ is close to $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big] > \psi^*$, so $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^z-z)}\big] > \psi^*$, which is a contradiction. Then, $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^z-z)}\big] > \psi^*$ for $z \in (\tau-\varepsilon,\tau)$, for $\varepsilon > 0$ small enough, which contradicts the optimality of the price process. ■

Finally, given that $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|\mathcal{F}_\tau\big] = \psi^*$ at all $\tau > 0$, $\tilde{t}_2^\tau$ is exponentially distributed with parameter λ^*, independent of past prices. Otherwise, there must be a $z > \tau$ and positive probability event $A_z = \big\{\omega \in \Omega : \tilde{t}_2^\tau \geq z\big\}$ such that $\mathbb{E}\big[e^{-(r+\rho)\left(\tilde{t}_2^\tau-z\right)}\big|A_z\big] \neq \psi^*$. Because, by assumption, $\tilde{t}_2^z = \tilde{t}_2^\tau$ on A_z, $\mathbb{E}\big[e^{-(r+\rho)\left(\tilde{t}_2^z-z\right)}\big|A_z\big] \neq \psi^*$, contradicting the previous claim. ■

Proof of Corollary 1

Proof. The first result follows because ψ^* is invariant to r and ρ. For the second, consider the objective (1), which can be written as

$$\chi = \beta\left(p_1(\psi v_1(\theta_2)) - c\right) + (1-\beta)\,\psi\,(\theta_2 - c)\,.$$

We have

$$\frac{\partial^2 \chi}{\partial\psi\partial\beta} = v_1(\theta_2)p_1'(\psi v_1(\theta_2)) - (\theta_2 - c) < 0.$$

Hence, the second result follows from standard monotone comparative statics arguments. ■

Proof of Proposition 3

Proof. A key difference relative to the setting in Section 3.3 and the one in Section 3.2 is that now the stopping times of θ_n-buyers arriving at a given time $\tau \geq 0$, for $n = 1, 2$, denoted $\tilde{t}_n^\tau$, cannot depend on information about the previous realizations of the price process. Hence, a stopping time $\tilde{t}_n^\tau$ must be adapted to the filtration generated by the price process from date τ onwards.

Consider a high type arriving at τ given that Conditions (3) and (4) hold. If he has not purchased at a date $s > \tau$ and there has been no sale in $[s, \tau)$, and if the price at time s is p_1^*, he obtains a payoff $v_1(p_1^*) = \psi^* v_1(\theta_2)$ from buying immediately. By instead delaying and purchasing at the next "sale", he expects the weakly larger payoff

$$\mathbb{E}\Big[e^{-(r+\rho)(\tilde{t}_2^\tau - s)}\Big|\tilde{t}_2^\tau > s\Big]\,v_1(\theta_2).$$

Hence, by Condition (4), if the buyer elects not to purchase upon arrival at date τ, his payoff from purchasing at some time $s > \tau$ (given that no sale occurs in $[\tau, s]$) is no greater than by purchasing at the next sale. Given Condition (3), it is then incentive compatible for the buyer to purchase on arrival at date τ.

Now, consider a τ such that Condition (3) holds while Condition (4) fails. Hence, there is $s > \tau$ such that $\Omega^s \equiv \{\omega | \tilde{t}_2^\tau(\omega) > s\}$ satisfies $\mathcal{P}(\Omega^s) > 0$ and $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau - s)}\big|\Omega^s\big] < \psi^*$. Then, since $\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau - \tau)}\big|\tilde{t}_2^\tau > \tau\big] = \mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau - \tau)}\big] = \psi^*$, we have

$$\psi^* = (1 - \mathcal{P}(\Omega^s))\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau - \tau)}\big|\Omega\backslash\Omega^s\big] + \mathcal{P}(\Omega^s)e^{-(r+\rho)(s-\tau)}\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau - s)}\big|\Omega^s\big].$$

So, necessarily,

$$(1 - \mathcal{P}(\Omega^s))\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau - \tau)}\big|\Omega\backslash\Omega^s\big] > \Big(1 - \mathcal{P}(\Omega^s)e^{-(r+\rho)(s-\tau)}\Big)\psi^*.$$

Then, the payoff of a high type arriving at τ and purchasing at the next sale or at date s, whichever comes first, is

$$\overbrace{(1-\mathcal{P}(\Omega^s))\mathbb{E}\big[e^{-(r+\rho)(\tilde{t}_2^\tau-\tau)}\big|\Omega\backslash\Omega^s\big]v_1(\theta_2)}^{>\left(1-\mathcal{P}(\Omega^s)e^{-(r+\rho)(s-\tau)}\right)\psi^* v_1(\theta_2)}+\overbrace{\mathcal{P}(\Omega^s)e^{-(r+\rho)(s-\tau)}v_1(p_1^*)}^{=\mathcal{P}(\Omega^s)e^{-(r+\rho)(s-\tau)}\psi^* v_1(\theta_2)},$$

which is strictly higher than $v_1(p_1^*)$. This shows that purchasing immediately with probability one is not an incentive-compatible strategy for the high type arriving at date τ. ■

Proof of Corollary 2

For any $\tau, s \in \mathbb{R}_+$ satisfying $\tau \leq s$, and any $x > 0$, we have that the probability of a sale by date $s + x$, given no sale in $[\tau, s)$ is

$$\Pr\left(s \leq \tilde{t}_2^\tau \leq s + x | \tilde{t}_2^\tau \geq s\right) = \frac{\int_{s-\tau}^{s+x-\tau}(1-F(y))\,dy}{\int_{s-\tau}^{\infty}(1-F(y))\,dy}.$$

Differentiating with respect to s, we find that this is weakly increasing in s if and only if

$$\frac{(F(s-\tau)-F(s+x-\tau))\int_{s-\tau}^{\infty}(1-F(y))\,dy+(1-F(s-\tau))\int_{s-\tau}^{s+x-\tau}(1-F(y))\,dy}{\left(\int_{s-\tau}^{\infty}(1-F(y))\,dy\right)^2} \geq 0,$$

or equivalently,

$$\int_{s-\tau}^{\infty}\frac{(y-(s-\tau))\,f(y)}{1-F(s-\tau)}dy \geq \int_{s+x-\tau}^{\infty}\frac{(y-(s+x-\tau))\,f(y)}{1-F(s+x-\tau)}dy.$$

This inequality must be satisfied if F satisfies DMRL. Conversely, if F satisfies SIMRL, then the previous inequality fails to hold for all τ, s and $x > 0$, which implies that $\Pr\left(s \leq \tilde{t}_2^\tau \leq s + x | \tilde{t}_2^\tau \geq s\right)$ is strictly decreasing in s.

Finally, suppose F does not satisfy NBUE. We want to compute $\mathbb{E}[e^{-(r+\rho)\tilde{t}_2^\tau}|\tilde{t}_2^\tau > s]$, for some fixed τ and $s \geq \tau$. This takes the form

$$\mathbb{E}[e^{-(r+\rho)(\tilde{t}_2^\tau-s)}|\tilde{t}_2^\tau > s] = \frac{\int_{s-\tau}^{\infty}(1-F(y))\,e^{-(r+\rho)(y-(s-\tau))}dy}{\int_{s-\tau}^{\infty}(1-F(y))\,dy}.$$

Integrating both the numerator and the denominator by parts we obtain

$$\mathbb{E}[e^{-(r+\rho)(\tilde{t}_2^\tau-s)}|\tilde{t}_2^\tau > s] = \frac{1}{r+\rho}\frac{1-\int_{s-\tau}^{\infty}e^{-(r+\rho)(y-(s-t))}\frac{f(y)}{1-F(s-\tau)}dy}{\int_{s-\tau}^{\infty}y\frac{f(y)}{1-F(s-\tau)}dy-(s-\tau)}.$$

Condition (4) requires that, for all $s \geq \tau$,

$$\mathbb{E}[e^{-(r+\rho)(\tilde{t}_2^\tau - s)}|\tilde{t}_2^\tau > s] \geq \mathbb{E}[e^{-(r+\rho)(\tilde{t}_2^\tau - \tau)}|\tilde{t}_2^\tau > \tau].$$

Or, equivalently, using the expression above,

$$\frac{\int_{s-\tau}^{\infty} y \frac{f(y)}{1-F(s-\tau)} dy - (s-\tau)}{1 - \int_{s-\tau}^{\infty} e^{-(r+\rho)(y-(s-t))} \frac{f(y)}{1-F(s-\tau)} dy} \leq \frac{\int_0^\infty y f(y) dy}{1 - \int_0^\infty e^{-(r+\rho)y} f(y) dy}. \tag{8}$$

Suppose that F fails to satisfy NBUE, i.e. there exists $x > 0$ such that $\int_x^\infty y \frac{f(y)}{1-F(x)} dy - x > \int_0^\infty y f(y) dy$. Then, for any τ and $s = \tau + x$, the inequality (8) is violated provided that r is taken sufficiently large. Nonetheless, because ψ^* varies continuously with β, there exist parameters such that (3) holds, implying the result.

Proof of Proposition 4

Proof. Using a notation consistent with the one used in the main text, for each date $\tau \in [0, T]$ we let $\tilde{t}_2^\tau \in [\tau, T] \cup \{+\infty\}$ be the random time at which the price next drops to θ_2 under the suggested strategy after τ, that is, the sale time where the price is equal to θ_2 after τ. Then, notice that, for any date $\tau \in [0, T)$

$$\begin{aligned}
\mathbb{E}\left[e^{-(r+\rho)(\tilde{t}_2^\tau - \tau)}|\tilde{t}_2^\tau > \tau\right] &= \int_\tau^T \lambda^* e^{-(r+\rho+\lambda^*)(s-t)} ds + e^{-(r+\rho+\lambda^*)(T-\tau)} \psi^* \\
&= \frac{\lambda^*}{\lambda^* + r + \rho}\left(1 - e^{-(r+\rho+\lambda^*)(T-\tau)}\right) + e^{-(r+\rho+\lambda^*)(T-\tau)} \psi^* \\
&= \psi^*.
\end{aligned}$$

Hence, the θ_1-buyers purchase immediately upon arrival at the same price as the one obtained in Proposition 1, and the same profits are obtained from a θ_2-buyer (equal to $\psi^*(\theta_2 - c)$). As a result, the stochastic price path maximizes profits generated from a cohort arriving at any time $\tau \in [0, T)$. On the other hand, buyers arriving at date T have mass zero.

Finally, that the conditions in Propositions 2 are satisfied for any optimal price path follows from the same argument as for that proposition. ■

Proof of Proposition 5

Proof. An argument analogous to Proposition 1 of Garrett (2016) shows that the set of sales dates S is discrete. To see this, suppose not and so, for any $\varepsilon > 0$, we can find a date t and $\varepsilon-$ball around t such that, at three dates t', t'', t''', $t' < t'' < t'''$, in such a ball there is a "sale" in the sense that any low-value buyer in the market purchases. Then, provided ε is sufficiently small, and given that

$\beta \in \left(\underline{\beta}, \bar{\beta}\right)$, omitting any sales between t' and t''', and setting the price equal to $p_1^d(t''' - t)$ at dates $t \in (t', t''')$ (while keeping prices unchanged at all other dates) increases expected profits for every arrival time $\tau \in (t', t''')$, and does so strictly for buyers arriving in (t', t'') (indeed, this follows when ε is sufficiently small because $\psi^* < 1$ and because expected profits for any arrival date, as given by (1), are strictly concave in ψ).

Let $\hat{\Pi}^*$ be the highest payoff that the seller can achieve using non-stochastic price processes. Suppose first that

$$\hat{\Pi}^* > \frac{\gamma}{r}\beta\left(\theta_1 - c\right),$$

implying that $\hat{\Pi}^*$ is not approached as $z \to +\infty$ in (5). In this case, there must exist a finite optimizer z^*. Then, note that, the seller has the option of setting the first sale at time $z > 0$, which implies

$$\hat{\Pi}^* \geq \int_0^z \gamma e^{-r\tau}\left(\beta\big(p_1\big(e^{-(r+\rho)(z-\tau)}v_1(\theta_2)\big) - c\big) + (1-\beta)\, e^{-(r+\rho)(z-\tau)}\left(\theta_2 - c\right)\right) d\tau + e^{-rz}\hat{\Pi}^*.$$

Optimality of $z = z^*$ implies that the derivative of the right hand side of the previous equation equals 0 at $z = z^*$. Rearranging the first order condition, we obtain the following equation

$$\begin{aligned}
&\frac{\gamma}{r}\left(\beta\left(p_1\big(e^{-(r+\rho)z^*}v_1\left(\theta_2\right)\big) - c\right) + (1-\beta)\, e^{-(r+\rho)z^*}\theta_2\right) \\
&= \hat{\Pi}^* \\
&= \frac{\gamma}{r}\int_0^{z^*}\frac{re^{-r\tau}}{1-e^{-rz^*}}\left(\beta\left(p_1\big(e^{-(r+\rho)(z^*-\tau)}v_1\left(\theta_2\right)\big) - c\right) + (1-\beta)\, e^{-(r+\rho)(z^*-\tau)}\left(\theta_2 - c\right)\right) d\tau. \qquad (9)
\end{aligned}$$

Since

$$\beta\big(p_1\left(\psi v_1\left(\theta_2\right)\right) - c\big) + (1-\beta)\,\psi\left(\theta_2 - c\right) \qquad (10)$$

is strictly concave in ψ, we have that the first equality in (9) can hold only for at most two values of z^*. However, notice that the integral in the last expression in (9) is a weighted average over (10) for $\psi \in \left[e^{-rz^*}, 1\right]$, and hence we cannot have $e^{-rz^*} > \psi^*$ (the maximizer of (10)). Indeed this would imply that the first line of (9) is strictly greater than the third. The same logic implies that $e^{-rz^*} < \psi^*$ and indeed z^* is unique.

Alternatively, $\hat{\Pi}^* = \frac{\gamma}{r}\beta\left(\theta_1 - c\right)$ and so $\hat{\Pi}^*$ is approached as $z \to +\infty$ in (5). It is then easy to see that a constant price equal to θ_1 is optimal. If there were an optimizer of (5) equal to $z' < +\infty$, then we would have $e^{-rz'} > \psi^*$ again by the concavity of (10), and again this would violate (9). ■

A.2 Proofs of Results in Sections 4 and 5

Proof of Lemma 2

Proof. Part 1. Suppose there are three types (i.e., $N = 3$) such that $\theta_1 < \bar{u} < \theta_2 + 1$. Let $v_n(p) = \ln(\theta_n - p + 1)$. Note then that

$$\begin{aligned} \frac{-v_1'(p)}{v_1(p)} &= \frac{1}{\ln(\theta_1 - p + 1)(\theta_1 - p + 1)} \\ &< \frac{1}{\ln(\theta_2 - p + 1)(\theta_2 - p + 1)} \\ &= \frac{-v_2'(p)}{v_2(p)} \end{aligned}$$

for $p < \theta_2$, so Assumption A1 holds. Also, Assumption A2 holds because, for any $\theta \geq \theta_2$, any $p \in U$,

$$\frac{d}{d\theta_n}\left(\frac{d^2(\ln(\theta_n - p + 1))/dp^2}{d(\ln(\theta_n - p + 1))/dp}\right) = -\frac{1}{1 - p + \theta_n} < 0.$$

Let $\varepsilon > 0$ and consider a price process P such that the date-ε price follows a non-degenerate distribution, with $P_\varepsilon \in (\theta_3, \theta_2)$ with probability one. Assume that P is such that, with probability one, P_0 is equal to $p_2\left(e^{-(r+\rho)\varepsilon}\mathbb{E}\left[v_2(P_\varepsilon)\right]\right)$ and P_t is equal to θ_1 for all $t \neq 0, \varepsilon$. Then type θ_2 optimally purchases at date zero. When ε is sufficiently small, type θ_1 optimally waits for date ε to purchase since

$$v_1\big(p_2\big(e^{-(r+\rho)\varepsilon}\mathbb{E}[v_2(P_\varepsilon)]\big)\big) < e^{-(r+\rho)\varepsilon}\mathbb{E}\big[v_1(P_\varepsilon)\big].$$

Hence, the skimming property fails to hold.

Part 2. Fix a price process P where each type θ_n pays a sure price when purchasing. Fix also some incentive-compatible purchasing times $(\tilde{t}_n)_{n=1}^N$.

Assume, for the sake of contradiction, that $\Pr(\tilde{t}_m < \tilde{t}_n) > 0$ for some n and m such that $m > n$. Let $\Omega' = \left\{\omega \in \Omega : \tilde{t}_m < \tilde{t}_n\right\}$. Note that incentive compatibility requires both that

$$\mathbb{E}\left[e^{-r\tilde{t}_n} v_n\left(P_{\tilde{t}_n}\right) \middle| \Omega'\right] \geq \mathbb{E}\left[e^{-r\tilde{t}_m} v_n\left(P_{\tilde{t}_m}\right) \middle| \Omega'\right]$$

(otherwise the stopping time $\tilde{t}_n \wedge \tilde{t}_m$ generates a strictly higher payoff for type θ_n) and

$$\mathbb{E}\left[e^{-r\tilde{t}_m} v_m\left(P_{\tilde{t}_m}\right) \middle| \Omega'\right] \geq \mathbb{E}\left[e^{-r\tilde{t}_n} v_m\left(P_{\tilde{t}_n}\right) \middle| \Omega'\right]$$

(otherwise the stopping time $\tilde{t}_n \vee \tilde{t}_m$ generates a strictly higher payoff for type θ_m). Moreover, it requires the "individual rationality" condition for type θ_m that $\Pr\left(P_{\tilde{t}_m} \leq \theta_m | \Omega'\right) = 1$.

The following are then true for almost all $\omega \in \Omega'$. First, by the first incentive-compatibility condition (and the absence of pricing risk), we have $P_{\tilde{t}_n} < P_{\tilde{t}_m}$; hence, given that prices are no less than θ_N by assumption, we must have have $m < N$. Second, by the first two incentive compatibility conditions jointly,

$$\frac{v_m\left(P_{\tilde{t}_m}\right)}{v_m\left(P_{\tilde{t}_n}\right)} \geq \frac{v_n\left(P_{\tilde{t}_m}\right)}{v_n\left(P_{\tilde{t}_n}\right)}.$$

However,

$$\begin{aligned}\frac{v_m\left(P_{\tilde{t}_m}\right)}{v_m\left(P_{\tilde{t}_n}\right)} &= \frac{v_n\left(P_{\tilde{t}_m}\right)}{v_n\left(P_{\tilde{t}_n}\right)} e^{\int_{P_{\tilde{t}_n}}^{P_{\tilde{t}_m}}\left(\left(\frac{-v_n'(p)}{v_n(p)}\right)-\left(\frac{-v_m'(p)}{v_m(p)}\right)\right)dp} \\ &< \frac{v_n\left(P_{\tilde{t}_m}\right)}{v_n\left(P_{\tilde{t}_n}\right)},\end{aligned}$$

with the inequality following by Assumption A1, contradicting the previous observation. ■

Proof of Proposition 6 and Corollary 3

Proof. We divide the proof into five steps.

Step 1. Existence of deterministic price process satisfying Condition 3. We first show that, under Assumption A, for any decreasing sequences $(p_n^*)_{n=1}^N \in \mathbb{R}_+^N$ and $(\psi_n^*)_{n=1}^N \in [0,1]^N$ satisfying the Condition 3 of Proposition 6, there is a deterministic price process where type n purchases at time $t_n^* \equiv -\log(\psi_n^*)/(r+\rho)$ at price p_n^* or $t_n^* = +\infty$ if $\psi_n^* = 0$ (i.e., type θ_n does not purchase). In this suggested deterministic price process, the price is θ_1 except for each time t_n^*, where it is equal to p_n^*.

Notice that, by Lemma 2, we only need to check the incentive compatibility constraints only for consecutive types. To check the downward-mimicking incentives, assume that $n < N$ is such that $\psi_n^* > 0$, the payoff of type n from purchasing at time t_{n+1}^* instead of at time t_n^* is

$$e^{-(r+\rho)t_{n+1}^*}v_n(p_{n+1}^*) = e^{-(r+\rho)t_n^*}\frac{\psi_{n+1}^*}{\psi_n^*}v_n(p_{n+1}^*) = e^{-(r+\rho)t_n^*}v_n(p_n^*)$$

where we used Part 3 of Proposition 6. Hence, type θ_n is indifferent between purchasing at t_n^* and instead waiting for t_{n+1}^*. If, instead, $n = N$, and we have $\psi_N^* > 0$, then $v_N(p_N^*) = 0$ (by Condition 3), so $p_N^* = \theta_N$, and the θ_N-buyer is willing to purchase at time t_N^*.

Now, let us check the upward-mimicking incentives. Consider any $n > 1$ such that $t_n^* < +\infty$ and note that if $p_{n-1}^* \geq \theta_n$ then type θ_n does not gain by deviating upwards. If $p_{n-1}^* < \theta_n$, then by Assumption A1,

$$\begin{aligned}\frac{\frac{v_{n-1}\left(p_{n-1}^*\right)}{v_{n-1}\left(p_n^*\right)}}{\frac{v_n\left(p_{n-1}^*\right)}{v_n\left(p_n^*\right)}} &= e^{\int_{p_n^*}^{p_{n-1}^*}\left(\left(\frac{-v_n'(p)}{v_n(p)}\right)-\left(\frac{-v_{n-1}'(p)}{v_{n-1}(p)}\right)\right)dp} \\ &> 1,\end{aligned}$$

which implies (using Condition 3) that

$$
\begin{aligned}
e^{-(r+\rho)t^*_{n-1}} v_n(p^*_{n-1}) &< e^{-(r+\rho)t^*_{n-1}} \frac{v_{n-1}\left(p^*_{n-1}\right) v_n\left(p^*_n\right)}{v_{n-1}\left(p^*_n\right)} \\
&= e^{-(r+\rho)t^*_{n-1}} \frac{\psi^*_n v_n\left(p^*_n\right)}{\psi^*_{n-1}} \\
&= e^{-(r+\rho)t^*_n} v_n\left(p^*_n\right).
\end{aligned}
$$

Hence, type θ_n prefers to purchase at t^*_n than at t^*_{n-1}.

Step 2. Necessity of the conditions for optimality of a deterministic price process. We now prove that a deterministic price process P that does not satisfy the conditions in Proposition 6 is not optimal. Notice first that if a deterministic price process induces purchases at prices $(p_n)_{n=1}^N \in \mathbb{R}_+^N$ and times $(t_n)_{n=1}^N \in (\mathbb{R}_+ \cup \{+\infty\})^N$, then necessarily the sequences $(p_n)_{n=1}^N$ and $(e^{-(r+\rho)t_n})_{n=1}^N$ are (weakly) decreasing. Incentive compatibility implies $e^{-(r+\rho)t_{n+1}} v_n(p_{n+1}) \leq e^{-(r+\rho)t_n} v_n(p_n)$ for all n (with the notation $t_{N+1} = +\infty$).

Conditions 1 and 2 of the proposition are satisfied by definition of a deterministic price path, except for the claim that $\psi_1 = 1$. Nevertheless, if $\psi_1 = 0$ we have that profits are equal to 0 (which is dominated by offering θ_1 at all times), and if $\psi_1 \in (0,1)$ (so $t_1 \in (0,\infty)$) then a deterministic price process inducing purchases at prices $(p_n)_{n=1}^N$ and times $(t_n - t_1)_{n=1}^N$ exists and gives the seller higher profits than than P. Hence, if the price process does not satisfy the conditions in Proposition 6, then necessarily Condition 3 is not satisfied for some n, that is, $e^{-(r+\rho)t_{n+1}} v_n(p_{n+1}) < e^{-(r+\rho)t_n} v_n(p_n)$ for some n.

Let $\bar{n}$ be the highest value of n with this property. Assume first that $t_{\bar{n}+1} < +\infty$. Now, define a deterministic price process $\hat{P}$ such that the price is equal to θ_1 for all times except $(\hat{t}_n)_{n=1}^N$ as follows. For all $n \leq \bar{n}$, $\hat{P}_{\hat{t}_n} = p_n$, and $\hat{t}_n = t_n$. Furthermore, $\hat{t}_{\bar{n}+1}$ satisfies

$$
e^{-(r+\rho)\hat{t}_{\bar{n}}} v_{\bar{n}}(p_{\bar{n}}) = e^{-(r+\rho)\hat{t}_{\bar{n}+1}} v_{\bar{n}}(p_{\bar{n}+1}).
$$

Finally, $\hat{t}_n = t_n + \hat{t}_{\bar{n}+1} - t_{\bar{n}+1}$ for all $n > \bar{n}+1$. It is clear that the adjacent incentive constraints (as considered above) are satisfied for all types. This follows from noticing that $\bar{\theta}_n$ is indifferent between purchasing at $\hat{t}_{\bar{n}}$ and at $\hat{t}_{\bar{n}+1}$ and that $\hat{t}_{\bar{n}+1} > \hat{t}_{\bar{n}}$ with $p_{\bar{n}+1} < p_{\bar{n}}$, so that, by Assumption A1, type $\theta_{\bar{n}+1}$ strictly prefers to purchase at $\hat{t}_{\bar{n}+1}$ than at $\hat{t}_{\bar{n}}$. Hence, by Lemma 2, it is incentive compatible for each type θ_n to purchase at time $\hat{t}_n$. Also, clearly, $\hat{t}_{\bar{n}+1} < t_{\bar{n}+1}$, so $\hat{t}_n < t_n$ for all $n \geq \bar{n}+1$. Given the definition of $\bar{n}$ and $t_{\bar{n}+1} < +\infty$, the lowest price charged is at least $\theta_N > c$. Hence, the seller's profits under the new price process are strictly higher than under the original.

Assume now that $t_{\bar{n}+1} = +\infty$. Then, the deterministic price process $\hat{P}$ can be defined instead as setting a price equal to θ_1 at all times except at $(t_n)_{n=1}^{\bar{n}}$, where the price charged is p_n for all $n < \bar{n}$ and and $\theta_{\bar{n}} > p_{\bar{n}}$ at time $t_{\bar{n}}$. In this case, each buyer is instructed to buy at his earliest optimal time, which for all types with $t_n < t_{\bar{n}}$ remains t_n, and for types such that $t_n = t_{\bar{n}}$ is either $t_{\bar{n}}$ or $t_{\bar{n}-1}$.

As a result, under the new price process, all types buy either at the same or at an earlier time, and all pay a weakly higher price, which is strictly higher for type $\theta_{\bar{n}}$. Again, prices are always strictly greater than c, so the seller's profits strictly increase.

Step 3. Necessity of the conditions for stochastic price processes. We now consider the necessity of Conditions 1-3 for stochastic price processes in general. Fix a stochastic price process P with corresponding purchasing times $(\tilde{t}_n)_{n=1}^N$. Define, for each $n = 1, ..., N$,

$$t_n^d \equiv -\frac{\log(\mathbb{E}[e^{-(\rho+r)\tilde{t}_n}])}{r+\rho} \in \mathbb{R}_+ \cup \{+\infty\}$$

(where we adopt the convention that $t_n^d = +\infty$ in case $\mathbb{E}[e^{-(\rho+r)\tilde{t}_n}] = 0$).

Suppose first that the second condition is satisfied. In this case, for every type θ_n with $t_n^d < +\infty$, we can associate the degenerate purchase price p_n^*, from the price process P. Then, define a deterministic price process $\hat{P}$ by setting price p_n^* at any date $t_n^d < +\infty$ (noting, by incentive compatibility of $\hat{P}$, that $t_n^d = t_{n'}^d$ implies $p_n^* = p_{n'}^*$). At all other times, the price is set at θ_1. Under $\hat{P}$, it is incentive compatible for all types θ_n to purchase at date t_n^d (or never if $t_n^d = +\infty$). This follows from noting that, if θ_n strictly prefers to purchase at time $t_{n'}^d \neq t_n^d$ for some index n', then type θ_n strictly gains by mimicking type $\theta_{n'}$ under P; i.e., $\tilde{t}_n$ is not incentive compatible under P. The price process $\hat{P}$ generates the same expected profits as P. Using the definition of t_n^d, it is then immediate that if P fails Condition 1 or 3, then the same is true for the deterministic process $\hat{P}$, which generates less than optimal profits by Step 2. Hence, the process P is not optimal.

Now we show that the second condition must be satisfied if P is optimal. Suppose for a contradiction that P is optimal but the condition does not hold. That is, there exists a type θ_n such that there is no p_n^* satisfying $\Pr\left((P_{\tilde{t}_n} \neq p_n^*) \wedge (\tilde{t}_n < +\infty)\right) = 0$. Let $\bar{n}$ denote the highest index such that $t_{\bar{n}}^d < +\infty$ (which exists, since P is profit maximizing). It is then clear that $t_n^d < +\infty$ for all $n \leq \bar{n}$ and $t_n^d = +\infty$ for all $n > \bar{n}$.

For each type θ_n with $n < \bar{n}$, let $\hat{p}_{t_n^d}$ be defined by

$$e^{-(r+\rho)t_n^d} v_n\big(\hat{p}_{t_n^d}\big) = \mathbb{E}\left[e^{-(r+\rho)\tilde{t}_n} v_n\left(P_{\tilde{t}_n}\right)\right],$$

so that type θ_n, by purchasing at price $\hat{p}_{t_n^d}$ at date t_n^d, earns the same payoff as under the original price process. Let $\hat{p}_{t_{\bar{n}}^d} = \theta_{\bar{n}}$ so that type $\theta_{\bar{n}}$ purchasing at this price earns payoff zero (i.e., weakly less than the expected payoff under P). It is then clear that, since v_n is strictly concave, $e^{-(r+\rho)t_n^d}\big(\hat{p}_{t_n^d} - c\big) \geq \mathbb{E}\big[e^{-(r+\rho)\tilde{t}_n}\big(P_{\tilde{t}_n} - c\big)\big]$ with a strict inequality in case $P_{\tilde{t}_n}$ is non-degenerate on $\big\{\omega \in \Omega : \tilde{t}_n < +\infty\big\}$.

Let $J \in \mathbb{N}$ and let $n_J = \bar{n}$ be the index for the lowest type such that $t_{n_J}^d < +\infty$. Proceeding inductively, once n_{j+1} is fixed, let n_j be highest index such that $n_j < n_{j+1}$ and $e^{-(r+\rho)t_{n_j}^d}\big(\hat{p}_{t_{n_j}^d} - c\big) > e^{-(r+\rho)t_{n_{j+1}}^d}\big(\hat{p}_{t_{n_{j+1}}^d} - c\big)$. We thus obtain a strictly-increasing finite sequence of indices with cardinality J, $(n_1, \ldots, n_J)$. Moreover, $e^{-(r+\rho)t_{n_j}^d}\big(\hat{p}_{t_{n_j}^d} - c\big)$ is strictly decreasing with j. Let M

denote the set of indices n_j.

Now, note that, for any $n_j \in M$, a buyer of type θ_{n_j} prefers to purchase at date $t^d_{n_j}$ and price $\hat{p}_{t^d_{n_j}}$ than at $t^d_{n_{j'}}$ and price $\hat{p}_{t^d_{n_{j'}}}$ for any $j' > j$. To see this, note that, by purchasing at date $t^d_{n_j}$, type θ_{n_j} earns the same payoff as under the original process P. If $j' < J$, then the claim follows by Assumption A2, which means that type $\theta_{n_{j'}}$ is more risk averse than θ_{n_j}. This implies $e^{-(r+\rho)t^d_{n_{j'}}} v_{n_j}\big(\hat{p}_{t^d_{n_{j'}}}\big) \le \mathbb{E}\big[e^{-(r+\rho)\tilde{t}_{n_{j'}}} v_{n_j}\big(P_{\tilde{t}_{n_{j'}}}\big)\big]$, where the latter is the expected payoff for θ_{n_j} when mimicking $\theta_{n_{j'}}$ under the original price process P. If instead $j' = J$, the claim follows since $P_{\tilde{t}_{n_J}} \le \theta_{n_J}$ with probability one on $\big\{\omega \in \Omega : \tilde{t}_{n_J} < +\infty\big\}$ (while $\hat{p}_{t^d_{n_J}} = \theta_{n_J}$).

Next, note that if $n_j, n_{j'} \in M$ with $n_j < n_{j'}$, then $t^d_{n_j} < t^d_{n_{j'}}$. Otherwise, $t^d_{n_j} \ge t^d_{n_{j'}}$ and by the previous observation, $\hat{p}_{t^d_{n_j}} \le \hat{p}_{t^d_{n_{j'}}}$. But this implies (given that profits at all offered prices are positive) $e^{-(r+\rho)t^d_{n_j}}\big(\hat{p}_{t^d_{n_j}} - c\big) \le e^{-(r+\rho)t^d_{n_{j'}}}\big(\hat{p}_{t^d_{n_{j'}}} - c\big)$, which contradicts the choice of M. We can then define a deterministic price process $\hat{P}$ by price $\hat{p}_{t^d_{n_j}}$ at date $t^d_{n_j}$ for $n_j \in M$, and price θ_1 at all other dates, noting that the dates $t^d_{n_j}$ are increasing with n_j.

Now, note that, under $\hat{P}$, for all $n_j \in M$, type θ_{n_j} purchases weakly earlier than $t^d_{n_j}$, generating weakly higher profits for the seller than the expected profits under P (and strictly higher in case $P_{\tilde{t}_{n_j}}$ is non-degenerate on $\big\{\omega \in \Omega : \tilde{t}_{n_j} < +\infty\big\}$). For $n \notin M$, we have two possibilities. If $n > \bar{n}$, then whether θ_n purchases or not, it is easily verified that profits are weakly higher than expected profits under P. Alternatively, $n < \bar{n}$ and, by Lemma 2, θ_n purchases at a date no later than $t^d_{n_j}$ where n_j is the smallest index greater than n in M. Profits are then no less than $e^{-(r+\rho)t^d_{n_j}}\big(\hat{p}_{t^d_{n_j}} - c\big)$, which are strictly positive and no less than $e^{-(r+\rho)t^d_n}\big(\hat{p}_{t^d_n} - c\big)$. The latter is at least the expected profits for type θ_n under P, and strictly higher in case $P_{\tilde{t}_n}$ is non-degenerate on the positive probability set $\big\{\omega \in \Omega : \tilde{t}_n < +\infty\big\}$.

Step 4. Existence an optimal sequences: For arbitrary pairs of decreasing sequences $(p_n, \psi_n)_{n=1}^N$ satisfying Conditions 1-3, the seller's expected profits per buyer may be written

$$\Pi = \sum_{n=1}^{N} \beta_n \psi_n (p_n - c) \ .$$

Since this is continuous in each component, and the set of sequences that satisfy the conditions is compact, an optimal pair of sequences $(p^*_n, \psi^*_n)_{n=1}^N$ exists.

Step 5. Uniqueness of optimal sequences: Assume that there are two different (decreasing) sequences $(p_n, \psi_n)_{n=1}^N$ and $(p'_n, \psi'_n)_{n=1}^N$ maximizing the profits of the seller. Let P and P' be two corresponding deterministic price processes, and let Π be the (same) value of the seller's expected profits. Also, let $(t_n)_{n=1}^N$ and $(t'_n)_{n=1}^N$ denote the purchasing times (i.e., $t_n = -\frac{\log(\psi_n)}{r+\rho}$ and $t'_n = -\frac{\log(\psi'_n)}{r+\rho}$ for all n). Without loss of generality, let prices be equal to θ_1 at all other times (so buyers are willing not to purchase at these times). Let n^* be the minimum index such that $(p_{n^*}, \psi_{n^*}) \neq$

(p'_{n^*}, ψ'_{n^*}). Assume that $\psi_{n^*} \geq \psi'_{n^*}$, so the corresponding purchasing times are such that $t_{n^*} \leq t'_{n^*}$.

Now, consider a new (stochastic) price process $\hat{P}$ where the price equals $(p_n)_{n=1}^{n^*-1}$ at times $(t_n)_{n=1}^{n^*-1}$ with certainty. At time t_{n^*}, the price offered is $\hat{p}_{n^*}$ that solves

$$\psi_{n^*} v_{n^*}(\hat{p}_{n^*}) = \tfrac{1}{2}\psi_{n^*} v_{n^*}(p_{n^*}) + \tfrac{1}{2}\psi'_{n^*} v_{n^*}(p'_{n^*}).$$

Finally, let $\varepsilon > 0$ and let the continuation prices after t_{n^*} be given by $(p_n)_{n=n^*+1}^{N}$ at times

$$\Big(\frac{-\log\left(e^{-(r+\rho)\varepsilon}\psi_n\right)}{r+\rho}\Big)_{n=n^*+1}^{N}$$

(with price set equal to θ_1 at other times) with probability $\frac{1}{2}$. With complementary probability, the continuation prices are given by $(p'_n)_{n=n^*+1}^{N}$ at times

$$\Big(\frac{-\log\left(e^{-(r+\rho)\varepsilon}\psi'_n\right)}{r+\rho}\Big)_{n=n^*+1}^{N}$$

(with price set equal to θ_1 at other times).

Let $\bar{p}$ be given by

$$\left(\tfrac{1}{2}\psi_{n^*} + \tfrac{1}{2}\psi'_{n^*}\right) v_{n^*}(\bar{p}) = \tfrac{1}{2}\psi_{n^*} v_{n^*}(p_{n^*}) + \tfrac{1}{2}\psi'_{n^*} v_{n^*}(p'_{n^*}),$$

and notice that optimality of P and P' imply $\bar{p} > c$. Because the buyer with type θ_{n^*} is risk averse, we have

$$\bar{p} - c \geq \frac{\psi_{n^*}}{\psi_{n^*} + \psi'_{n^*}}\left(p_{n^*} - c\right) + \frac{\psi'_{n^*}}{\psi_{n^*} + \psi'_{n^*}}\left(p'_{n^*} - c\right),$$

with strict inequality if $p_{n^*} \neq p'_{n^*}$. This implies that

$$2\psi_{n^*}\left(\bar{p} - c\right) \geq \left(\psi_{n^*} + \psi'_{n^*}\right)\left(\bar{p} - c\right) \geq \psi_{n^*}\left(p_{n^*} - c\right) + \psi'_{n^*}\left(p'_{n^*} - c\right),$$

where the first inequality is strict if $\psi_{n^*} > \psi'_{n^*}$ and the second is strict if $p_{n^*} \neq p'_{n^*}$. Since $\hat{p}_{n^*} \geq \bar{p}$, we can conclude that

$$\psi_{n^*}\left(\hat{p}_{n^*} - c\right) > \frac{\psi_{n^*}\left(p_{n^*} - c\right) + \psi'_{n^*}\left(p'_{n^*} - c\right)}{2}.$$

Given that type θ_{n^*} finds it incentive compatible to purchase at time t_{n^*} at price $\hat{p}_{n^*}$ instead of waiting, the new price process generates the same profits for types $\theta_1, \ldots, \theta_{n^*-1}$, *strictly* higher profit than the average of profits from θ_{n^*} across P and P' (by an amount independent of ε), and expected profits from types $\theta_{n^*+1}, \ldots, \theta_N$ that approach the average from P and P' as $\varepsilon \searrow 0$. Hence, for sufficiently small ε, profits increase, a contradiction. ■

Proof of Proposition 7

Proof. Let $(p_n^*, \psi_n^*)_{n=1}^N$ be the unique optimal sequence characterized in Proposition 6. Let $(n_j)_{j=1}^J$ be the (unique) increasing sequence containing all indices satisfying $\psi_{n_j}^* > \psi_{n_j+1}^*$, with the convention that $\psi_{N+1}^* = 0$.

Consider the following price process, characterized as a process with $J-1$ states $\{\sigma_j\}_{j=1}^{J-1}$ (analogous to the one described in the main text for three types) for some values $\Lambda \in \mathbb{R}_{++}$ and $(\mu_j^\Lambda)_{j=2}^{J-1} \in (0,1]^{J-2}$ to be determined. Initializing the state at $\tau = 0$ to σ_1, the price process is characterized by:

1. In state σ_1 the price is $p_{n_1}^*$, and the state changes to state σ_2 at rate λ_2^* satisfying
$$v_{n_1}(p_{n_1}^*) = \frac{\lambda_2^*}{\lambda_2^* + r + \rho} v_{n_1}(p_{n_2}^*) \ .$$
2. In state σ_j, for $j = 2, ..., J-2$, the price is $p_{n_j}^*$. At rate $\mu_j^\Lambda \Lambda$ the state changes to state σ_1. At rate $(1-\mu_j^\Lambda)\Lambda$, the state changes to σ_{j+1}.
3. In state σ_{J-1}, the price is $p_{n_{J-1}}^*$, and the state changes to state σ_1 at rate $\mu_{J-1}^\Lambda \Lambda$ and, at rate $(1-\mu_{J-1}^\Lambda)\Lambda$ the price $p_{n_J}^*$ is offered and the state subsequently changes to σ_1.

Let ψ_j^Λ denote the expected discounting until the price is $p_{n_j}^*$ if the current state is σ_1 (i.e., it is equal to $\mathbb{E}\big[e^{-(r+\rho)\tilde{s}_{n_j}} \big| \sigma_1\big]$, where $\tilde{s}_{n_j}$ is the random time until price $p_{n_j}^*$ is charged when in state σ_1). As in the example in the main text, we require any type θ_{n_j}, $j = 2, ..., J-1$, to be indifferent between purchasing at price $p_{n_j}^*$ at state σ_j and waiting for the price to be $p_{n_{j+1}}^*$. That is

$$v_{n_j}(p_{n_j}^*) = \left(\frac{\Lambda \mu_j^\Lambda}{r+\rho+\Lambda} \psi_{j+1}^\Lambda + \frac{\Lambda(1-\mu_j^\Lambda)}{r+\rho+\Lambda} \right) v_{n_j}(p_{n_{j+1}}^*). \tag{11}$$

For each $j = 2, ..., J$, ψ_j^Λ solves

$$\psi_j^\Lambda = \begin{cases} \frac{\lambda_2^*}{\lambda_2^*+r+\rho} & \text{if } j = 2, \\ \psi_{j-1}^\Lambda \left(\frac{\Lambda \mu_j^\Lambda}{r+\rho+\Lambda} \psi_j^\Lambda + \frac{\Lambda(1-\mu_j^\Lambda)}{r+\rho+\Lambda} \right) & \text{if } j \in \{3, ..., J\}. \end{cases} \tag{12}$$

Now using (11) and the third condition of Proposition 6, we have $\psi_2^\Lambda = \psi_2^*$ and

$$\frac{\Lambda \mu_j^\Lambda}{r+\rho+\Lambda} \psi_{j+1}^\Lambda + \frac{\Lambda(1-\mu_j^\Lambda)}{r+\rho+\Lambda} = \frac{\psi_{n_{j+1}}^*}{\psi_{n_j}^*}.$$

As a result, using (12), we have $\frac{\psi_j^\Lambda}{\psi_{j-1}^\Lambda} = \frac{\psi_{n_j}^*}{\psi_{n_{j-1}}^*}$ for each $j = 3, ..., J$, so, since $\psi_2^\Lambda = \psi_2^*$, we have $\psi_j^\Lambda = \psi_{n_j}^*$ for all $j = 2, ..., J$.

Now, suppose buyers of type θ_n purchase as soon as the price reaches p_n^* (as is incentive compatible). Then, the expected profits from type θ_n arriving at a time τ when the state is σ_1 are $\psi_n^*(p_n^* - c)$ (as in the fixed arrival case). As $\Lambda \to +\infty$, the probability that the state at time τ is σ_1 becomes arbitrarily close to 1, uniformly over $\tau \geq 0$. This implies the result. ■

Proof of Proposition 8

Proof. Part 1. This is a consequence of the following lemma, which ensures, for any revelation mechanism $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$, the existence of a deterministic price process (and hence revelation mechanism) that achieves the same expected payoffs as the original revelation mechanism both for the seller and all buyer types. (Given the following result, the construction of such a deterministic price process is the same as described in Section 4.1.)

Lemma 4 *Assume that Assumptions A1 and 12* hold. Consider a profit-maximizing revelation mechanism $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$. There are two (weakly) decreasing sequences, $(p_n^*)_{n=1}^N \in \mathbb{R}_+^N$ and $(\psi_n^*)_{n=1}^N \in [0,1]^N$ such that, for any n:*

1. *The purchasing time $\tilde{t}_n$ of a θ_n-buyer satisfies $\mathbb{E}[e^{-(r+\rho)\tilde{t}_n}] = \psi_n^*$ and $\psi_1^* = 1$.*
2. *If a type θ_n buys the good, he pays a certain price p_n^*; that is, $\Pr\left((P^n \neq p_n^*) \wedge (\tilde{t}_n < +\infty)\right) = 0$.*
3. *Downward incentive constraints bind; i.e., $\psi_n^* v_n(p_n^*) = \psi_{n+1}^* v_n(p_{n+1}^*)$ for all n, where $\psi_{N+1}^* \equiv 0$.*

Proof of Lemma 4. Consider a profit-maximizing revelation mechanism $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$ (with a suitable probability space) such that Property 2 holds, but 1 or 3 fails. Let p_n^* the purchasing price of the θ_n-buyer. As in the proof of Proposition 6 (Step 3), we can let, for each $n = 1, ..., N$,

$$t_n^d \equiv -\frac{\log(\mathbb{E}[e^{-(\rho+r)\tilde{t}_n}])}{r+\rho} \in \mathbb{R}_+ \cup \{+\infty\}.$$

Notice that, by the incentive compatible constraints, $t_n^d = t_{n'}^d$ implies $p_n^* = p_{n'}^*$. We can then define a deterministic price process $\hat{P}$ (noting that this defines, as described in the main text, a revelation mechanism) by setting price p_n^* at any date $t_n^d < +\infty$ and price θ_1 otherwise, and note that this price process induces purchase by types θ_n on dates t_n^d when $t_n^d < +\infty$ (and no purchase otherwise), generating profits equal to the expected profits in the revelation mechanism. However, this deterministic price process fails Conditions 1 or 3 of Proposition 6, implying the existence of a price process with higher expected profits. Hence, the revelation mechanism $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$ cannot be profit maximizing.

The proof that Property 2 must hold also follows closely Step 3 in the proof of Proposition 6. Again, we can define t_n^d as above. Let $\bar{n}$ denote the highest index such that $t_{\bar{n}}^d < +\infty$ (which exists, since $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$ is profit maximizing). For each type θ_n with $n < \bar{n}$, let $\hat{p}_{t_n^d}$ be defined by

$$e^{-(r+\rho)t_n^d} v_n\left(\hat{p}_{t_n^d}\right) = \mathbb{E}\left[e^{-(r+\rho)\tilde{t}_n} v_n\left(P^n\right)\right],$$

so that type θ_n, by purchasing at price $\hat{p}_{t_n^d}$ at date t_n^d, earns the same payoff as under the original revelation mechanism $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$. Let $\hat{p}_{t_{\bar{n}}^d} = \theta_{\bar{n}}$ so that type $\theta_{\bar{n}}$ purchasing at this price earns payoff zero (i.e., weakly less than the expected payoff under $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$). It is then clear that, since v_n is strictly concave, $e^{-(r+\rho)t_n^d}\left(\hat{p}_{t_n^d} - c\right) \geq \mathbb{E}\left[e^{-(r+\rho)\tilde{t}_n}\left(P^n - c\right)\right]$, with a strict inequality in case P^n is non-degenerate on $\left\{\omega \in \Omega : \tilde{t}_n < +\infty\right\}$.

One can then construct the index set M and deterministic price process process $\hat{P}$ precisely as in Step 3 in the proof of Proposition 6 (recalling that this also determines a revelation mechanism). We can then note that, for any $n_j \in M$, a buyer of type θ_{n_j} prefers to purchase at date $t_{n_j}^d$ and price $\hat{p}_{t_{n_j}^d}$ than at $t_{n_{j'}}^d$ and price $\hat{p}_{t_{n_{j'}}^d}$ for any $j' > j$. To see this, note that, by purchasing at date $t_{n_j}^d$, type θ_{n_j} earns the same payoff as under the original process P. For any $j' > j$, because $n_{j'} > n_j$ (by construction of M), the claim follows by Assumption A2*. In particular, since type $\theta_{n_{j'}}$ is more risk averse than θ_{n_j}, $e^{-(r+\rho)t_{n_{j'}}^d} v_{n_j}\left(\hat{p}_{t_{n_{j'}}^d}\right) \leq \mathbb{E}\left[e^{-(r+\rho)\tilde{t}_{n_{j'}}} v_{n_j}\left(P^{n_{j'}}\right)\right]$, where the latter is the expected payoff for θ_{n_j} when mimicking $\theta_{n_{j'}}$ under the revelation mechanism $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$.

The rest of the proof follows the same lines as in Step 3 in the proof of Proposition 6, establishing that profits strictly increase under the price process $\hat{P}$. As noted in the main text, the key difference in the proof, and the reason we need Assumption A2* rather than Assumption A2 relates to the argument in the previous paragraph. Because an incentive-compatible and individually-rational revelation mechanism $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$ need not satisfy individual rationality ex-post, a type θ_N need not be charged a price less than θ_N under such mechanisms. As a result, assuming such types purchase under $\left(\tilde{t}_n, P^n\right)_{\theta_n \in \Theta}$, replacing the random price P^N, when $\tilde{t}_N < +\infty$, with the sure price θ_N does not necessarily make mimicry of type θ_N by higher types less attractive. This is true, in particular, if θ_N is less risk averse than higher types, as ruled out by Assumption A2*. ■

Part 2. This can be viewed as a result of the example in Maskin and Riley (1984, p. 1503), which we modify slightly here to fit the assumptions of our model. Let $\theta_2 > 0$ and $\theta_1 \in (\theta_2, \theta_2 + 1)$. Let $\varepsilon \in (0, \theta_1 - \theta_2)$, $U = [0, 1 + \theta_1 - \varepsilon]$, $v_1(p) = \log(1 + \theta_1 - p)$ and $v_2(p) = \phi(\theta_2 - p)$ for some function ϕ that is strictly increasing and concave on $\{\theta_2 - p : p \in U\}$.

Now, let $\eta \in \left(0, \frac{\theta_2}{1+\theta_1-\varepsilon}\right)$. Consider a revelation mechanism that is static (as in Maskin and Riley) and puts $\Omega = \{\omega_a, \omega_b\}$, with $\tilde{t}_1(\omega) = \tilde{t}_2(\omega) = 0$ for all $\omega \in \Omega$. Specify payments by $P^2(\omega_a) = 0$, $P^2(\omega_b) = 1 + \theta_1 - \varepsilon$, and $P^1(\omega_a) = P^1(\omega_b) = \theta_1$. Finally, put $\Pr(\{\omega_b\}) = \frac{\theta_2}{1+\theta_1-\varepsilon} - \eta$ and $\Pr(\{\omega_a\}) = 1 - \Pr(\{\omega_b\})$.

Now, note that, taking ϕ sufficiently close to the identity, uniformly over $\{\theta_2 - p : p \in U\}$, there exists $\eta > 0$ such that the expected payoff of type θ_2 from reporting truthfully is equal to zero. Indeed, such η vanishes as ϕ approaches the identity. For this value of η, both types expect payoff zero from reporting truthfully. However, provided ε is chosen sufficiently small, the expected payoff to type θ_1 from mimicking type θ_2 is negative. Hence, the revelation mechanism defined here is incentive compatible and attains profits that approach first best as ϕ approaches the identity. (Maskin and Riley's example posits that the low type is risk neutral, and hence profits equal to the first best are attained.)

Conversely, our results in Section 3.1 imply that the optimal stochastic price process exhibits no price uncertainty and the seller's expected profits are bounded away from the first best. ■