A Dynamic Theory of Random Price Discounts*

Francesc Dilmé[†] **Daniel Garrett**[‡]

February, 2021

Abstract

A seller with commitment power sets prices over time. Risk-averse buyers arrive to the market and decide when to purchase. We obtain that the optimal price path is a "regular" price, with occasional episodes of sequential discounts that occur at random times. The optimal price path has the property that the price a buyer ends up paying is independent of his arrival and purchase times, and only depends on his valuation. Our theory accommodates empirical findings on the timing of discounts.

JEL classification: D82

Keywords: dynamic pricing, sales, random mechanisms

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. Such patterns have been noticed across a range of empirical work; e.g., Warner and Barsky (1995), Pesendorfer (2002), Eichenbaum, Jaimovich, and Rebelo (2011), Kehoe and Midrigan (2015), and Chevalier and Kashyap (2017).

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 some dates, then the price

*For helpful discussions and comments, we would like especially to thank Pierre Dubois, Drew Fudenberg, Matthew Gentry, Holger Gerhardt, Paul Heidhues, Stephan Lauermann, Volker Nocke, Alessandro Pavan, Martin Pesendorfer, Anton Sobolev, Philipp Strack, and Curtis Taylor. We would also like to thank seminar participants at the CSIO/IDEI joint conference (2018; Toulouse), CRC TR 224 Internal Conference Bonn-Mannheim, EARIE in Maastricht, U. of Bath, U. of Mannheim, UNC Chapel Hill, and the Workshop on Dynamic Pricing, Santiago de Chile (2017). This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 714147). This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through grant CRC TR 224 and under Germany's Excellence Strategy — EXC 2126/1– 390838866 and EXC-2047/1 – 390685813.

[†]**University of Bonn. fdilme@uni-bonn.de**

[‡]**Toulouse School of Economics. dfgarrett@gmail.com**

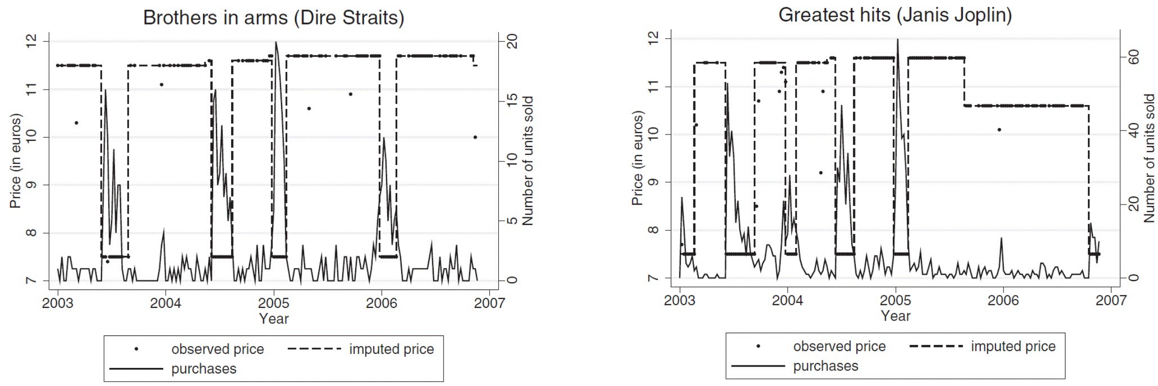

(a) (b)

Figure 1: Illustration of typical price and quantity patterns (Figure 1 from Février and Wilner, 2016) for two albums – 'Brothers in arms' by Dire Straits, and 'Greatest hits' by Janis Joplin – featuring two focal prices (a high price and a discounted or sale price).

discounts are *predictable*. Strategic and forward-looking buyers therefore 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 draws near. Stokey (1979), Conlisk, Gerstner, and Sobel (1984), Sobel (1991), Board (2008), and Garrett (2016) are but a few instances.

For an example of common empirical price patterns, consider Février and Wilner's (2016) analysis of a French music retailer in the early 2000s. They observe that price discounts are typically discrete rather than gradual and that purchases do not decline immediately before sizeable price reductions (see Figure 1). Février and Wilner (2016) interpret the latter observation as indicating that buyers are unable to foresee the timing of discounts. They find that demand at the regular price is nonetheless sensitive to the frequency and size of price reductions, which is taken as evidence consumers are forward-looking. The view of consumers as forward-looking but uncertain about future prices is in common with much of the literature on dynamic demand estimation (see the discussions in Gowrisankaran and Rysman, 2012, who consider camcorders, and in Hendel and Nevo, 2013, who consider soft-drinks).

In this paper, we propose a novel theory of buyers' failure to predict the timing of price reductions based on optimal price discrimination by sellers. We show that setting random discounts is optimal for a seller with commitment power who faces buyers that are forward-looking and *risk averse*, and who arrive to the market over time.[1,2] This contrasts with with the optimality of constant prices in

[1]The assumption of full commitment seems useful for shedding light on pricing patterns adopted by sellers. Our view is in line with Board and Skrzypacz (2016) who suggest that commitment "is reasonable with applications such as retailing, online ads, and concerts in which the seller automates the pricing scheme and uses it repeatedly."

[2]While not all price reductions are difficult to predict in practice (e.g., Black Friday and Christmas specials), many

important benchmarks with *risk-neutral* buyers (see Stokey, 1979, and Conlisk, Gerstner, and Sobel, 1984).

While risk aversion is a commonly studied assumption in other allocation problems such as auctions, its role has largely been overlooked in relation to dynamic pricing. There is, nevertheless, evidence of risk aversion in markets for durable goods. For instance, Chen, Kalra, and Sun (2009) document the widespread sale of warranties for electronic goods at much worse than fair prices. Also, the presence of a "buy-it-now" price in online auctions on platforms such as Yahoo and eBay has often been associated with buyer risk aversion (see Budish and Takeyama, 2001, and Reynolds and Wooders, 2009).[3]

The seller in our model commits to the price path offered to buyers who arrive over time. We show that a virtually optimal price path involves a constant regular price, with short-lived episodes of discounting that are randomly timed, and which buyers in particular find unpredictable.[4] Within each discounting episode, the initial discount is small, and after each further discount there is a positive probability that the price goes back to the regular price. The pricing policy is stationary in that the future process for prices depends only on the current price (and not, for instance, on calendar time).

An important feature of our pricing policy is that it implies virtually all the buyers with a given valuation purchase at the same price. This is possible because buyers have the same beliefs regarding the future evolution of prices. For instance, all highest valuation buyers have the same incentive to accept the constant regular price instead of waiting for lower ones. Similarly, buyers with intermediate values wait for the price to drop enough to purchase. Delaying purchase involves the risk that the discounting episode ends and the price returns to the regular level. Each type of the buyer arriving at a time where the regular price is offered ends up buying at a predictable price but at a random time. The importance of buyers with the same valuation purchasing at the same price is that this is efficient given buyer risk aversion. In essence, buyers are protected from pricing risk associated with their time of arrival to the market, increasing the surplus the seller can extract.

Our analysis of the seller's problem proceeds in two main steps. The first step (in Section 3) involves analyzing a static allocation problem with a single (representative) buyer, with payments made only in case the buyer receives the good. This analysis is closely connected to work on auctions with risk-averse bidders such as Matthews (1983), Maskin and Riley (1984), and Moore (1984),

retailers discount products throughout the year but do not inform customers about the timing in advance. Since timely advance information *could* be made available at little cost, it may be reasonable to infer that its absence is often part of a deliberate policy.

[3]Some caution may be needed in interpreting buyers as risk averse when there are small stakes, given the implications for gambles with larger stakes. However, risk-averse preferences might also be understood as capturing consumer loss aversion; see Rabin (2000) who discusses loss aversion as a leading explanation of "small-scale" risk aversion that has been observed empirically.

[4]The reason for considering price paths that are only "virtually optimal" relates to the impossibility of offering different price discounts "within the same instant of time". We show that this means there are cases where no optimal price path exists, and we look at virtually optimal policies in these cases.

although a key difference is the restriction to "winner pays" which necessitates separate analysis. The unique optimal mechanism involves a type-dependent probability of receiving the good and a non-stochastic payment for allocation.

The second step (in Section 4) is to consider a setting where buyers arrive over time, and where the profits from the static mechanism provide a natural upper bound on the available profits per buyer. We show that this upper bound can be attained — or approximated arbitrarily closely — by a stochastic price path. As we explain below, the random price processes can be understood in terms of a dynamic implementation of the optimal static allocations. The properties of the dynamic format are therefore intimately related to those of the optimal static mechanism. For instance, the above result that all buyers with the same valuation purchase at the same price in the dynamic format follows from the same result for the static mechanism.

A final part of our analysis (in Section 5) relaxes the assumption that buyers observe the prices posted before their arrival. We argue that this permits the theory to accommodate observed discount patterns. Such analysis is relevant to empirical investigations of the topic in the macroeconomics literature on price stickiness (see, e.g., Nakamura and Steinsson, 2008, Eichenbaum, Jaimovich, and Rebelo, 2011, and Kehoe and Midrigan, 2015), and in industrial organization (e.g., Pesendorfer, 2002, Berck, Brown, Perloff, and Villas-Boas, 2008, and Février and Wilner, 2016). Other relevant theories of price discounts are reviewed in Section 6.

2 Set-up

We consider a setting with a single seller and either one or many buyers. Buyers have unit demand and are risk averse. The seller faces no capacity constraints, zero production costs, and is risk neutral.

Our setting is static in Section 3 and dynamic in Section 4. In the dynamic setting, time is continuous and the horizon infinite, with time indexed by $t \in [0, \infty)$. Both the seller and buyers then have a common discount rate $r > 0$. The arrival process for buyers is delayed to Section 4.

Buyers' enjoyment of the good depends on their "types". We label these types $\{\theta_n | n = 1, ..., N\}$, with $\theta_N > ... > \theta_1 > 0$. Values $\beta_n > 0$ denote the probability of each type θ_n, with $\sum_{n=1}^{N} \beta_n = 1$. In the dynamic setting, types will be independent of the arrival date.

Any buyer can transact at most once with the seller; that is, allocation of the good and payment must occur on the same date. Payments are made only if the buyer obtains the good, and the payoff from not receiving the good is set to zero. For each type θ_n and payment $p \in \mathbb{R}_+$, we let $v_n(p) \in \mathbb{R}$ denote the utility of a purchase for this type and payment. Note that, in the dynamic setting of Section 4, our assumptions imply that a type θ_n buyer's intertemporal payoff is $e^{-rt} v_n(p_t)$ if the good is purchased at price p_t on date t, while it is equal to zero in case of never purchasing.

We restrict buyer preferences as follows. For each n, $v_n(\cdot)$ is a strictly decreasing, strictly concave, and twice continuously-differentiable function. We normalize by setting $v_n(\theta_n) = 0$ for each

n, i.e. a buyer's type corresponds to his valuation for the good. We make the following additional assumptions.

Condition A.

A1 *Higher types are "more eager":* For any $n=1,...,N-1$ and $p<\theta_n$, $\frac{-v'_{n+1}(p)}{v_{n+1}(p)} < \frac{-v'_n(p)}{v_n(p)}$.

A2 *Higher types are less risk averse:* For any $n=1,...,N-1$ and $p \in \mathbb{R}_+$, $\frac{v''_{n+1}(p)}{v'_{n+1}(p)} \leq \frac{v''_n(p)}{v'_n(p)}$.

The role of Assumptions A1 and A2 will be explained further below. For now, note that they will have important implications for the form of optimal static mechanism in Section 3, and consequently for optimal stochastic price processes in Section 4. For instance, Assumptions A1 and A2 together will ensure that higher types are more likely to receive the good and pay higher prices in the optimal mechanism. A natural interpretation of higher types (see, for instance, Maskin and Riley, 1984) is that they represent wealthier individuals, since risk aversion is generally believed to be decreasing in wealth.

3 Optimal mechanisms in static environments

This section considers static mechanisms for a single buyer, anticipating the relevance for dynamic pricing problems in the following section. The key reason for the connection will be that allocating the good at price p after delay t generates the same expected payoffs for the players as immediate allocation at the same price with probability e^{-rt}.

By the revelation principle, it will be without loss of generality to consider direct mechanisms which allocate a unit to each type θ_n with probability x_n. In addition, these mechanisms stipulate a potentially random price conditional on assignment, with distribution $H_n : \mathbb{R}_+ \to [0,1]$ for each type θ_n. In case $x_n = 0$, we might as well set the payment deterministically to zero and we do so below. A static mechanism can then be written as $M = (x_n, H_n)_{n=1}^N$.

To define incentive compatibility, note that type θ_n's expected payoff when reporting θ_k is

$$U_{n,k} \equiv x_k \int v_n(p) dH_k(p).$$

An *incentive compatible* direct mechanism is one where, for all n and k, $U_{n,n} \geq U_{n,k}$. Apart from being incentive compatible, the static mechanism should be *individually rational*, which requires $U_{n,n} \geq 0$ for all n. We say that a mechanism has *deterministic payments* if H_n is degenerate for each n. In this case, with an abuse of notation, we may write the mechanism as $M^D = (x_n, p_n)_{n=1}^N$. The following result implies monotonicity of the allocation in mechanisms with deterministic prices.

Lemma 1. *Consider any two types θ_k and θ_l with $k<l$, and consider two allocation probabilities and (sure) prices (x', p') and (x'', p'') with $x' < x''$ and $p'' \leq \theta_k$. If $x''v_k(p'') \geq x'v_k(p')$, then $x''v_l(p'') > x'v_l(p')$.*

Lemma 1 follows from Assumption A1, which provides a sense in which higher types are less price sensitive or "more eager" to purchase at higher prices. The result also assumes deterministic payments. Our next result is that this is the relevant case for optimal mechanisms, where we use now both Assumptions A1 and A2.

Lemma 2. *Any optimal mechanism has deterministic payments.*

The proof of Lemma 2 involves finding, for any mechanism with random payments, a mechanism with deterministic payments and higher profits. This takes place in two steps. First, we consider the mechanism in which each type is charged the certainty equivalent price, i.e. the sure price that gives each type the same expected payoff when receiving the good as under the original mechanism. While this mechanism is more profitable than the original if the buyer reports the truth, truth-telling may not be incentive compatible. The second step then involves suppressing the option to report certain types that are less profitable for the seller, determining a kind of indirect mechanism.[5] Precisely, the new mechanism permits a report equal to the lowest type that obtains the good with positive probability under the original mechanism. Other available reports are then determined in an increasing sequence, allowing a type to be reported if and only if profits are greater than for all lower available reports. We then establish that every type sends a message from the available options which generates higher seller profits than for the truthful message in the original mechanism.

Let us briefly position Lemma 2 in the literature. First note that, in an auctions setting, Maskin and Riley (1984) establish a sufficient condition for the optimality of deterministic mechanisms using an optimal control argument (see their Theorem 9). As Moore (1984) observes, this sufficient condition appears difficult to evaluate and depends on an endogenous variable (see Equation (45) in Maskin and Riley). Moore therefore proposes, in a single-buyer setting with discrete types, sufficient conditions on primitives guaranteeing optimality of deterministic payments. Unlike our model, those of Maskin and Riley and Moore feature a payment also in case the buyer is not awarded the good. Moore's argument (see his Theorem 1) depends on payments by the losing buyer, so it does not apply to our setting.[6]

Lemmas 1 and 2 permit further characterization of the optimal mechanism. We show the following result.

Proposition 1. *The optimal mechanism is unique. It is fully characterized by a weakly increasing sequence* $(x_n^*, p_n^*)_{n=1}^N$ *of allocation probabilities and prices for each type such that* $x_N^* = 1$. *Downward incentive constraints bind: for all* $n = 1, \ldots, N$, $x_n^* v_n(p_n^*) = x_{n-1}^* v_n(p_{n-1}^*)$, *where we put* $x_0^* = p_0^* = 0$.

[5] Some types may send the same truthful message as under the original direct mechanism, but clearly others may not.

[6] Note that Lemma 3 of Matthews (1983) establishes the optimality of deterministic payments for the winning bidder in an auction, but does so only under preferences defined by constant absolute risk aversion.

It is worth highlighting here the proof of uniqueness, which is new to the literature. It supposes that there are two distinct optimal mechanisms, and then constructs a randomization over them where a buyer reporting to the mechanism plays each of the original distinct mechanisms with a fixed probability. The new mechanism is also optimal, and it involves randomized payments, contradicting Lemma 2. Uniqueness is of interest here because it will be important for our discussion of optimal price paths in the dynamic environment of Section 4.

Optimality of random mechanisms. We do not attempt a full characterization of the optimal allocations $(x_n^*)_{n=1}^N$, but note that concavity of the agent's preferences v_n can imply the optimality of random allocations: that is, it may be that $x_n^* \in (0,1)$ for some n. To illustrate, consider first the case where $N=2$ so there is a low type and a high type. By Proposition 1, it is optimal to set the probability of allocation to the high type equal to one. Also, letting x_1 denote the probability of allocation to the low type and p_2 the price charged to the high type, we may assume $v_2(p_2) = x_1 v_2(\theta_1)$ (i.e., indifference of the high type to the low type's option). We then have that θ_1 is the price charged to the low type, and $v_2^{-1}(x_1 v_2(\theta_1))$ the price charged to the high type, where v_2^{-1} is the inverse of v_2. We can therefore write the seller's profits as:

$$\beta_1 x_1 \theta_1 + \beta_2 v_2^{-1}(x_1 v_2(\theta_1)) \tag{1}$$

The optimal mechanism is then determined by maximizing the expression in Equation (1).[7]

Proposition 2. *Suppose $N=2$ and consider the allocation probability to the low type in the optimal mechanism, x_1^*, which is the value maximizing the expression in Equation (1). There is an interval $(\underline{\beta}, \bar{\beta})$, with $0 < \underline{\beta} < \bar{\beta} < 1$, such that x_1^* is in $(0,1)$ if and only if $\beta_2 \in (\underline{\beta}, \bar{\beta})$. If $\beta_2 \leq \underline{\beta}$, then $x_1^* = 1$, and if $\beta_2 \geq \bar{\beta}$, then $x_1^* = 0$.*

Note that it is the concavity of $v_2(\cdot)$, or equivalently the concavity of $v_2^{-1}(\cdot)$, that explains why we find an interior solution for a range of probabilities β_2 of the high type, different to the case where $v_2(\cdot)$ is linear.[8] Intuitively, when the probability of allocation to the low type (i.e., x_1) is low, the price charged to the high type is high, and so the high type is more price sensitive. Therefore, raising x_1 above the lower bound of zero requires reducing the price of the high type relatively little, suggesting the profitability of doing so. Conversely, when x_1 is high, the price charged to the high type is low, and so the high type is less price sensitive. Lowering x_1 below the upper bound of one permits increasing the price to the high type by a relatively large amount, which suggests the profitability of doing so. Indeed, for intermediate values of the probability of the high type (namely

[7] A previous version of the paper established Proposition 2 without Assumptions A1 and A2.

[8] That is, when payoffs are linear in prices, we obtain the usual "no-haggling" result that it is optimal to make a take-it-or-leave it offer to the buyer (see Riley and Zeckhauser, 1983, for this result in a dynamic setting with many buyers).

$\beta_2 \in (\underline{\beta}, \bar{\beta})$), both the above adjustments are profitable, explaining why the optimal choice of x_1 is interior.

When it comes to consider many types, a key question which will be relevant for dynamic pricing in Section 4 is whether we can now have a range of types receiving a random allocation. We aim here to illustrate this possibility in a case inspired by the auctions literature, namely taking the constant absolute risk aversion (CARA) preferences considered by Matthews (1983) and a continuum of types.[9] We introduce preferences in the continuum setting by writing the valuation of type θ purchasing at price p as $v(p;\theta) = 1 - e^{-R(\theta-p)}$, where $R > 0$ is the coefficient of absolute risk aversion. For any discrete set of types, $\{\theta_1, ..., \theta_N\}$, these preferences satisfy all the conditions of the model set-up, including Assumptions A1 and A2.

We suppose types θ are distributed on a bounded interval $[\underline{\theta}, \bar{\theta}]$, with $\underline{\theta} \geq 0$, according to a twice continuously differentiable distribution F. We denote the corresponding density by f. At a posted price p, all types $\theta \geq p$ find it optimal to purchase, so expected profit from a price p is $p(1-F(p))$. Suppose F is such that this profit is uniquely maximized by a price $p^* \in (\underline{\theta}, \bar{\theta})$. Noting that the derivative of profits with respect to price is $-f(p)p+1-F(p)$, a sufficient condition is that $\frac{1-F(p)}{f(p)}$ is decreasing, and $\lim_{p\downarrow\underline{\theta}} \frac{1}{f(p)} > \underline{\theta}$ while $\lim_{p\uparrow\bar{\theta}} \frac{1-F(p)}{f(p)} < \bar{\theta}$.

Proposition 3. *Suppose the buyer has CARA preferences as just introduced, with types distributed according to a cdf F satisfying the assumptions just stated. Then there is $\bar{R} > 0$ such that, for all $R > \bar{R}$, the seller obtains higher profit with a random mechanism than with the optimal deterministic mechanism.*

By Proposition 3, the seller may prefer a random mechanism, with random allocations for a positive measure of types, to the optimal deterministic mechanism, which involves a posted price to obtain a unit for sure. This occurs whenever the buyer is sufficiently risk averse. The random mechanism we consider in the proof is a simple perturbation of the optimal posted price, introducing an additional option for the buyer to obtain the good with probability one half. We show the perturbed mechanism is more profitable whenever R is large enough. The same basic argument can be applied to a fine grid of discrete types, with the comparison again being between the optimal posted price (where profit approaches that in the continuum setting with the fineness of the grid) and a perturbed mechanism. It is easy to see that existence of an optimal mechanism is obtained in the setting with discrete types, so we have that the unique optimal mechanism can involve randomization. The discrete environment is also amenable to numerical analysis, where we can observe the optimality of assigning random allocations to a range of types for particular cases.[10]

[9] The difference in our environment relative to Matthews (1983) is that the buyer in his setting can make a payment when not receiving the good. In the optimal mechanism of his model, indeed, buyers pay when they do not get the good (see his Theorem 1).

[10] We approximate the uniform distribution on $[0, 10]$ with $N = 101$ equi-probable and evenly spaced types. Letting $R = 1$, we find that all types below 2.0 are given the good with probability zero, and all types above 8.5 are given the good with probability one. All other types are given the good with an interior probability.

4 Optimal price mechanisms in dynamic arrivals

This section considers dynamic pricing mechanisms, initially for a fixed arrival date and then for buyers arriving over time. The optimal profits from the static mechanism studied in the previous section, denoted Π^*, will be an upper bound on the profits attainable from each cohort. Our main question is whether and how a dynamic price path can generate profits equal or close to this bound.

The restriction implicit in the consideration of price paths is that one price is offered at any instant. We view the seller as being able to fully commit to the path of prices, including to random price paths, particular instances of which will be described below. Throughout this section, we refer frequently to the values that characterize the optimal static mechanism, namely $(x_n^*, p_n^*)_{n=1}^N$.

Deterministic price paths for fixed arrival date. Consider first a single buyer arriving at a fixed date, say $t = 0$. There is a *deterministic* price path achieving profits Π^*. In this sense, the restriction to a deterministic price path comes at no cost to the seller.

To determine an optimal deterministic price path, suppose the buyer purchases whenever indifferent. We then choose prices so that the buyer makes any purchase at a date in $\mathcal{T} = \{-\log(x_n^*)/r : x_n^* > 0\}$. For each $x_n^* > 0$, we set the price at date $t_n^* = -\log(x_n^*)/r$ to p_n^*. The prices at other dates are immaterial, provided they are sufficiently high; for instance, it is enough to set them to p_N^* (the highest payment in the static mechanism). Prices fall over dates in $\mathcal{T}$. If the buyer has type θ_n, he waits for the price to fall to p_n^* at date t_n^* to purchase (assuming such a date exists). The optimality of purchasing at date t_n^* follows from the same incentive constraints respected by the static mechanism $(x_n^*, p_n^*)_{n=1}^N$. That is, for each n and k with $n \neq k$, $U_{n,n} \geq U_{n,k}$ guarantees type θ_n does not gain by purchasing at date t_k^*. Using Proposition 1, the set of purchase dates $\mathcal{T}$, as well as the prices at these dates, are unique across any optimal deterministic price path.

Suboptimality of deterministic price paths for dynamic arrivals. The above observations hold equally for a single buyer and for a unit mass of buyers arriving at a fixed date. Suppose now that a large population of buyers instead arrives to the market over time. Viewing buyers as infinitesimal, we specify a finite arrival rate $\gamma > 0$. Normalize demand by setting $\int_0^\infty \gamma e^{-rt} dt = 1$ (i.e., $\gamma = r$), so that the seller's total profits correspond to a per-buyer (weighted) average. The value Π^* is then an upper bound on the seller's expected profits.

Consider now whether profits Π^* are attainable with a deterministic price path. This turns out not be possible if the optimal static mechanism involves $x_n^* \in (0, 1)$ for some n.[11] To see this, note that to maximize profit from buyers arriving at date $t = 0$, the price path must induce these buyers to purchase at dates in $\mathcal{T}$ at the prices described above. However, using the aforementioned uniqueness of the optimal static mechanism, profits cannot then be maximized for later cohorts.

The reason profits must be less for later cohorts can be seen by considering buyers who arrive a

[11] When there is no n with $x_n^* \in (0, 1)$, there is an optimal price path which is constant with the price equal to the lowest type of buyer receiving the good in the optimal static mechanism.

short time before the lowest price occurs. For any type arriving at such a time, the buyer optimally purchases when the low price is charged or earlier, rather than with the delays specified above which are necessary to achieve optimal profits.

Convenient dynamic revelation mechanisms. To begin understanding how a stochastic price path can help, we next introduce a class of stochastic dynamic mechanisms where buyers are sometimes asked to report their types, and in which the seller obtains expected profits Π^*. Because this is a combination of static revelation and dynamic formats, we refer to this as the *Hybrid Mechanism*. From hereon, we assume there is some type θ_n such that the allocation probability in the static mechanism is $x_n \in (0,1)$, as otherwise a deterministic and constant price path is optimal (see Footnote 11).

Let $\theta_{\bar{n}}$ then be the highest type for which $x^*_{\bar{n}} < 1$. The mechanisms we examine set a constant "buy-it-now" price that is to be accepted immediately by all types strictly above $\theta_{\bar{n}}$. Buyers who have not purchased are then asked to play static and memoryless revelation mechanisms at random times. That is, at randomly determined times, buyers are asked to send reports of their types and the mechanism determines the allocation probability and price on this basis, but not on the basis of previous reports. Buyers leave the market if they are awarded the good, but otherwise remain and can report to future revelation mechanisms each time they occur.[12]

In selecting the Hybrid Mechanism, we are guided by the optimal static mechanism of the previous section. Let the buy-it-now price equal p^*_N, the price paid for allocation with certainty in the static mechanism. Suppose revelation mechanisms occur at a Poisson rate $\lambda_{\bar{n}} = \frac{rx^*_{\bar{n}}}{1-x^*_{\bar{n}}}$. The revelation mechanism awards the good with certainty to types at least $\theta_{\bar{n}}$ at price $p_{\bar{n}}$, and awards it with probability $\frac{x^*_n}{x^*_{\bar{n}}}\frac{1-x^*_{\bar{n}}}{1-x^*_n}$ to types θ_n below $\theta_{\bar{n}}$ at prices p^*_n. The implication is that a buyer of type $\theta_n \leq \theta_{\bar{n}}$ who continues to report truthfully in the revelation mechanisms receives the good at a Poisson rate $\lambda_n = \frac{rx^*_n}{1-x^*_n}$, and so the expected discounting until purchase is x^*_n.[13] Assuming types greater than $\theta_{\bar{n}}$ purchase immediately at the buy-it-now price, and all other types report truthfully in the revelation mechanisms, the expected payoff of each type is the same as in the optimal static mechanism, and the seller earns expected profits Π^*.

It remains to check that buyers are willing to behave as prescribed. This is straightforward, however, because stationarity of the mechanism implies the optimality of a stationary strategy for each buyer. In a stationary strategy, a buyer of any type θ_n either purchases immediately at price p^*_N, or never takes the buy-it-now price and instead either (a) makes the same report θ_k in every revelation mechanism, or (b) never reports to any revelation mechanism. The payoff from the buy-it-now price is $v_n(p^*_N)$ and that from continuing to report θ_k (given that the revelation mechanism is not currently available) is $x^*_k v_n(p^*_k)$. Therefore, verifying incentive compatibility and individual rationality of the behavior prescribed for the buyer above involves checking the same inequalities as

[12] Since buyers have unit demand, they cannot benefit from participating after receiving the good.

[13] The calculation of the Poisson rate λ_n follows from standard formulae for Poisson thinning.

guaranteed by the satisfaction of the incentive constraints for the static mechanism $(x_n^*, p_n^*)_{n=1}^N$.

(Virtual) optimality of random price paths. Next, note that there is a strategically equivalent implementation of the Hybrid Mechanism which is a format closer in appearance to a price path. At each instant a revelation mechanism would take place, there is instead an episode of "price discounting" that takes place instantaneously. At such a time, the good is initially offered for sale at price $p_{\bar{n}}$. Then, for $n \leq \bar{n}$, if the price p_n has been offered, the price drops to p_{n-1} with probability $\frac{x_{n-1}^*}{x_n^*}\frac{1-x_n^*}{1-x_{n-1}^*}$, where recall that we put $x_0^* \equiv 0$.[14] With complementary probability, the episode of price discounting ends and the good remains available only at the buy-it-now price p_N^* until the next discounting episode arrives at Poisson rate $\lambda_{\bar{n}}$. Buyers with types above $\theta_{\bar{n}}$ purchase immediately, while buyers with types $\theta_n \leq \theta_{\bar{n}}$ either wait until the price hits p_n^*, or never purchase if $x_n^* = 0$.

Recall that a price path requires one price to be offered at each instant. The dynamic pricing format just described generally fails to be a price path, because different discounted prices may be offered on the same date. It turns out then that a price path can nevertheless approximate the suggested mechanism arbitrarily closely, so we obtain the following result.

Proposition 4. *Suppose buyers arrive over time as specified above. For any $\varepsilon > 0$, there is a (possibly random) price path such that the seller's expected profits are at least $\Pi^* - \varepsilon$.*

Our approach to show the approximation of profits Π^* is, in essence, to spread episodes of sequential discounts over intervals of short duration. The proof of Proposition 4 constructs a Markov price process where the price equals p_N^* in the "regular" state. At arrival rate $\lambda_{\bar{n}} = \frac{r x_{\bar{n}}^*}{1-x_{\bar{n}}^*}$ the state transitions to the first-discount state, where the price is $p_{\bar{n}}^*$. Shortly after, either the state transitions back to the regular state or transitions to the second-discount state. The transition probability is such that, in the regular state, the second-discount state is first reached at a Poisson rate $\lambda_{\bar{n}'} = \frac{r x_{\bar{n}'}^*}{1-x_{\bar{n}'}^*}$, where $x_{\bar{n}'}^*$ is the largest value satisfying $x_{\bar{n}'}^* < x_{\bar{n}}^*$. The price at the second-discount state is $p_{\bar{n}'}^*$. Importantly, the second-discount state is only reached after the first-discount state. The process features as many discount states as different values of $x_n^* \in (0, 1)$, each lower discount state reached with some probability shortly after the previous discount state. Figure 2 depicts a realization of the price path for $N = 3$ when $0 < x_1^* < x_2^* < x_3^* = 1$.

By restricting attention to stationary strategies (as for the Hybrid Mechanism above) we show that a θ_n-buyer arriving at the regular state buys at the same price p_n^* and with the same expected discounting x_n^* as in the dynamic revelation mechanisms described above. Buyers resolve a trade-off between purchasing and waiting for further discounts. For example, in Figure 2, type θ_2 accepts the discounted price p_2^* because, even though he knows that the price may be further discounted soon, it may also go back to the regular price, and hence he will have to wait a while for another discount.

When the duration of discount states is short, most buyers arrive when the state is regular. Hence, the expected profits from cohorts arriving at these times equal Π^*. The only loss in profits

[14] If $x_{n-1}^* = x_n^*$, the same price p_n^* is offered again with certainty. This "re-offering" is clearly immaterial and could equivalently be dropped from the description of the mechanism.

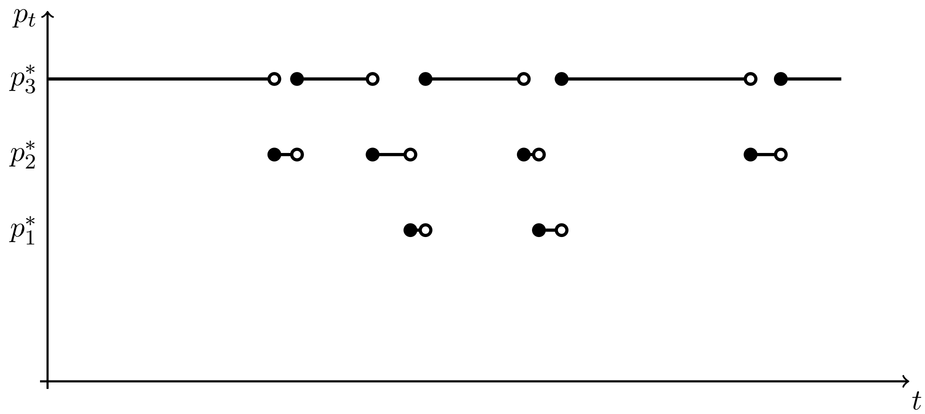

Figure 2: Example of a realization of a price path for $N = 3$ when $0 < x_1^* < x_2^* < x_3^* = 1$. In the example, there are four discount episodes from the regular price p_3^*. In the first and the last, only p_2^* is offered before p_3^* is offered again. In the other two discount episodes a second discount at price p_1^* occurs after the first discount.

arises from buyers arriving *during* the episodes of price reductions. For instance, a buyer of type θ_N arriving during these episodes purchases at a price below the price p_N^* prescribed by the optimal static mechanism. By the uniqueness of the optimal static mechanism, such a purchase is inconsistent with profit maximization. Still, the episodes of price reductions can be chosen as short as desired. Hence, there are price processes with profits arbitrarily close to Π^*.

Proposition 4 leaves open the question of whether profits Π^* can be exactly attained. In Appendix II, we prove that there exists a price process attaining Π^* if and only if there are no two values $x_{n'}^*, x_{n''}^* \in (0,1)$ with $x_{n'}^* \neq x_{n''}^*$. An example is where $N = 2$. If β_2 belongs to the interval $(\underline{\beta}, \bar{\beta})$ identified in Proposition 2 so that $x_1^* \in (0,1)$, then x_1^* is the only allocation probability in $(0,1)$, since recall $x_2^* = 1$. Then the seller achieves the optimal profits Π^* through a constant regular price of p_2^* that is occasionally discounted to price $p_1^* = \theta_1$ at Poisson rate $\frac{rx_1^*}{1-x_1^*}$.

5 Timing of discounts

As mentioned in the Introduction, work in macroeconomics and industrial organization has been interested to understand patterns of price discounts. For instance, Pesendorfer (2002), Février and Wilner (2016) and Lan, Lloyd, and Morgan (2016) find an increasing hazard rate for discounts: a long spell without a price discount predicts that one will occur relatively soon. In contrast, our findings above predict Poisson discounts, meaning a constant hazard rate. In this section, we limit the price information available to consumers by supposing that they observe prices *only from their arrival to the market* and argue that this permits our theory to better accommodate the empirical patterns.

We specialize for convenience to the case with two types, assuming that β_2 belongs to the interval $(\underline{\beta}, \bar{\beta})$ identified in Proposition 2. An initial observation is that, when the purchase decisions of buyers are measurable only with respect to the information generated by prices since arrival, Poisson discounting remains optimal. As before, the price is set at p_2^* (identified in the optimal static mechanism) except at discount dates when it drops to $p_1^* = \theta_1$. Also, the rate of price discounting is $\frac{rx_1^*}{1-x_1^*}$, as identified in the previous section (recalling that x_1^* is the probability of allocation to low types in the optimal static mechanism). Since profits from each cohort are already maximized by this policy, there is no scope for the seller to profit from buyers' ignorance of past prices.

There are now also other optimal price processes, with the regular price set at p_2^* and discounted to price θ_1. Suppose these discount dates are determined by a simple point process. Low types can be assumed to purchase at discount dates, while optimality requires that all but possibly measure zero of high types are induced to purchase immediately at price p_2^*. For any arrival time t, let $\tilde{\tau}_1^t$ be the random purchase time of a low type arriving to the market at date t. As observed in the previous section, optimality calls for the expected discounting until time of purchase for the low type to equal the probability of allocation in the static mechanism. In particular, for almost all t, we must have

$$\mathbb{E}\left[e^{-r(\tilde{\tau}_1^t - t)}\right] = \mathbb{E}\left[e^{-r(\tilde{\tau}_1^t - t)} \middle| \tilde{\tau}_1^t > t\right] = x_1^*. \tag{2}$$

When (2) holds, a high type who arrives to the market at date t and has no information about past prices expects the same payoff purchasing immediately at arrival (paying price p_2^* almost surely) or, alternatively, waiting and purchasing at the next price discount. A buyer who delays purchase, however, *is not restricted to purchasing at a discounted price*, so the condition (2) is not sufficient to guarantee immediate purchase. The condition for immediate purchase can be expressed as one on the point process for price discounting, which we state next.

Proposition 5. *Suppose that buyers arrive over time, but observe prices only since arrival to the market. Suppose a "regular price" p_2^* is posted except at moments when the price is discounted to θ_1 with the first discount date after t given by $\tilde{\tau}_1^t$. For any t such that* (2) *holds, high types arriving at t are willing to purchase immediately if and only if, for all $s > t$ such that $\tilde{\tau}_1^t > s$ with positive probability,*

$$\mathbb{E}\left[e^{-r(\tilde{\tau}_1^t - s)} \middle| \tilde{\tau}_1^t > s\right] \geq x_1^*. \tag{3}$$

To interpret Condition (3), suppose it is satisfied for a buyer arriving at a given date t. If he delays purchase until $s > t$ and observes no price discount (i.e., $\tilde{\tau}_1^t > s$), then this absence of a price discount is effectively "good news" in that he expects the next discount relatively sooner at date s than at date t (as measured in terms of expected discounting). This means that if it is optimal for a high type buyer to delay purchase, it is optimal to keep on delaying. Condition (2) then ensures that the high type is not willing to delay purchase. Note that the condition accommodates the empirical patterns discussed above, where the hazard rates for price discounts are increasing.

The multiplicity of optimal price processes when buyer information is restricted suggests sellers might pick among them according to different criteria. One consideration may be limiting the maximum inventory size to avoid stocking out. More evenly spaced discounts lead to less accumulation of low types, which limits demand peaks. An extreme case is where price discounts occur a fixed time $\Delta > 0$ apart. A price process satisfying our conditions is obtained by supposing the initial discount is uniformly distributed on $[0, \Delta]$. Buyers arriving at date t with no information on past prices then believe the next discount is uniformly distributed on $[t, t + \Delta]$. The appropriate choice of Δ satisfies

$$\int_0^{\Delta} \frac{e^{-rs}}{\Delta} ds = x_1^*.$$

Recall our demand normalization that buyers arrive at rate r. We then observe that no more than measure $\beta_1 \Delta r$ buyers purchase at any date, so the capacity required to store the inventory is limited.

6 Alternative theories for discounting and further discussion

We conclude by comparing our theory of price discounting with several others in the literature, highlighting possible advantages of our theory. An early explanation for price discounting is mixed-strategy pricing by competing firms, as for instance in Shilony (1977) and Varian (1980). Such theories assume that not all consumers in the market can access the same price offers. Hence, any offered price may be the only one a given customer receives, or it may be one of many among which the customer can choose. When taken to a dynamic context (as, for instance, in Fershtman and Fishman, 1992), these theories typically generate too much price variation relative to the considerable price stability that has been observed empirically.[15] The reason is that the more frequently firms are permitted to choose prices, the more often they randomize. Our theory generates in contrast a regular price that is only sporadically discounted, even though the continuous-time environment permits price changes at any time. While Heidhues and Kőszegi (2014) also account for the prevalence of regular prices in a theory of price dispersion based on consumer loss aversion, their theory is not explicitly dynamic which forestalls an investigation of the timing of discounts such as the one in the previous section.

A leading set of competing theories of price discounting (to which we contribute) are those based on intertemporal price discrimination. For instance, it could be that buyer values change deterministically over time as in Stokey (1979), or they could change randomly over time as in Garrett (2016). Another idea is that different cohorts of buyers have different demand elasticities as in Board (2008). Or, it could be that buyers are more impatient than the seller as in Landsberger

[15]This point is emphasized in a recent paper by Ronayne and Myatt (2019) who suggest a dynamic variant with a first-stage choice of a “list price” and a subsequent discounting decision. While this theory accounts for stickiness of the list price, discounting of the list price does not occur in equilibrium.

and Meilijson (1985).[16] None of these papers, however, predict that the seller can profit from random price discounting. In the work of Garrett and Board, respectively, the optimality of deterministic pricing can be seen from the maximization of virtual surpluses which are linear in the probability of allocation, so that optimal allocations are "bang bang" (i.e., allocation probabilities are corner solutions and thus either zero or one). Related to Landsberger and Meilijson's setting, Correa, Escobar, and Perlroth (2019) consider a model with multiple buyers and different discount rates and argue that deterministic dynamic mechanisms are optimal (they do not consider dynamic arrivals, however).[17]

A possible message regarding work on intertemporal price discrimination, then, is that it is not enough to suggest a rationale for price discounting. There should also be a theory to account for randomization. In our theory based on buyer risk aversion and dynamic arrivals, random price discounts are needed, paradoxically, to shield buyers from pricing risk. Randomization is needed to ensure that (virtually) all consumers with the same valuation purchase at the same price irrespective of when they arrive. Some other papers on risk-averse buyers such as Liu and van Ryzin (2008) and Bansal and Maglaras (2009) have noted that risk aversion could give rise to theories of price discrimination. But these papers do not consider dynamic arrivals and so did not uncover the important role of random price discounts.

A still further class of theories worth mentioning is where random prices arise due to exogenous uncertainty that is realized over time. Examples include Hörner and Samuelson (2011), Board and Skrzypacz (2016), Gershkov, Moldovanu, and Strack (2017) and Dilmé and Li (2019) (which feature demand uncertainty) and Ortner (2017) (which features cost uncertainty). The patterns of price fluctuations are varied, and mainly quite different from the patterns uncovered in this paper. A further crucial difference is that the environment of the present paper is deterministic: in the dynamic setting on which we focus, there are infinetismal buyers and so no demand uncertainty. Corroborating our theory therefore suggests looking for evidence of deliberate randomization by sellers that is not simply a response to shocks in the environment.

[16] Related, it could also be that buyers with high valuations are myopic as in Pesendorfer (2002). See Chevalier and Kashyap (2017), for an application of Pesendorfer (2002)'s model. The myopia assumption, however, seems too strong in many markets; see the evidence that buyers are forward looking as suggested by Chevalier and Goolsbee (2009) and Février and Wilner (2016).

[17] Öry (2017) shows that, when buyers cannot observe their arrival time and can be contacted through some costly (email/text) notification, a seller without commitment sets a constant regular price with evenly spaced discrete discounts. Our work emphasizes that holding sales at random times (at least from the perspective of buyers) can be a fully-optimizing choice for sellers with commitment, even when consumers can fully monitor prices.

A Proofs

Proof of Lemma 1

Proof. Suppose that $x''v_k(p'') \geq x'v_k(p')$. Note that $x''v_l(p'') > x'v_l(p')$ is immediate if $p'' \leq p'$ or if $x' = 0$. Hence, we may assume $x' > 0$ and $p' < p'' < \theta_k$. Then,

$$\begin{aligned} x''v_l(p'') &\geq x'v_l(p')\frac{v_k(p')}{v_k(p'')}\frac{v_l(p'')}{v_l(p')} \\ &= x'v_l(p')e^{\int_{p'}^{p''}\left(\left(-\frac{v_k'(p)}{v_k(p)}\right)-\left(-\frac{v_l'(p)}{v_l(p)}\right)\right)dp} \\ &> x'v_l(p'), \end{aligned} \tag{4}$$

where the strict inequality follows from Assumption A1. □

Proof of Lemma 2.

Proof. Consider an arbitrary (incentive compatible and individually rational) mechanism $M = (x_n, H_n)_{n=1}^N$ that assigns a non-degenerate price distribution H_n to at least one type θ_n. Then, for each n with $x_n > 0$, let p_n be the unique (certainty equivalent) price satisfying

$$v_n(p_n) = \int v_n(p)dH_n(p).$$

This determines a mechanism with deterministic prices $M^D = (x_n, p_n)_{n=1}^N$. Note that (by Jensen's inequality and strict concavity of each v_n) this is strictly more profitable than the original mechanism M if the buyer reports the truth. As mentioned in the main text, however, truth-telling may not be incentive compatible. The remainder of the proof then involves constructing an indirect mechanism which, when the buyer follows an optimal strategy, generates profits at least as high as if the buyer were truthful in M^D.

For each n, let $\pi_n = p_n x_n$ be the seller's expected profit if type θ_n reports truthfully in M^D. We can construct a set of types $\mathcal{J}$ of cardinality $J \equiv |\mathcal{J}|$ along which expected profit strictly increases. We begin by letting θ_{n_1} be the lowest type assigned the good with strictly positive probability in M^D. Then, having determined θ_{n_j}, we let $\theta_{n_{j+1}}$ be the next smallest type such that profits exceed those for θ_{n_j}. That is, for each $j \geq 1$, we let $n_{j+1} = \min\{n : n > n_j,\ \pi_n > \pi_{n_j}\}$ if the set is non-empty, and stop otherwise so that $J = j$. This determines $\mathcal{J} = \{\theta_{n_j} : j = 1, ..., J\}$. We then denote M^R the "restricted" indirect mechanism which is the same as M^D except that the buyer is permitted to choose only among messages in $\mathcal{J}$.

Consider now the reporting decision of any type θ_{n_j} in M^R, with $\theta_{n_j} \in \mathcal{J}$. Because the original mechanism M was individually rational, $p_{n_j} \leq \theta_{n_j}$. Because higher types are less risk averse in the sense of Assumption A2, type θ_{n_j} prefers message θ_{n_j} to $\theta_{n_{j'}}$ with $j' < j$. This implies two important observations. First, by asking type θ_{n_j} to send a message at least his true type, such a type generates expected profit at least π_{n_j} in M^R. Second, because $\pi_{n_j} > \pi_{n_{j'}}$ for any $j' < j$, we must have $x_{n_j} > x_{n_{j'}}$ for all such j'.[18]

Finally, consider a type θ_n with $n \neq n_j$ for any j. If $n < n_1$ then, whether θ_n participates in M^R or not, profits are higher for this type than in the original mechanism M. Suppose instead $n_j < n < n_{j+1}$ for some j, or that $n > n_j$ for $j = J$. Then we recall that for any $j' < j$, we have $x_{n_j} > x_{n_{j'}}$, and also θ_{n_j} prefers message θ_{n_j} to $\theta_{n_{j'}}$. Therefore, by Lemma 1, θ_n strictly prefers message θ_{n_j} to $\theta_{n_{j'}}$. Hence, θ_n can be asked to report a message at least θ_{n_j}, generating profit at least π_{n_j}, which in turn is at least π_n by construction of $\mathcal{J}$. □

Proof of Proposition 1

Proof. **Initial observations.** Note that, following a "taxation principle" (see for instance Rochet, 1985), any mechanism with deterministic payments can be viewed as presenting a choice to the buyer among pairs of strictly positive allocation probabilities and payments.[19] There is no loss in supposing that all combinations are chosen by *some* type, so there is one price for each allocation probability. By Lemma 1, higher types choose weakly higher allocation probabilities. Also, because all combinations are chosen by some type, the prices associated with allocation must be weakly increasing in the allocation probability. Considering momentarily the corresponding direct mechanism, this shows that it can be represented as a weakly increasing sequence $(x_n^*, p_n^*)_{n=1}^N$.

Downward incentive constraints bind. Now consider why downward incentive constraints bind, and continue to view the mechanism as a set of options of (strictly positive) allocation probabilities and accompanying payments. We can first use our initial observations to show that the seller's profits are strictly increasing with the allocation probability for any optimal mechanism. Because prices are weakly increasing, it is enough to observe that, in an optimal mechanism, every purchase is at a strictly positive price. In fact, we can show that no type pays a price less than θ_1. A mechanism that does charge a price less than θ_1 to some types can be adjusted by revising upwards the price of every lower-priced option to θ_1. Every type that chooses an option with a price higher than θ_1 in the original mechanism remains willing to choose the same option, while the other types can be taken to choose an allocation probability that is at least the highest one associated with price θ_1. The seller then makes strictly higher profits for every type that obtained a price below θ_1 in the original mechanism. That the adjusted mechanism is strictly more profitable contradicts the optimality of the original.

We now claim that, if θ_k is the lowest type making some choice (x, p) in an optimal mechanism, then this type must be indifferent to the alternative (x', p') which has the next highest allocation

[18] Note that $\pi_{n_j} > \pi_{n_{j'}}$ requires that $x_{n_j} > x_{n_{j'}}$ or $p_{n_j} > p_{n_{j'}}$ (or both). If $p_{n_j} > p_{n_{j'}}$ then, given that type θ_{n_j} prefers message θ_{n_j} to $\theta_{n_{j'}}$, it must be that $x_{n_j} > x_{n_{j'}}$. Hence, in either case, $x_{n_j} > x_{n_{j'}}$.

[19] The buyer can also choose not to participate, which means zero allocation probability and zero payment.

probability, or to not participating if there is no such alternative. Suppose for a contradiction this is not true for some choice (x, p) and lowest type choosing this option, θ_k. The first case is where there is an alternative (x', p') with the next highest allocation probability relative to (x, p). Then (x', p') is chosen by type θ_{k-1}, and θ_k strictly prefers (x, p) to (x', p'). Since θ_{k-1} prefers (x', p') to any option with a lower allocation probability, θ_k strictly prefers (x', p') to any such option by Lemma 1. Therefore, if we raise the price of the option (x, p) to some $\widetilde{p}$ where type θ_k is indifferent between $(x, \widetilde{p})$ and (x', p'), type θ_k then prefers $(x, \widetilde{p})$ to any smaller allocation probability. Again by Lemma 1, any type higher than θ_k then strictly prefers $(x, \widetilde{p})$ to any option with a smaller allocation probability. It follows that it is incentive compatible for any type choosing the original option (x, p) to choose at least the probability of allocation x when the price p is changed to $\widetilde{p}$. Because profits are strictly increasing with the allocation probability, these types now generate strictly higher profits than before. Also, types not choosing the original option (x, p) continue to make the same choice as before. Thus we arrive at a new mechanism that is strictly more profitable than the original, contradicting the optimality of the original.

The second and remaining case is where there is no allocation probability lower than (x, p). By assumption, then, $p < \theta_k$ where θ_k is the lowest type receiving the good with positive probability. Analogous to the previous case, we consider raising this price to $\widetilde{p} = \theta_k$. Any type willing to participate in the original mechanism remains willing to participate. Because (x, p) represents the least profitable option for the seller in the original mechanism, it follows that profits strictly increase in the adjusted mechanism. This again contradicts the optimality of the original mechanism.

Finally, note that we have shown each type is indifferent to mimicking the choice of the downward adjacent type, or to not participating in the case of the lowest type, θ_1. This is either because the lower type makes the same choice, or because of the indifference to the next highest allocation probability, or to non-participation, as shown above. Therefore, considering the direct mechanism, we have for all $n = 1, \ldots, N$, $x_n^* v_n(p_n^*) = x_{n-1}^* v_n(p_{n-1}^*)$, where we put $x_0^* = p_0^* = 0$.

Highest type receives allocation probability one. Now consider the highest allocation probability. If this is less than one, the probability can be increased to one and the payment adjusted (weakly) upwards so that the lowest type that chooses this option in the original mechanism remains indifferent to the next highest allocation probability. Since this type prefers the highest allocation probability to all other options, all higher types also prefer the highest probability by Lemma 1, and profits in the mechanism strictly increase. Considering direct mechanisms, this shows that optimality requires $x_N^* = 1$.

Existence and uniqueness of the optimal mechanism. Existence of an optimal mechanism can be seen from the following observations. Given that downward incentive constraints bind, profits can be determined simply from the choice of allocations $(x_n)_{n=1}^N$, and are continuous in these allocations. Also, the allocations themselves are from the compact set $\left\{(x_1, ..., x_N) \in [0,1]^N : x_1 \leq ... \leq x_N\right\}$.

Let us therefore now show that the optimal mechanism is unique. Suppose for a contradiction that there are distinct mechanisms $(x_n^A, p_n^A)_{n=1}^N$ and $(x_n^B, p_n^B)_{n=1}^N$, both of which are optimal.

We show first that there is a type θ_n such that $x_n^A, x_n^B > 0$ and $p_n^A \neq p_n^B$. Suppose for a contradiction this is not true; that is, assume that for all types θ_n with $x_n^A, x_n^B > 0$ we have $p_n^A = p_n^B$. Consider the smallest value $\underline{n}^A$ such that $x_{\underline{n}^A}^A > 0$ and the smallest value $\underline{n}^B$ such that $x_{\underline{n}^B}^B > 0$. If $\underline{n}^A > \underline{n}^B$ then, from the above characterization of an optimal mechanism, we have $p_{\underline{n}^A}^A = \theta_{\underline{n}^A} > p_{\underline{n}^A}^B$, contradicting our previous assumption which implies $p_{\underline{n}^A}^A = p_{\underline{n}^A}^B$. Given that a contradiction can also be reached for the case $\underline{n}^A < \underline{n}^B$, it must be that $\underline{n}^A = \underline{n}^B = \underline{n}$. Since downward constraints bind in both mechanisms, for all $n > \underline{n}$, we have

$$\frac{x_n^A}{x_{n-1}^A} = \frac{v_n(p_{n-1}^A)}{v_n(p_n^A)} = \frac{v_n(p_{n-1}^B)}{v_n(p_n^B)} = \frac{x_n^B}{x_{n-1}^B}.$$

Therefore, for all $n > \underline{n}$,

$$\frac{x_n^A}{x_{\underline{n}}^A} = \frac{x_n^B}{x_{\underline{n}}^B}.$$

Since $x_N^A = x_N^B = 1$, we have $x_n^A = x_n^B$ for all n, but then the mechanisms $(x_n^A, p_n^A)_{k=1}^N$ and $(x_n^B, p_n^B)_{k=1}^N$ are not distinct.

Now, consider the mechanism determined as follows. The buyer reports his type θ_n, then the allocation probability and payment is determined by one of the two distinct mechanisms according to a 50/50 randomization. This can be described by the "reduced" mechanism that has allocation probability $x_n^C = \frac{1}{2}x_n^A + \frac{1}{2}x_n^B$ for report θ_n. For θ_n such that $x_n^C > 0$, it specifies H_n^C to put mass $\frac{x_n^A}{x_n^A + x_n^B}$ on p_n^A and the remaining mass on p_n^B. Incentive compatibility of the new mechanism is equivalent to the requirement that, for all n, k,

$$\begin{aligned}&\left(\tfrac{1}{2}x_n^A + \tfrac{1}{2}x_n^B\right)\left(\tfrac{x_n^A}{x_n^A+x_n^B}v_n(p_n^A) + \tfrac{x_n^B}{x_n^A+x_n^B}v_n(p_n^B)\right)\\ &\quad \geq \left(\tfrac{1}{2}x_k^A + \tfrac{1}{2}x_k^B\right)\left(\tfrac{x_k^A}{x_k^A+x_k^B}v_n(p_k^A) + \tfrac{x_k^B}{x_k^A+x_k^B}v_n(p_k^B)\right)\end{aligned}$$

or

$$x_n^A v_n(p_n^A) + x_n^B v_n(p_n^B) \geq x_k^A v_n(p_k^A) + x_k^B v_n(p_k^B).$$

This inequality holds by incentive compatibility of $(x_n^A, p_n^A)_{n=1}^N$ and $(x_n^B, p_n^B)_{n=1}^N$, so $(x_n^C, H_n^C)_{n=1}^N$ is incentive compatible. Individual rationality similarly is inherited from $(x_n^A, p_n^A)_{n=1}^N$ and $(x_n^B, p_n^B)_{n=1}^N$. Moreover, it is readily checked that the new mechanism $(x_n^C, H_n^C)_{n=1}^N$ attains the same optimal profit as $(x_n^A, p_n^A)_{n=1}^N$ and $(x_n^B, p_n^B)_{n=1}^N$. However, it does not have deterministic payments, which contradicts Lemma 2. $\square$

Proof of Proposition 2

Proof. Because the expression in Equation (1) is continuous and strictly concave in x_1, it has a unique maximizer $x_1^* \in [0,1]$. Using that $\beta_1 = 1 - \beta_2$, the derivative of this expression with respect to x_1 at $x_1 = 0$ is

$$(1-\beta_2)\theta_1 + \beta_2 \frac{v_2(\theta_1)}{v_2'(\theta_2)} \ .$$

Hence, strict concavity of the expression implies that $x_1^* = 0$ if and only if $\beta_2 \geq \overline{\beta} \equiv \frac{\theta_1}{\theta_1 - \frac{v_2(\theta_1)}{v_2'(\theta_2)}} > \frac{\theta_1}{\theta_2}$.[20] Similarly, the derivative of the expression in Equation (1) with respect to x_1 at $x_1 = 1$ is

$$(1-\beta_2)\theta_1 + \beta_2 \frac{v_2(\theta_1)}{v_2'(\theta_1)}.$$

Strict concavity implies $x_1^* = 1$ if and only if $\beta_2 \leq \underline{\beta} \equiv \frac{\theta_1}{\theta_1 - \frac{v_2(\theta_1)}{v_2'(\theta_1)}} < \frac{\theta_1}{\theta_2}$.[21] It is then necessarily the case that $x_1^* \in (0,1)$ if and only if $\beta_2 \in (\underline{\beta}, \bar{\beta})$. □

Proof of Proposition 3

Proof. Using that higher types receive a higher allocation (by Lemma 1), we may assume that, in a deterministic mechanism, all types above a threshold obtain a unit with probability one, while all types below the threshold obtain it with probability zero. This outcome is induced by a posted price equal to the threshold type. The optimal posted price is $p^* \in (\underline{\theta}, \bar{\theta})$, the maximizer of $p(1-F(p))$ mentioned in the text. Letting θ^* be the threshold type, the first-order condition for the optimal threshold is

$$1 - F(\theta^*) - \theta^* f(\theta^*) = 0. \tag{5}$$

Now, we introduce mechanisms in which some types also purchase with probability 1/2. By Lemma 1, the buyer will follow a threshold rule. Letting $\varepsilon \in (0, \theta^* - \underline{\theta})$, we consider mechanisms in which types $\theta \geq \theta^*$ obtain the good with certainty, types $\theta \in [\theta^* - \varepsilon, \theta^*)$ obtain it with probability 1/2, and types $\theta < \theta^* - \varepsilon$ do not obtain the good at all. Incentive compatibility implies there is one price associated with each probability of obtaining the good.

Incentive compatibility requires the threshold type $\theta^* - \varepsilon$ to be indifferent between not receiving the good and receiving it with probability 1/2. Therefore the payment for receiving the good with probability 1/2, made conditional on receiving it, must be $\theta^* - \varepsilon$. Incentive compatibility also requires indifference of type θ^* between receiving the good with certainty and with probability 1/2, so the

[20] The last inequality follows from the concavity of v_2, which implies the inequality $v_2'(\theta_2)(\theta_2 - \theta_1) < v_2(\theta_2) - v_2(\theta_1)$, together with our normalization $v_2(\theta_2) = 0$.

[21] The last inequality follows from the concavity of v_2, which implies the inequality $v_2'(\theta_1)(\theta_2 - \theta_1) > v_2(\theta_2) - v_2(\theta_1)$, together with our normalization $v_2(\theta_2) = 0$.

payment to receive it with certainty must be $p(\varepsilon)$ satisfying

$$1 - e^{-R(\theta^* - p(\varepsilon))} = \tfrac{1}{2}\big(1 - e^{-R\varepsilon}\big).$$

We can solve for $p(\varepsilon)$ to obtain

$$p(\varepsilon) = \theta^* + \tfrac{1}{R}\log\big(\tfrac{1}{2} + \tfrac{1}{2}e^{-R\varepsilon}\big).$$

We refer to the mechanism which offers the options of purchase at probability $1/2$ at price $\theta^* - \varepsilon$ and purchase with probability 1 at price $p(\varepsilon)$ as the "new" mechanism. This is a perturbation of the optimal deterministic mechanism.

We can compute profits in the new mechanism as

$$\tfrac{1}{2}(F(\theta^*) - F(\theta^* - \varepsilon))(\theta^* - \varepsilon) + (1 - F(\theta^*))\big(\theta^* + \tfrac{1}{R}\log\big(\tfrac{1}{2} + \tfrac{1}{2}e^{-R\varepsilon}\big)\big).$$

The first derivative of the profits of the new mechanism with respect to ε is

$$\tfrac{1}{2}f(\theta^* - \varepsilon)(\theta^* - \varepsilon) - \tfrac{1}{2}(F(\theta^*) - F(\theta^* - \varepsilon)) - (1 - F(\theta^*))\frac{e^{-R\varepsilon}}{1 + e^{-R\varepsilon}}.$$

Using the first-order condition in Equation (5), it is easily seen that the first derivative evaluated at $\varepsilon = 0$ is 0. Similarly, the second derivative evaluated at $\varepsilon = 0$ can be written as

$$-\tfrac{1}{2}f'(\theta^*)\theta^* - f(\theta^*) + (1 - F(\theta^*))\tfrac{R}{4}.$$

Then notice that this is positive whenever R is taken large enough, noting θ^* does not depend on R. In this case, we conclude that $\varepsilon = 0$ is a local minimum and so profits can be improved strictly by taking $\varepsilon > 0$. That is, a random mechanism is strictly more profitable. □

Proof of Proposition 4

Proof. **Attaining or approaching expected profits Π^*.** Recall that $(x_n^*, p_n^*)_{n=1}^N$ denotes the optimal static mechanism. Let $x_0^* \equiv 0$, and let $(n_j)_{j=1}^J$ be the (unique) increasing sequence containing *all* indices satisfying $x_{n_j-1}^* < x_{n_j}^*$ (thus θ_{n_1} is the smallest type obtaining the good with positive probability).

When $J = 1$ the seller can attain expected profits Π^* through a constant price path with price equal to $p_{n_1}^*$. All buyers with type weakly above θ_{n_1} buy upon arrival, and all buyers with type strictly below θ_{n_1} never buy. If $J = 2$ the seller can attain expected profits Π^* with a price process consisting of a regular price $p_{n_2}^*$, with random discounts with price $p_{n_1}^*$ arriving at a constant Poisson rate $\lambda_{n_1} = rx_{n_1}^*/(1 - x_{n_1}^*)$. In this case, all buyers with type weakly above θ_{n_2} buy upon arrival at the regular price, all buyers with type strictly below θ_{n_2} and weakly above θ_{n_1} wait and buy at the discounted price, and all buyers with type strictly below θ_{n_1} never buy.

Assume, for the rest of the proof, that $J > 2$. Consider the following price process, characterized as a process with J states $\{\sigma_j\}_{j=1}^J$ and by some value $\Lambda > 0$. Initializing the state at $t = 0$ to σ_J, the price process is described as follows:

1. In state σ_J the price is $p^*_{n_J}$, and the state changes to state σ_{J-1} at rate $\lambda_{n_{J-1}} = \frac{r x^*_{n_{J-1}}}{1 - x^*_{n_{J-1}}}$.
2. In state σ_j, for $j = 1, ..., J-1$, the price is $p^*_{n_j}$. At rate Λ the state changes to state σ_J and, if $j > 1$, at rate $m^\Lambda_j \Lambda$ the state changes to state σ_{j-1}, where m^Λ_j is obtained below.

For each $j = 2, ..., J-1$, we choose m^Λ_j so that the expected discounting for the first time the state becomes σ_j, starting in state σ_J, is $x^*_{n_j}$. This is achieved if the following equation is satisfied:

$$x^*_{n_{j-1}} = x^*_{n_j} \left(\frac{\Lambda}{\Lambda + m^\Lambda_j \Lambda + r} x^*_{n_{j-1}} + \frac{m^\Lambda_j \Lambda}{\Lambda + m^\Lambda_j \Lambda + r} \right) \quad \Rightarrow \quad m^\Lambda_j = \frac{((1 - x^*_{n_j})\Lambda + r) x^*_{n_{j-1}}}{\Lambda (x^*_{n_j} - x^*_{n_{j-1}})} > 0 \ .$$

We will show that the profits generated by this stochastic price process approach Π^* as we take $\Lambda \to \infty$. To do so, it will be enough to show that each type θ_n that arrives in state σ_J purchases as soon as the price falls to p^*_n, or never purchases if $x^*_n = 0$. By construction, the expected discounting until this time is x^*_n. We do not need to analyze the behavior of cohorts arriving in states other than σ_J since their contribution to expected profits vanishes as $\Lambda \to \infty$.

Consider then why the specified strategy for each type θ_n of waiting to purchase at p^*_n (assuming $x^*_n > 0$) is incentive compatible. As observed for the Hybrid Mechanism, stationarity of the price path (with the future evolution summarized at any point by the state σ_j) implies the optimality of a stationary strategy. Also, for any such strategy there is a highest state $\sigma_{j'}$ in which the buyer is willing to purchase, if the strategy specifies any purchase at all. This state, if any, completely characterizes the buyer's purchase decision starting in state σ_J, as states fall in sequence (a state lower than $\sigma_{j'}$ cannot be reached without first passing through this state itself). As with the Hybrid Mechanism, any stationary strategy that involves purchase then induces (starting in state σ_J) an expected discounting $x_{n_{j'}}$ for some j', with purchase at price $p_{n_{j'}}$. Incentive compatibility for type θ_n then evaluates the willingness to purchase as soon as the price is p^*_n. That is, it requires $U_{n,n} \geq U_{n,n_{j'}}$ for all $n_{j'}$, which is the same incentive constraint as for the static mechanism.

Condition to exactly attain profits Π^*. Now consider why profits Π^* are not exactly attainable when $J > 2$. Achieving total expected profits Π^* would require achieving these profits almost surely for almost every cohort t. This implies that expected discounting to each price $p^*_{n_j}$ must be given almost surely by $x^*_{n_j}$ for almost every cohort t. This is only possible if, with probability one, the first purchase by types θ_{n_2} and θ_{n_1} after date zero are at prices $p^*_{n_2}$ and $p^*_{n_1}$, respectively. Denote the corresponding purchase dates $\tilde{\tau}_{n_2}$ and $\tilde{\tau}_{n_1}$. Expected discounting must satisfy $\mathbb{E}[\tilde{\tau}_{n_2}] = x^*_{n_2}$ and $\mathbb{E}[\tilde{\tau}_{n_1}] = x^*_{n_1}$. Were this not the case, we could find a positive measure of cohorts in a neighborhood of date zero which, with positive probability, do not purchase at prices $p^*_{n_2}$ and $p^*_{n_1}$ for types θ_{n_2}

and θ_{n_1} with expected discounting to purchase of $x^*_{n_2}$ and $x^*_{n_1}$, and therefore profits would be less than Π^*. Our aim will be to show that the expected discounting $\mathbb{E}[\tilde{\tau}_{n_2}] = x^*_{n_2}$ and $\mathbb{E}[\tilde{\tau}_{n_1}] = x^*_{n_1}$ is, nonetheless, incompatible with obtaining total profits Π^*.

We will use that incentive compatibility requires $\tilde{\tau}_{n_1} > \tilde{\tau}_{n_2}$ almost surely. This is because otherwise type θ_{n_2} can earn a higher payoff by purchasing at $p^*_{n_1} = \theta_{n_1} < p^*_{n_2}$ whenever this price is offered first. Let K_2 be the event that $\tilde{\tau}_{n_2} < \infty$. A consequence of the previous claim is then that $\Pr(\tilde{\tau}_{n_1} - \tilde{\tau}_{n_2} < \varepsilon | K_2) \to 0$ as ε tends to zero.

Note that attaining total expected profits Π^* requires that, with probability one, almost all cohorts arriving after $\tilde{\tau}_{n_2}$ generate expected profits Π^*. This requires, in particular, that expected discounting to $\tilde{\tau}_{n_1}$ is almost surely equal to $x^*_{n_1}$. This implies that we must have $\mathbb{E}_{\tilde{\tau}_{n_1}}\left[e^{-r(\tilde{\tau}_{n_1}-\tilde{\tau}_{n_2})} | \tilde{\tau}_{n_2}\right] = x^*_{n_1}$ almost surely on K_2. Otherwise we would have that, for a positive measure of cohorts immediately following $\tilde{\tau}_{n_2}$, with positive probability, the expected discounting to date $\tilde{\tau}_{n_1}$ would differ from $x^*_{n_1}$. This conclusion can be obtained using that $\Pr(\tilde{\tau}_{n_1} - \tilde{\tau}_{n_2} < \varepsilon | K_2) \to 0$ as ε tends to zero, as noted above.

Now, by the law of iterated expectations, we have

$$\begin{aligned}
\mathbb{E}_{\tilde{\tau}_{n_1}}\left[e^{-r\tilde{\tau}_{n_1}}\right] &= \Pr(K_2)\mathbb{E}_{\tilde{\tau}_{n_1},\tilde{\tau}_{n_2}}\left[e^{-r\tilde{\tau}_{n_2}} e^{-r(\tilde{\tau}_{n_1}-\tilde{\tau}_{n_2})} \big| K_2\right] \\
&= \Pr(K_2)\mathbb{E}_{\tilde{\tau}_{n_2}}\left[e^{-r\tilde{\tau}_{n_2}} \mathbb{E}_{\tilde{\tau}_{n_1}}\left[e^{-r(\tilde{\tau}_{n_1}-\tilde{\tau}_{n_2})} \big| \tilde{\tau}_{n_2}\right] \big| K_2\right] \\
&= \mathbb{E}_{\tilde{\tau}_{n_2}}\left[e^{-r\tilde{\tau}_{n_2}}\right] x^*_{n_1} \\
&= x^*_{n_2} x^*_{n_1},
\end{aligned}$$

which is strictly less than $x^*_{n_1}$. This indeed violates that expected discounting to $\tilde{\tau}_{n_1}$ is $x^*_{n_1}$, which is what we wanted to show.

□

Proof of Proposition 5

Proof. **Sufficiency.** Consider now a high type arriving at t such that Conditions (2) and (3) hold. At any date $s > t$ such that he has not yet purchased and there has not been a discount in $[t, s]$, if the price at time s is p^*_2, he obtains a payoff $v_2(p^*_2)$ from buying immediately. By Proposition 1, this payoff is equal to $x^*_1 v_2(\theta_1)$. By instead delaying and purchasing at the next price discount, he expects the weakly larger payoff

$$\mathbb{E}\left[e^{-r(\tilde{\tau}^t_1 - s)} \big| \tilde{\tau}^t_1 > s\right] v_2(\theta_1).$$

Hence, if the buyer elects not to purchase upon arrival at date t, his payoff from purchasing at some time $s > t$ (given that no price discount occurs in $[t, s]$) is no greater than by purchasing at the next price discount. Given Condition (2), it is then incentive compatible for the buyer to purchase on arrival at date t.

Necessity. Now, consider a t such that Condition (2) holds while Condition (3) fails. For any $s > t$, let $K_{t,s}$ denote the event that $\tilde{\tau}_1^t > s$, and let $K'_{t,s}$ be its complement. Then there is $s > t$ such that $\Pr(K_{t,s}) > 0$ and $\mathbb{E}\big[e^{-r(\tilde{\tau}_1^t - s)}\big|K_{t,s}\big] < x_1^*$.

Since $\mathbb{E}\big[e^{-r(\tilde{\tau}_1^t - t)}\big] = x_1^*$, we have

$$x_1^* = (1 - \Pr(K_{t,s}))\mathbb{E}\big[e^{-r(\tilde{\tau}_1^t - t)}\big|K'_{t,s}\big] + \Pr(K_{t,s})e^{-r(s-t)}\mathbb{E}\big[e^{-r(\tilde{\tau}_1^t - s)}\big|K_{t,s}\big].$$

So, necessarily,

$$\big(1 - \Pr(K_{t,s})\big)\mathbb{E}\big[e^{-r(\tilde{\tau}_1^t - t)}\big|K'_{t,s}\big] > \big(1 - \Pr(K_{t,s})e^{-r(s-t)}\big)x_1^*. \tag{6}$$

The payoff of a high type arriving at t and purchasing at the next price discount or at date s, whichever comes first, is

$$(1 - \Pr(K_{t,s}))\mathbb{E}\big[e^{-r(\tilde{\tau}_1^t - t)}\big|K'_{t,s}\big]v_2(\theta_1) + \Pr(K_{t,s})e^{-r(s-t)}v_2(p_2^*) \tag{7}$$

The first term of Equation (7) is strictly greater than $\big(1 - \Pr(K_{t,s})e^{-r(s-t)}\big)x_1^* v_2(\theta_1)$ by the previous inequality (i.e., Equation (6)), while the second term is equal to $\Pr(K_{t,s})e^{-r(s-t)}x_1^* v_2(\theta_1)$. Therefore, the expression in Equation (7) is strictly greater than $x_1^* v_2(\theta_1)$, which is equal (by Lemma 1) to $v_2(p_2^*)$. This shows that purchasing immediately with probability one gives the buyer a strictly lower payoff than waiting and purchasing at the next discount date, or at date s, whichever comes first. In particular, immediate purchase at date t is not incentive compatible. $\square$

References

Bansal, M., and C. Maglaras, 2009, "Dynamic pricing when customers strategically time their purchase: Asymptotic optimality of a two-price policy," *Journal of Revenue and Pricing management*, 8(1), 42–66.

Berck, P., J. Brown, J. M. Perloff, and S. B. Villas-Boas, 2008, "Sales: tests of theories on causality and timing," *International Journal of Industrial Organization*, 26(6), 1257–1273.

Board, S., 2008, "Durable-goods monopoly with varying demand," *The Review of Economic Studies*, 75(2), 391–413.

Board, S., and A. Skrzypacz, 2016, "Revenue management with forward-looking buyers," *Journal of Political Economy*, 124(4), 1046–1087.

Budish, E. B., and L. N. Takeyama, 2001, "Buy prices in online auctions: irrationality on the internet?," *Economics letters*, 72(3), 325–333.

Chen, T., A. Kalra, and B. Sun, 2009, "Why do consumers buy extended service contracts?," *Journal of Consumer Research*, 36(4), 611–623.

Chevalier, J., and A. Goolsbee, 2009, "Are durable goods consumers forward-looking? Evidence from college textbooks," *The Quarterly Journal of Economics*, 124(4), 1853–1884.

Chevalier, J. A., and A. K. Kashyap, 2017, "Best Prices: Price Discrimination and Consumer Substitution," working paper, National Bureau of Economic Research.

Conlisk, J., E. Gerstner, and J. Sobel, 1984, "Cyclic pricing by a durable goods monopolist," *The Quarterly Journal of Economics*, 99(3), 489–505.

Correa, J., J. F. Escobar, and A. Perlroth, 2019, "Revenue maximization with heterogeneous discounting: auctions and pricing," .

Dilmé, F., and F. Li, 2019, "Revenue Management without Commitment: Dynamic Pricing and Periodic Flash Sales," *The Review of Economic Studies*, 86(5), 1999–2034.

Eichenbaum, M., N. Jaimovich, and S. Rebelo, 2011, "Reference prices, costs, and nominal rigidities," *American Economic Review*, 101(1), 234–62.

Fershtman, C., and A. Fishman, 1992, "Price cycles and booms: dynamic search equilibrium," *The American Economic Review*, pp. 1221–1233.

Février, P., and L. Wilner, 2016, "Do consumers correctly expect price reductions? Testing dynamic behavior," *International Journal of Industrial Organization*, 44, 25–40.

Garrett, D. F., 2016, "Intertemporal price discrimination: Dynamic arrivals and changing values," *American Economic Review*, 106(11), 3275–99.

Gershkov, A., B. Moldovanu, and P. Strack, 2017, "Revenue-Maximizing Mechanisms with Strategic Customers and Unknown, Markovian Demand," *Management Science*, 64(5), 2031–2046.

Gowrisankaran, G., and M. Rysman, 2012, "Dynamics of consumer demand for new durable goods," *Journal of political Economy*, 120(6), 1173–1219.

Heidhues, P., and B. Kőszegi, 2014, "Regular prices and sales," *Theoretical Economics*, 9(1), 217–251.

Hendel, I., and A. Nevo, 2013, "Intertemporal price discrimination in storable goods markets," *American Economic Review*, 103(7), 2722–51.

Hörner, J., and L. Samuelson, 2011, "Managing strategic buyers," *Journal of Political Economy*, 119(3), 379–425.

Kehoe, P., and V. Midrigan, 2015, "Prices are sticky after all," *Journal of Monetary Economics*, 75, 35–53.

Lan, H., T. Lloyd, and W. Morgan, 2016, "Price Reductions Predictable? The Hazard Function of Sales," *mimeo U. of Bournemouth, U. of East Anglia and U. of Sheffield.*

Landsberger, M., and I. Meilijson, 1985, "Intertemporal price discrimination and sales strategy under incomplete information," *The RAND Journal of Economics*, pp. 424–430.

Liu, Q., and G. J. van Ryzin, 2008, "Strategic capacity rationing to induce early purchases," *Management Science*, 54(6), 1115–1131.

Maskin, E., and J. Riley, 1984, "Optimal auctions with risk averse buyers," *Econometrica: Journal of the Econometric Society*, pp. 1473–1518.

Matthews, S. A., 1983, "Selling to risk averse buyers with unobservable tastes," *Journal of Economic Theory*, 30(2), 370–400.

Moore, J., 1984, "Global incentive constraints in auction design," *Econometrica: Journal of the Econometric Society*, pp. 1523–1535.

Nakamura, E., and J. Steinsson, 2008, "Five facts about prices: A reevaluation of menu cost models," *The Quarterly Journal of Economics*, 123(4), 1415–1464.

Ortner, J., 2017, "Durable goods monopoly with stochastic costs," *Theoretical Economics*, 12(2), 817–861.

Öry, A., 2017, "Consumers on a Leash: Advertised Sales and Intertemporal Price Discrimination," *mimeo Yale University*.

Pesendorfer, M., 2002, "Retail sales: A study of pricing behavior in supermarkets," *The Journal of Business*, 75(1), 33–66.

Rabin, M., 2000, "Diminishing marginal utility of wealth cannot explain risk aversion," in *The Oxford Handbook of Innovation*, ed. by D. Kahneman, and A. Tversky. Cambridge University Press, Cambridge, pp. 202–208.

Reynolds, S. S., and J. Wooders, 2009, "Auctions with a buy price," *Economic theory*, 38(1), 9–39.

Riley, J., and R. Zeckhauser, 1983, "Optimal selling strategies: When to haggle, when to hold firm," *The Quarterly Journal of Economics*, 98(2), 267–289.

Rochet, J.-C., 1985, "The taxation principle and multi-time Hamilton-Jacobi equations," *Journal of Mathematical Economics*, 14(2), 113–128.

Ronayne, D., and D. P. Myatt, 2019, "A Theory of Stable Price Dispersion," working paper.

Shilony, Y., 1977, "Mixed pricing in oligopoly," *Journal of Economic Theory*, 14(2), 373–388.

Sobel, J., 1991, "Durable goods monopoly with entry of new consumers," *Econometrica: Journal of the Econometric Society*, pp. 1455–1485.

Stokey, N. L., 1979, "Intertemporal price discrimination," *The Quarterly Journal of Economics*, pp. 355–371.

Varian, H. R., 1980, "A model of sales," *The American Economic Review*, 70(4), 651–659.

Warner, E. J., and R. B. Barsky, 1995, "The timing and magnitude of retail store markdowns: evidence from weekends and holidays," *The Quarterly Journal of Economics*, 110(2), 321–352.

B Appendix II: A non-existence proof

The objective of this Appendix is to formally show that there is no price process achieving the static optimal profits Π^* whenever there exist some n and m such that $0 < x_m^* < x_n^* < 1$. We first introduce incentive-compatible price processes and their optimality. Then we state and show the result.

Let $\langle \Omega, (\mathcal{F}_t)_{t\in[0,\infty)}, \mathcal{P}\rangle$ be some filtered probability space.

Definition 1. *A price process is a pair (P, τ) satisfying that:*

1. *P is a stochastic process defined on $\langle \Omega, (\mathcal{F}_t)_{t\in[0,\infty)}, \mathcal{P}\rangle$.*
2. *For each pair (n, t), $\tau_{n,t}$ is a stopping time predictable with respect to the filtration generated by P and satisfying that $\tau_{n,t}(\omega) \geq t$ for all $\omega \in \Omega$.*[22]

Note that this definition specifies not only a stochastic process for prices, but also the stopping time $\tau_{n,t}$ which should be interpreted as the purchase time of a buyer of type θ_n who arrives to the market at date t. Incentive compatibility of a price process is defined as follows.

Definition 2. *We say that (P, τ) is* incentive compatible *if, for all n and all t,*

$$\mathbb{E}[e^{-r(\hat{\tau}-t)}v_n(P_{\hat{\tau}})] \leq \mathbb{E}[e^{-r(\tau_{n,t}-t)}v_n(P_{\tau_{n,t}})]$$

for all stopping times $\hat{\tau}$ predictable with respect to the filtration generated by P satisfying that $\hat{\tau}(\omega) \geq t$ for all $\omega \in \Omega$.

Optimality is then defined as follows.

Definition 3. *We say that (P, τ) is* optimal *if it maximizes*

$$\int_0^\infty \sum_{n=1}^N \beta_n \mathbb{E}[e^{-r\hat{\tau}_{n,t}} \hat{P}_{\hat{\tau}_{n,t}}] r\mathrm{d}t$$

among all incentive compatible price processes $(\hat{P}, \hat{\tau})$.

Define $\Pi_t \equiv \sum_{n=1}^N \beta_n e^{-r(\tau_{n,t}-t)} P_{\tau_{n,t}}$ to be the realization of profits from a buyer arriving at date t. Recall that expected profits for any cohort are bounded above by the optimal static profits Π^*. Therefore, the principal's discounted profits are equal to $\int_0^\infty \mathbb{E}[\Pi_t] e^{-rt} r\mathrm{d}t \leq \Pi^*$. The result to be shown is then the following.

Proposition 6. *If there exist n and m such that $0 < x_m^* < x_n^* < 1$, there is no incentive compatible price process giving Π^* to the principal.*

Proof. The proof will be by contradiction. From now on, we assume with a view to contradiction that there is a price process (P, τ) that gives the principal a payoff equal to Π^*; that is, $\int_0^\infty \mathbb{E}[\Pi_t] e^{-rt} r\mathrm{d}t = \Pi^*$. We will assume, without loss of generality, that $(\mathcal{F}_t)_t$ is the filtration generated by P. We divide the proof into three steps:

[22] A stopping time may be finite or infinite valued. A value $\tau_{n,t}(\omega) = \infty$ indicates that the buyer of type θ_n does not purchase.

Step 1: In this step, we define the set $\mathcal{T}$ of dates t such that, for some n, either (i) there is $A_t \in \mathcal{F}_t$ with $\mathcal{P}(A_t) > 0$ such that

$$\mathbb{E}[e^{-r(\tau_{n,t}-t)}|A_t] \neq x_n^* \tag{8}$$

or (ii) there is $A_t \in \mathcal{F}_t$ with $\mathcal{P}(A_t) > 0$ such that $\mathcal{P}(P_{\tau_{n,t}} = p_n^*|A_t) < 1$. The result to be established is that the set $\mathcal{T}$ has Lebesgue measure zero. To see this, note that by Proposition 1 (and the equivalence between allocation probabilities of the static mechanism and expected discounting in a dynamic setting), for dates $t \in \mathcal{T}$ we can find $A_t \in \mathcal{F}_t$ (chosen so that $\mathbb{E}[e^{-r(\tau_{n,t}-t)}|A_t] \neq x_n^*$ or $\mathcal{P}(P_{\tau_{n,t}} = p_n^*|A_t) < 1$ for some n) for which $\mathbb{E}[\Pi_t|A_t] < \Pi^*$. In addition, we know that, for all t, and any $B \in \mathcal{F}_t$, $\mathbb{E}[\Pi_t|B] \leq \Pi^*$. Therefore, taking $B = \Omega \backslash A_t$, we obtain

$$\mathbb{E}[\Pi_t] = \mathcal{P}(A_t)\mathbb{E}[\Pi_t|A_t] + \mathcal{P}(\Omega\backslash A_t)\mathbb{E}[\Pi_t|\Omega\backslash A_t] < \Pi^*$$

for all $t \in \mathcal{T}$. It follows that, if the integral defining the seller's discounted payoff $\int_0^\infty \mathbb{E}[\Pi_t]e^{-rt}r dt$ is well-defined, it is strictly less than Π^*. This contradicts our assumption that the seller obtains Π^*.

We then observe that, for any $t \notin \mathcal{T}$, any $A \in \mathcal{F}_t$ with $\mathcal{P}(A) > 0$, we have

$$\mathcal{P}\big(\,\tau_{n,t} < \tau_{m,t} \text{ or } \tau_{n,t} = \tau_{m,t} = \infty \,\big|\, A\,\big) = 1$$

for all n and all m such that $x_n^* > x_m^*$. This is immediate when $x_m^* = 0$, since then $\mathcal{P}\left(\tau_{m,t} = \infty\right) = 1$. If instead $x_m^* > 0$, we argue that $\mathcal{P}(\,\tau_{n,t} < \tau_{m,t}$ or $\tau_{n,t} = \tau_{m,t} = \infty \,|\, A\,) < 1$ would imply that the type θ_n buyer has a profitable deviation to stopping time $\tau_{n,t} \wedge \tau_{m,t}$. Higher profits under this stopping time can be explained by observing that either there is positive probability that the buyer purchases at price $p_m^* < \theta_n$ whereas he does not purchase under $\tau_{n,t}$, or there is a positive probability that the buyer purchases earlier and hence at price p_m^* rather than at p_n^*, where $p_m^* < p_n^*$.

Finally note that there is no loss in profits for the seller if we assume that all buyers with the same type play the same continuation strategy. That is, $\tau_{n,t'}(\omega) = \tau_{n,t}(\omega)$ for all $t' \in [t, \tau_{n,t}(\omega)]$.

Step 2: In this step we introduce the following notation. For any $t, t', t'' \in \mathbb{R}_+$, any n, let $K_{n,t}^{t',t''} \equiv \{\omega | t' \leq \tau_{n,t}(\omega) < t''\}$, and let $K_{n,t}^{t',\infty} \equiv \{\omega | t' \leq \tau_{n,t}(\omega)\}$. We then show the following result.

Lemma 3. *Fix some n and m, with $n > m$ and with $x_n^* \in (0,1)$. Then, for $t, t', t'' \in \mathbb{R}_+ \backslash \mathcal{T}$ with $t \leq t' < t'' \leq \infty$, we have $\mathcal{P}(K_{n,t}^{t',t''}) > 0$ and*

$$\frac{x_m^*}{x_n^*} = \frac{\mathbb{E}\left[e^{-r(\tau_{m,t}-t')}\middle|K_{n,t}^{t',t''}\right]}{\mathbb{E}\left[e^{-r(\tau_{n,t}-t')}\middle|K_{n,t}^{t',t''}\right]}\ . \tag{9}$$

Proof. We first want to show that, for $t, t', t'' \in \mathbb{R}_+ \backslash \mathcal{T}$ with $t \leq t' < t'' \leq \infty$, $\mathcal{P}(K_{n,t}^{t',t''}) > 0$. This will be a result of what we call Claim A: For $t, t_1, t_2 \in \mathbb{R}_+ \backslash \mathcal{T}$ with $t \leq t_1 < t_2 < t_1 - \log(x_n^*)/r$ and with $\mathcal{P}(K_{n,t}^{t_1,\infty}) > 0$, we have $\mathcal{P}(K_{n,t}^{t_1,t_2}), \mathcal{P}(K_{n,t}^{t_2,\infty}) > 0$. Given $t, t', t'' \in \mathbb{R}_+ \backslash \mathcal{T}$ with $t \leq t' < t'' \leq \infty$, we then arrive at $\mathcal{P}(K_{n,t}^{t',t''}) > 0$ by applying Claim A iteratively along a sequence of dates $((t_1^{(i)}, t_2^{(i)}))_{i \in \mathbb{N}}$, requiring $t_1^{(1)} = t$, $t_1^{(i)} = t_2^{(i-1)}$ for all $i > 1$, and $t_1^{(i)} < t_2^{(i)} < t_1^{(i)} - \log(x_n^*)/r$ for all i. The first iteration is with $t_1 = t_1^{(1)}$ (observing that then $K_{n,t}^{t_1,\infty} = \Omega$) and $t_2 = t_2^{(1)}$; then the i^{th} iteration is with $t_1 = t_1^{(i)}$ and $t_2 = t_2^{(i)}$. This establishes that each event $K_{n,t}^{t_1^{(i)},t_2^{(i)}}$ has strictly positive probability.

To show Claim A, consider any $t_1, t_2 \in \mathbb{R}_+ \backslash \mathcal{T}$ with $t \leq t_1 < t_2 < t_1 - \log(x_n^*)/r$ and with $\mathcal{P}(K_{n,t}^{t_1,\infty}) > 0$. Applying Step 1, after noting $K_{n,t}^{t_1,\infty} \in \mathcal{F}_{t_1}$, we have:[23]

$$
\begin{aligned}
x_n^* &= \mathbb{E}[e^{-r(\tau_{n,t}-t_1)} | K_{n,t}^{t_1,\infty}] \\
&= \frac{\mathcal{P}(K_{n,t}^{t_1,t_2})}{\mathcal{P}(K_{n,t}^{t_1,\infty})} \underbrace{\mathbb{E}\big[e^{-r(\tau_{n,t}-t_1)} \big| K_{n,t}^{t_1,t_2}\big]}_{(*)} + \frac{\mathcal{P}(K_{n,t}^{t_2,\infty})}{\mathcal{P}(K_{n,t}^{t_1,\infty})} \underbrace{e^{-r(t_2-t_1)}}_{(**)} \underbrace{\mathbb{E}\big[e^{-r(\tau_{n,t}-t_2)} \big| K_{n,t}^{t_2,\infty}\big]}_{(***)} \ .
\end{aligned} \tag{10}
$$

Note first that, if $\mathcal{P}(K_{n,t}^{t_1,t_2}) > 0$, then the term $(*)$ is no smaller than $e^{-r(t_2-t_1)}$. It then follows that, because $t_2 < t_1 - \log(x_n^*)/r$, $\mathcal{P}(K_{n,t}^{t_1,t_2})/\mathcal{P}(K_{n,t}^{t_1,\infty}) < 1$ and hence $\mathcal{P}(K_{n,t}^{t_2,\infty}) > 0$. Note also that, since $K_{n,t}^{t_2,\infty} \in \mathcal{F}_{t_2}$ and since $\tau_{n,t}(\omega) = \tau_{n,t_2}(\omega)$ for all $\omega \in K_{n,t}^{t_2,\infty}$, it follows that $(***)$ is equal to x_n^* whenever $\mathcal{P}(K_{n,t}^{t_2,\infty}) > 0$ (from Step 1). Furthermore, given $t_2 > t_1$, we have that $(**)$ is strictly smaller than 1; hence it must be that $\mathcal{P}(K_{n,t}^{t_1,t_2}) > 0$. We conclude that the probabilities of both $K_{n,t}^{t_1,t_2}$ and $K_{n,t}^{t_2,\infty}$ are strictly positive, which establishes the claim.

To establish the lemma, let then $t, t', t'' \in \mathbb{R}_+ \backslash \mathcal{T}$ with $t \leq t' < t'' \leq \infty$. Observe that

$$
\begin{aligned}
x_m^* &= \mathbb{E}[e^{-r(\tau_{m,t'}-t')} | K_{n,t}^{t',\infty}] \\
&= \mathbb{E}[e^{-r(\tau_{m,t}-t')} | K_{n,t}^{t',\infty}] \\
&= \frac{\mathcal{P}(K_{n,t}^{t',t''})}{\mathcal{P}(K_{n,t}^{t',\infty})} \mathbb{E}\big[e^{-r(\tau_{m,t}-t')} \big| K_{n,t}^{t',t''}\big] + \frac{\mathcal{P}(K_{n,t}^{t'',\infty})}{\mathcal{P}(K_{n,t}^{t',\infty})} e^{-r(t''-t')} \underbrace{\mathbb{E}\big[e^{-r(\tau_{m,t}-t'')} \big| K_{n,t}^{t'',\infty}\big]}_{(**)} \ .
\end{aligned} \tag{11}
$$

The second equality holds because $\tau_{m,t}(\omega) \geq t'$ for almost all $\omega \in K_{n,t}^{t',\infty}$ (since $\mathcal{P}(\tau_{n,t} < \tau_{m,t}$ or $\tau_{n,t} = \tau_{m,t} = \infty) = 1$), and so $\tau_{m,t}(\omega) = \tau_{m,t'}(\omega)$ for almost all $\omega \in K_{n,t}^{t',\infty}$. Since $K_{n,t}^{t'',\infty} \in \mathcal{F}_{t''}$, and since $\tau_{m,t} = \tau_{m,t''}$ on $K_{n,t}^{t'',\infty}$, we have that $(**)$ is equal to x_m^*. Considering Equation (11) for distinct m and n as in the lemma, as well as m taken equal to n, generates two equations which together imply Equation (9). $\square$

(End of proof of Lemma 3, proof of Proposition 6 continues.)

Step 3: We now assume that $0 < x_m^* < x_n^* < 1$ and conclude the argument. For $t, t', t'' \in \mathbb{R}_+ \backslash \mathcal{T}$ with $t \leq t' < t'' \leq \infty$, we have

$$
\begin{aligned}
\overbrace{\mathbb{E}\big[e^{-r(\tau_{m,t}-t')} \big| K_{n,t}^{t',t''}\big]}^{(*)} = &\mathcal{P}\big(K_{m,t}^{t',t''} \big| K_{n,t}^{t',t''}\big) \mathbb{E}\big[e^{-r(\tau_{m,t}-t')} \big| K_{n,t}^{t',t''} \cap K_{m,t}^{t',t''}\big] \\
&+ e^{-r(t''-t')} \mathcal{P}\big(\overline{K}_{m,t}^{t',t''} \big| K_{n,t}^{t',t''}\big) \underbrace{\mathbb{E}\big[e^{-r(\tau_{m,t}-t'')} \big| K_{n,t}^{t',t''} \cap \overline{K}_{m,t}^{t',t''}\big]}_{(**)} \ .
\end{aligned}
$$

Since $K_{n,t}^{t',t''} \cap \overline{K}_{m,t}^{t',t''} \in \mathcal{F}_{t''}$, we have that $(**)$ is equal to x_m^*. Now, let $\bar{\delta} = \frac{-1}{r} \log\left(\frac{1}{2}x_n^* + \frac{1}{2}\right)$, and suppose additionally that $t'' - t' < \bar{\delta}$ so that $e^{-r(t''-t')} > \frac{1}{2}x_n^* + \frac{1}{2}$. Then, using Equation (9) (to replace $(*)$) and that

[23] Note that, in expressions such as equation (10), if the conditioning event for a conditional expectation has probability zero, we take the conditional expectation to equal zero.

$\mathcal{P}\big(\overline{K}^{t',t''}_{m,t}\big|K^{t',t''}_{n,t}\big) = 1-\mathcal{P}\big(K^{t',t''}_{m,t}\big|K^{t',t''}_{n,t}\big)$, we obtain

$$\mathcal{P}\big(K^{t',t''}_{m,t}\big|K^{t',t''}_{n,t}\big) = \frac{\frac{x_m^*}{x_n^*}\overbrace{\mathbb{E}\big[e^{-r(\tau_{n,t}-t')}\big|K^{t',t''}_{n,t}\big]}^{\geq e^{-r(t''-t')}} - \overbrace{e^{-r(t''-t')}}^{\leq 1} x_m^*}{\underbrace{\mathbb{E}\big[e^{-r(\tau_{m,t}-t')}\big|K^{t',t''}_{n,t}\cap K^{t',t''}_{m,t}\big]}_{\leq 1} - e^{-r(t''-t')}x_m^*} \geq \frac{x_m^*}{x_n^*}\frac{e^{-r(t''-t')}-x_n^*}{1-e^{-r(t''-t')}x_m^*}\ .$$

The term on the right-hand side of this inequality is decreasing in $t''-t'$. Using the specification of $\bar{\delta}$, given $t,t',t''\in\mathbb{R}\backslash\mathcal{T}$ with $t''-t'\in(0,\bar{\delta})$, we have

$$\mathcal{P}\big(K^{t',t''}_{m,t}\big|K^{t',t''}_{n,t}\big) \geq \frac{x_m^*(1-x_n^*)}{x_n^*(2-(1+x_n^*)x_m^*)} > 0. \tag{12}$$

Now let $\delta\in(0,\bar{\delta})$ and specify $K^{\delta}_{n,m,t} \equiv \big\{\omega\,:\,\tau_{m,t}(\omega)\in(\tau_{n,t}(\omega),\tau_{n,t}(\omega)+\delta]\big\}$. Note that, for any t, $\lim_{\delta\to 0}\mathcal{P}\big(K^{\delta}_{n,m,t}\big)=0$. Now, pick any $t\in\mathbb{R}\backslash\mathcal{T}$ and choose a strictly increasing sequence $(t_k)_{k=0}^{\infty}$ in $\mathbb{R}\backslash\mathcal{T}$ with $t_0=t$ and such that $t_{k+1}-t_k\in(\delta/2,\delta)$ for all k. Then

$$\begin{aligned}\mathcal{P}\big(K^{\delta}_{n,m,t}|\tau_{n,t}(\omega)<\infty\big) &\geq \sum_{k=0}^{\infty}\mathcal{P}\big(K^{t_k,t_{k+1}}_{n,t}|\tau_{n,t}(\omega)<\infty\big)\mathcal{P}\big(K^{t_k,t_{k+1}}_{m,t}\big|K^{t_k,t_{k+1}}_{n,t}\big)\\ &\geq \frac{x_m^*(1-x_n^*)}{x_n^*(2-(1+x_n^*)x_m^*)}\ .\end{aligned}\tag{13}$$

The first inequality holds because $K^{t_k,t_{k+1}}_{m,t}\cap K^{t_k,t_{k+1}}_{n,t}\subset K^{\delta}_{n,m,t}\cap K^{t_k,t_{k+1}}_{n,t}$. The second inequality follows from Equation (12). Considering Equation (13) as $\delta\to 0$, we obtain a contradiction to the previous observation that $\lim_{\delta\to 0}\mathcal{P}\big(K^{\delta}_{n,m,t}\big)=0$.. $\square$