A Dynamic Theory of Random Price Discounts

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

23rd September 2022

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 et al. (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 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

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 agreements No 714147 and No 949465). Dilmé thanks the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for research support through grant CRC TR 224 (Project B02) and under Germany's Excellence Strategy – EXC 2126/1–390838866 and EXC 2047 – 390685813.

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

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

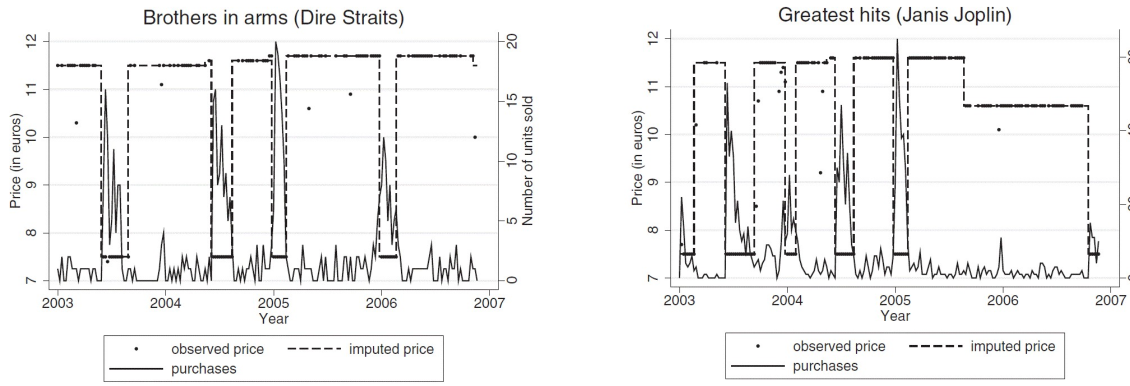

(a) (b)

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).

that the seller gradually reduces prices as the date with the steepest discount draws near. Stokey (1979), Conlisk et al. (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 abrupt rather than gradual and that purchases do not decline immediately before sizeable price reductions (see Figure 1). Février and Wilner 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] This contrasts with the optimality of constant prices in important benchmarks with *risk-neutral* buyers (see Stokey, 1979, and Conlisk et al., 1984). That our approach assumes full commitment is in contrast to the received work on Coasian dynamics, but is in line with a number of other papers studying intertemporal price discrimination in durable goods markets.[2]

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

[2]Some of these papers are reviewed in Section 9 at the end of the paper. 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

While risk aversion has been studied in other allocation problems such as auctions, its role has been given less attention in relation to dynamic pricing. Risk aversion has often been observed for situations with small stakes and we review some instances in the literature in Section 7. A prominent interpretation is that small-stakes risk aversion is reflective of agent loss aversion and Section 7 discusses how we can adapt our theory when viewing buyers as loss averse. There is a range of evidence of small-scale risk aversion, some of which relates to durable goods markets. One example is the sale of warranties for electronic goods at worse than fair prices (see Chen et al., 2009). Possibly another is the presence of "buy-it-now" prices in online auctions on platforms such as Yahoo and eBay which has often been associated with buyer risk aversion (see Budish and Takeyama, 2001, and Reynolds and Wooders, 2009).

The seller's problem is to choose the price path offered to buyers who arrive over time. We show that there is a virtually optimal price path involving a constant regular price, with short-lived episodes of discounting that are randomly timed, and which buyers find unpredictable.[3] 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, independently of their arrival time. For instance, all highest valuation buyers have the same incentive to accept the constant regular price instead of waiting for lower ones because the arrival of discounting episodes is history independent under the optimal policy. Similarly, buyers with intermediate values purchase at intermediate prices within discounting episodes because delaying purchase to obtain a lower price involves the risk that the discounting episode ends and the price returns to the regular level. Buyers with the lowest values obtain no rents and only buy if the price reaches their valuation in a discounting episode. Hence, 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), 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

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."

[3]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.

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 further 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. In particular, in our baseline model where buyers observe past prices, the prediction is that price discounts arrive at a Poisson rate, i.e. there is a constant hazard rate for price discounts. Several empirical studies instead find an increasing hazard rate: price discounts become more likely the longer since the last discount. This makes price discounts somewhat predictable to a buyer with access to historical price data. We show however that, when buyers only observe prices after their arrival to the market, there is a range of optimal price processes, with this range determined by an incentive compatibility condition for buyers. This condition requires that buyers do not become more pessimistic about the arrival of new discounts the longer they wait. Our result on random discounting when buyers do not observe prices before arrival therefore offers a possible reconciliation with empirical observations on the timing of discounts. The analysis is therefore relevant to empirical investigations of the topic in the macroeconomics literature on price stickiness (see, e.g., Nakamura and Steinsson, 2008, Eichenbaum et al., 2011, and Kehoe and Midrigan, 2015), and in industrial organization (e.g., Pesendorfer, 2002, Berck et al., 2008, and Février and Wilner, 2016).

The rest of the paper then examines in different ways the scope of our theory. Section 6 examines the extent to which random allocations are optimal in settings with many buyer types, and so provides conditions under which our theory supports the use of random price discounts. As noted above, Section 7 explores how our theory can be extended to loss-averse buyer preferences. Section 8 discusses the relevance of our theory for settings with more than one seller. Section 9 reviews other relevant theories of price discounts.

2 Set-up

Our principal interest is in a dynamic setting with a single seller and a continuum of buyers who arrive over time, and we describe this model here. In order to characterize pricing in this dynamic environment (see Section 4), Section 3 will first consider a static model with a single buyer.

Buyers in our model have unit demand and are risk averse. The seller faces no capacity constraints, zero production costs, and is risk neutral. In the dynamic setting of interest, 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 $r > 0$. Buyers are taken to arrive to the market at a fixed rate $\gamma > 0$. This is normalized by setting $\int_0^\infty \gamma e^{-rt} dt = 1$ (i.e., $\gamma = r$). This normalization will conveniently imply that the seller's total profits correspond to a per-buyer (weighted) average.[4]

Buyers' enjoyment of the good depends on their "types", labeled $\{\theta_n | n=1,...,N\}$, with $\theta_N > ... > \theta_1 > 0$. Each buyer's type θ_n will represent his willingness to pay for a unit. For each cohort of buyers, a fixed proportion $\beta_n > 0$ has type θ_n, where $\sum_{n=1}^N \beta_n = 1$.

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 $v(p;\theta_n) \in \mathbb{R}$ denote the utility of a purchase for this type and payment. We assume that $v(\theta_n;\theta_n) = 0$ for each θ_n which ensures that types θ_n have the interpretation of willingness to pay.

In our dynamic environment of Section 4, our assumptions imply that a type θ_n buyer's intertemporal payoff is $e^{-rt} v(p_t;\theta_n)$ if the good is purchased at price p_t on date t, while it is equal to zero in case of never purchasing. As an alternative notation, we find it convenient to let $v_n(p) = v(p;\theta_n)$. A natural possibility is to restrict v_n so that buyer types have an equivalent monetary value, that is to set $v_n(p) = u(\theta_n - p)$ for some function u. This possibility corresponds to Case 1 of Maskin and Riley (1984). Note, however, that our results will extend beyond the case where the buyer's type has a monetary interpretation.

We restrict buyer preferences as follows. For each n, $v_n(\cdot)$ is a strictly decreasing, strictly concave, and twice continuously-differentiable function. We also 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.

Before turning to results on static mechanisms, it is useful to summarize the above restrictions on preferences when $v_n(p) = u(\theta_n - p)$. In this case, we require u to be a function $u : \mathbb{R} \to \mathbb{R}$ a that is strictly increasing, strictly concave, twice continuously differentiable, and satisfies $u(0) = 0$. Assumption A1 is then the requirement that $\frac{u'(y)}{u(y)}$ is strictly decreasing in y, which is met without further restrictions given the assumed properties of u.[5] Assumption A2 is the requirement that the coefficient of absolute risk aversion $\frac{-u''(y)}{u'(y)}$ is weakly decreasing in y. An example for u is that specifying CARA preferences, where $u(y) = 1 - e^{-Ry}$ for a coefficient of absolute risk aversion $R > 0$.

[4]While a constant arrival rate is a convenient simplification, all our arguments and results extend also to settings with time-varying arrival rates.

[5]It follows, in particular, because $\frac{\partial}{\partial\theta}\frac{u'(\theta-p)}{u(\theta-p)} = \frac{u''(\theta-p)u(\theta-p) - u'(\theta-p)^2}{u(\theta-p)^2} < 0$ for $\theta > p$.

3 Selling to a risk-averse buyer

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}. As in the dynamic model, the seller has no costs. The buyer has unit demand and a willingness to pay of θ_n with a probability β_n as introduced above. The seller can only require payments when the buyer receives the good.

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 conditional on award 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 at some p_n 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 (shown in Appendix I together with the other proofs) 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.[6] 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.[7] Another point of comparison is the mechanism design analysis by Bansal and Maglaras (2009) in a model where buyers have CRRA utility and only pay if they get the good. A key point of difference beyond our general specification for the utility function is that their work *assumes* each type pays a sure price, rather than deriving the implication. Their results thus leave open that the seller could reach higher profits through a randomization of payments.

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* $x_0^* = p_0^* = 0$*.*

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 buyer'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. We establish this here in the case where $N = 2$, and delay establishing results for many types until Section 6.

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

[7]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.

We refer to θ_2 as the "high type" and θ_1 as the "low 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) with respect to x_1.[8]

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.[9] 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 the change. 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 the change. Indeed, for intermediate values of the probability of the high type (namely $\beta_2 \in (\underline{\beta}, \bar{\beta})$), both the above adjustments are profitable, explaining why the optimal choice of x_1 is interior.

4 Optimal price mechanisms in dynamic arrivals

This section considers dynamic pricing mechanisms, initially for a fixed arrival date and then for the model set-up, as described in Section 2, where buyers arrive 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 per buyer. 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,

[8] Proposition 2 holds without Assumptions A1 and A2. This can be seen in a previous version of the paper.

[9] 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).

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 $t=0$. We argue that 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. The idea behind our price path specification builds on the equivalence in payoffs (both for the seller and any buyer type) between an allocation of the good with probability $x_n^* > 0$ at price p_n^* without delay (as in the static mechanism), and an allocation with probability one at the same price after a delay t_n^* for an appropriately determined t_n^*. Here, t_n^* is obtained by setting the time discounting $e^{-rt_n^*}$ equal to the probability of allocation x_n^*; that is, $t_n^* = -\log(x_n^*)/r \geq 0$. This observation will mean there are payoff-equivalent dynamic formats where the buyer makes any purchase at dates in $\mathscr{T} = \{-\log(x_n^*)/r : x_n^* > 0\}$, with the x_n^* the allocation probabilities available to the buyer in the static mechanism.

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 so high as to be irrelevant for the buyer's problem; for instance, it is enough to set them to p_N^* (the highest payment in the static mechanism). Prices fall over dates in $\mathscr{T}$. If the buyer has type θ_n, he waits for the price to fall to p_n^* at date t_n^* to purchase. 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 $\mathscr{T}$, as well as the prices at these dates, are unique across any optimal deterministic price path.

The argument above derives a dynamic price path for a buyer with a fixed arrival date such that profits equal those in the optimal static mechanism. Our argument for the equivalence in profits hinges on the results in the previous section. In particular, it relies on Lemma 2 establishing that it is profit maximizing for each type to make a sure payment. Suppose that in the static mechanism a type θ_n were to receive the good, say with probability $x_n^* = 1$, making instead a *random* payment which has a low realization with positive probability. If we replicate the logic of the previous paragraph, then we specify a price path that sells to θ_n at this low price at date $t_n^* = 0$ with the same probability. When the low price is realized, types who according to the above argument will be specified to purchase at later dates if they receive lower allocation probabilities in the static mechanism, may not find it incentive compatible to wait. This suggests that the extent to which dynamic price paths can generate the same payoffs as mechanisms with random payments is a more difficult question even when the buyer arrives at a fixed date.

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. Now consider the arrival process in the model set-up, which introduces an inflow of infinitesimal buyers at a constant rate normalized to r. Note that, as for the single-buyer case, the value Π^* is 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.[10] 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 $\mathscr{T}$ at the prices described above. However, using the uniqueness of the optimal static mechanism (see Proposition 1), 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 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 stochastic dynamic mechanism in which the seller obtains expected profits Π^*. Because this will be 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 10).

Let $\theta_{\bar{n}}$ then be the highest type for which $x_{\bar{n}}^* < 1$. The mechanism we examine sets 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 at random times to play static and memoryless revelation mechanisms. 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.[11]

In defining 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. Let revelation mechanisms occur at a Poisson rate $\lambda_{\bar{n}} = \frac{r x_{\bar{n}}^*}{1-x_{\bar{n}}^*}$. 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 $\theta_n \le \theta_{\bar{n}}$ who continues to report truthfully in the revelation mechanisms receives the good at a Poisson rate $\lambda_n = \frac{r x_n^*}{1-x_n^*}$, and so the expected discounting until purchase is x_n^*.[12] 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 makes the same report θ_k in every revelation mechanism. The payoff from

[10] 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.

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

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

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 $x_k^* v_n(p_k^*)$. Verifying incentive compatibility and individual rationality of the behavior prescribed for the buyer above involves checking the same inequalities as 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 $x_0^* \equiv 0$.[13] 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 3. *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 3 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 $p_{\bar{n}}^*$ (so the expected discounting until $p_{\bar{n}}^*$ is offered is $x_{\bar{n}}^*$). 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 one. 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 Hybrid Mechanism. 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

[13] 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.

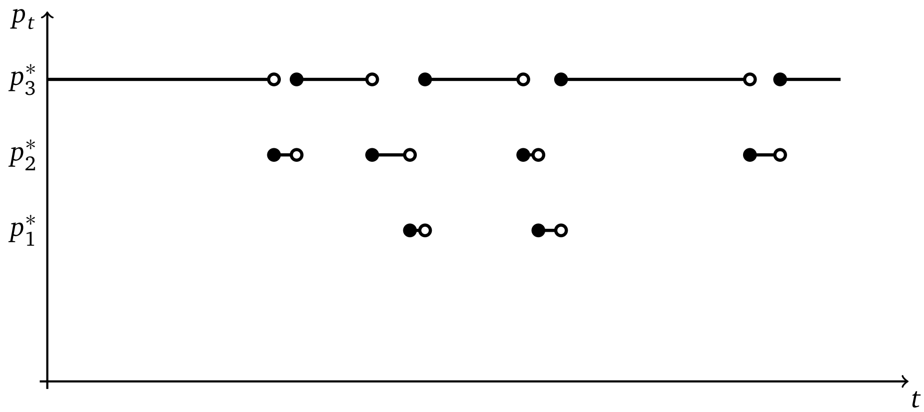

Figure 2: Example of a realization of a price path for $N=3$ when $0 < x_1^* < x_2^* < x_3^* = 1$. In the example, there are four discount episodes from the regular price p_3^*. In the first and the last, only p_2^* is offered before p_3^* is offered again. In the other two discount episodes a second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

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. The expected profits per buyer from cohorts arriving at these times equal Π^*. The only loss in profits 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 3 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 et al. (2022) 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. From the uniqueness of the static mechanism (Proposition 1), and since the optimal profits are achievable through Poisson discounting, we have that almost all high types must purchase immediately at price p_2^* and almost all low types must purchase at price θ_1 with expected discount x_1^* in an optimal price process. We consider for the rest of the section price processes that involve the price being constant at p_2^* except at discount dates where the price is θ_1 and which are determined according to a simple point process. Then, for almost all t, we must have

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

where $\tilde{\tau}_1^t$ is the date of the next discount after time t.

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 timing of price discounting, which we state next.

Proposition 4. *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 in the first place. 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 time of the first 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 Many types

We now consider settings with many types. This will permit us to relate to previous work on auctions and intertemporal price discrimination. Also, we will show that the result in Proposition 2 that there can be random allocations in the two-type model is not an artifact of the discreteness or sparseness of the type space. We revert to the static setting with a single buyer as in Section 3.

Continuum model. We begin by introducing a model with a continuum of types, establishing that the seller can profit from random allocations. We then translate this result into implications for a discrete-types model which is a particular case of the model we introduced in Section 2. For the continuum model, we suppose that 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 with density f. We revert to the notation that buyer preferences are determined by $v(p; \theta)$. Recall that the type θ is the buyer's willingness to pay (i.e., $v(\theta; \theta) = 0$). We will denote by $v'(p; \theta)$ and $v''(p; \theta)$ the first and second derivatives with respect to p.

We assume that $v(p; \theta)$ has the same properties as in the model set-up, except that we now impose a continuous-type version of Condition A. In particular, we assume that, for any $p \geq 0$, $-v'(p; \theta)/v(p; \theta)$ is strictly decreasing in θ for $\theta > p$. Also, for any $p \geq 0$, $v''(p; \theta)/v'(p; \theta)$ is weakly decreasing in θ. We impose the additional condition that $v(\cdot; \cdot)$ is continuous.[14]

If the seller decides to use a posted-price mechanism with price p then the buyer purchases whenever his type θ exceeds the price. The seller's profit is therefore $p(1 - F(p))$. We assume F is such that this

[14]This additional continuity will be used to establish Proposition 6 below, where we translate our finding for the continuum model into a result for a setting with discrete types.

profit is uniquely maximized by an interior price $p^* \in (\underline{\theta}, \bar{\theta})$.[15] Denote the threshold type purchasing at this price by $\theta^* = p^*$.

Our first result establishes a condition under which the optimal deterministic mechanism, which is implementable using a posted price p^*, can be improved upon by a mechanism with random allocations.

Proposition 5. *Consider the continuum-type static model with a single buyer and a unique optimal posted price $\theta^* \in \left(\underline{\theta}, \bar{\theta}\right)$. Then there is a mechanism with a random allocation that is more profitable than the optimal deterministic mechanism provided that*

$$\frac{v''(\theta^*;\theta^*)}{v'(\theta^*;\theta^*)} > \frac{f'(\theta^*)\theta^* + 2f(\theta^*)}{1-F(\theta^*)}. \tag{4}$$

This result is established by a simple perturbation to the optimal deterministic mechanism. In particular, in addition to the options presented to the buyer to purchase the good with certainty and not at all, we introduce an option to purchase with an interior probability. Some types prefer this intermediate option and, when Condition (4) is satisfied, the perturbation increases seller profits.

To understand Condition (4), note that the left-hand side of the inequality is the buyer's coefficient of absolute risk aversion when his type is equal to the price in the optimal deterministic mechanism (θ^*), evaluated at a price equal to his type. It thus represents a measure of local risk aversion for the marginal type in the optimal deterministic mechanism. How large this must be for the result to apply depends on the distribution of types. Note, for instance, that if F is the uniform distribution on $\left[0, \bar{\theta}\right]$ for $\bar{\theta} > 0$, then it is enough that the coefficient of absolute risk aversion at $\theta^* = \bar{\theta}/2$ is greater than $4/\bar{\theta}$. For illustration, consider CARA preferences where $v(p;\theta) = 1 - e^{-R(\theta - p)}$ for a parameter $R > 0$ and consider varying the top of the support $\bar{\theta}$. Then the coefficient of absolute risk aversion (the left-hand side of (4)) remains constant at R, while the right-hand side of (4) decreases with $\bar{\theta}$.

Discrete approximation. Now let us explain how the finding in Proposition 5 sheds light on certain discrete-type models with many types. Fix the continuous distribution F on $[\underline{\theta}, \bar{\theta}]$. Then, for a discrete-type approximation, we can consider a sequence of models with $N_m = 2^m$ types, $m \in \mathbb{N}$. We let $\theta_n^m = \underline{\theta} + (n-1)\frac{\bar{\theta}-\underline{\theta}}{N_m}$ and let $\beta_n^m = F(\theta_{n+1}^m) - F(\theta_n^m)$ for $n = 1, \ldots, N_m$, where we set $\theta_{N_m+1}^m = \bar{\theta}$. Here, β_n^m is the probability of type θ_n^m in the discrete-type model. The discrete types can be viewed as partitioning the set $[\underline{\theta}, \bar{\theta}]$, and higher values of m correspond to finer partitions. For each m, we let E^m represent the environment with the specified types, payoffs, and distribution over types. For each environment E^m, we let the corresponding optimal mechanism be given by $(x_n^m, p_n^m)_{n=1}^{N_m}$. We then have the following characterization of optimal mechanisms when m is large so that the discrete-types model closely approximates the continuum.

Proposition 6. *Suppose that the hypothesis of Proposition 5 is satisfied; in particular, that the inequality (4) holds. Then there exists $\varepsilon > 0$ and K sufficiently large that, for all $m > K$, the following is true: There*

[15]The derivative of profit with respect to p is $-f(p)p + 1 - F(p)$, so a sufficient condition that implies our assumption is that (i) $-p + (1-F(p))/f(p)$ is strictly decreasing, while (ii) $\lim_{p \downarrow \underline{\theta}} f(p)^{-1} > \underline{\theta}$, and (iii) $\lim_{p \uparrow \bar{\theta}} f(p)^{-1}(1-F(p)) < \bar{\theta}$.

exist adjacent types $\theta^m_{n'}, \theta^m_{n'+1}, \ldots, \theta^m_{n''}$ *with* $x^m_{n'}, x^m_{n'+1}, \ldots, x^m_{n''} \in [\varepsilon, 1-\varepsilon]$ *and* $\theta^m_{n''} - \theta^m_{n'} \geq \varepsilon$.

The result states that there is always a positive value ε such that, as the number of types N_m becomes large, all types in an interval of length ε receive an allocation that is bounded (by ε) away from both zero and one. The probability of types that receive allocations with non-negligible randomness therefore does not vanish as the discrete-type model approximates the continuum setting. In this sense, our finding in Proposition 5 is informative also about certain settings with many discrete types.

Comparison to existing literature. Now let us relate the above findings to the literature, beginning with Liu and van Ryzin (2008). They consider a two-period model with CRRA utility and a continuum of types θ. For $\gamma \in (0,1)$ and for $p \leq \theta$, their specification of utility is $v(p;\theta) = (\theta - p)^\gamma$ (note that our assumptions do not accommodate their case since we require utilities to be defined at all payments). Their Proposition 9 implies that, irrespective of the value taken by γ in $(0,1)$, there is a mechanism with random allocations that outperforms the optimal deterministic mechanism.[16] To understand more their finding, recall that the optimal deterministic mechanism is implemented by posting a price p^*. The marginal type, who is indifferent to receiving the good at price p^*, is $\theta^* = p^*$. Under CRRA utility, we have $\frac{v''(p;\theta^*)}{v'(p;\theta^*)} \to \infty$ as p approaches θ^* from below, providing a sense in which the marginal type is locally infinitely risk averse. In view of our results, this infinite risk aversion would seem to underpin the finding of Liu and van Ryzin which implies that random allocations are profitable for all values of γ. Relative to them, Proposition 5 (and its discrete analogue, Proposition 6) establishes that arbitrarily high risk aversion is however *not needed* for the optimality of random allocations, and it gives an easily evaluated sufficient condition. Section 9 further discusses the relationship between our results and those in Liu and van Ryzin.

The results in Propositions 5 and 6 can also be compared to the findings in auctions models where the buyer may make payments if not receiving the good, as for instance in Matthews (1983) and Maskin and Riley (1984). Both papers provide conditions under which optimal mechanisms prescribe random allocations to an interval of types. Matthews studies a model with CARA preferences, where type θ has payoff $1-e^{-R(\theta-p)}$ when receiving the good for a payment p. Here, $R > 0$ is the coefficient of absolute risk aversion. The bidder instead has payoff $1-e^{Rp}$ if paying p and not receiving the good. Under regularity conditions, Matthews shows that there is a non-degenerate interval of types receiving random allocations no matter how small the coefficient of absolute risk aversion (see his Theorem 1).[17] We provide in the Online Appendix a contrasting result that also concerns CARA preferences, but assumes no payments are made when the buyer does not receive the good; i.e. the result is for our static model, restricted to

[16] Their result is that selling to consumers over two distinct dates (intertemporal price discrimination) is always more profitable than selling on only a single date. This result readily translates to the stated implication for static mechanisms.

[17] Matthews (1983) establishes this result in a setting where the seller has no capacity constraints, so the key difference compared to our static model is the possibility of payments when the buyer does not receive the good. In his optimal mechanism, each type that receives the good with positive probability makes a fixed payment that does not depend on whether he receives the good; the payment is lower for types that receive the good with a smaller probability. Note the finding that payments are independent of whether the bidder receives the good appears specific to the case of CARA preferences; see Moore (1984) for a related analysis.

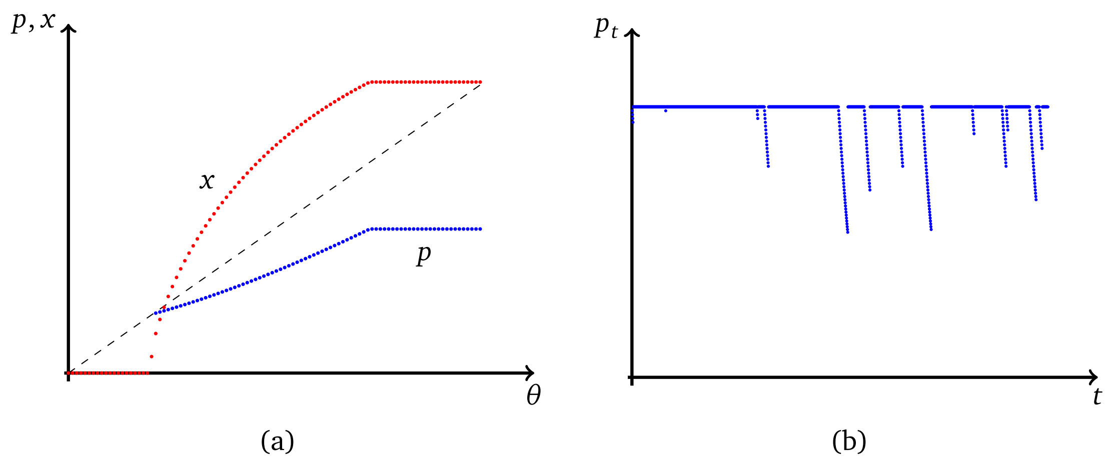

Figure 3: (a) Illustration the numerical optimal policies for a CARA utility function with $R = 1$. We assume that there are 100 types evenly spaced between 0 and 9.9, each with the same probability. In particular, $\theta_n = (n-1)/10$ and $\beta_n = 1/100$ for $n = 1, ..., 100$. (b) Simulation of the random price path under an near-optimal price process.

CARA preferences. Under our own mild regularity conditions, we show that the seller does better with a deterministic mechanism whenever the risk aversion coefficient is sufficiently small. This means that, while Proposition 5 establishes a sufficient condition on the level of risk aversion for optimal allocations to be random, we are able to show in the CARA case that at least a certain degree of risk aversion is also necessary.

An observation that helps to explain the difference between the finding of Matthews (1983) and our own is as follows. The optimal deterministic mechanism is the same whether or not the buyer can make payments when not receiving the good. In both settings, this deterministic mechanism can be implemented as a posted price (the buyer pays the price if and only if he gets the good). The range of random mechanisms available to the seller in the environment of Matthews is broader, however, as the seller can choose mechanisms that involve payments for not receiving the good. Moreover, the seller has access to any of the mechanisms in our environment where payments are instead restricted. For this reason, random mechanisms are favored more in the environment of Matthews than with the restriction to winner pays.

Price process with many types. Finally, in light of the results and discussion in this section, it is of interest to understand how our near-optimal price process behaves when there are many discrete types. Here, we will consider the same price process described after Proposition 3 when there are many discrete types. When discrete types approximate a continuum model (as described earlier in the section), the price process defined with respect to the discrete-type approximation is also approximately optimal in the continuum model.

Figure 3(a) depicts the numerical solution to the static problem. Like the other numerical exercises we performed, the solution to the static problem features a lower pool of types that get the good with probability zero, a region where the allocation probability increases from zero to one, and an upper pool of types which get the good allocated for sure. Figure 3(b) reproduces a realization of the price path of a near-optimal price mechanism derived from the optimal static mechanism as described in Section 4. As can be seen, there are different events where a discount episode is initiated, and such episodes end at random prices. Each time one such episode starts the price goes down fast but in small increments. In the simulation, there are eight such episodes of different lengths. The discounts range from 5% to 30% of the regular price.

7 Loss aversion

Aversion to risk with small stakes. A possible concern for our theory is that discounts in many of the relevant markets involve small stakes. It has been argued that significant risk aversion over small stakes would imply implausible risk preferences over larger gambles. For instance, Rabin (2000) shows how non-negligible risk aversion with respect to small-stakes gambles considered at a range of wealth levels implies implausible levels of risk aversion for modestly larger gambles, at least when one asserts expected utility preferences over final wealth and a single utility function. Nonetheless, risk aversion is often exhibited in small-stakes settings. Evidence from experimental settings includes Read et al. (1999), Fehr and Goette (2007) and Gächter et al. (2022). For instance, Read et al. find that 68% of their sample of students reject a gamble with a 50-50 chance to win $40 or lose $25. Fehr and Goette find that 27 of their sample of 42 bike messengers reject a gamble with a 50-50 chance to win 8 or lose 5 Swiss francs. Similar observations have been made in the field; e.g. Cicchetti and Dubin (1994), where telephony consumers were found to pay 45 cents per month to insure against an expected loss of 26 cents per month (roughly, consumers were insuring against a 1/200 chance each month of losing $55). Such examples are often seen as demonstrating "loss aversion", the idea that losses loom larger to decision makers than equally-sized gains (see Kahneman and Tversky, 1979).

Given these observations, we see two main arguments in favor of our theory. The first is that an expected utility model with risk aversion may capture well the trade-offs faced by consumers and therefore the problem of the seller, irrespective of Rabin's (2000) critique. This argument is similar to those made by Rubinstein (2002) and Cox and Sadiraj (2006), who do not see Rabin's arguments as a reason to dispense with expected utility models with risk-averse agents, even for settings with smaller stakes. For instance, Rabin shows how the expected utility model yields absurd conclusions by applying a unified view of individual decision-making in light of a sequence of choices that an individual would be expected to take. But individual behavior often in fact exhibits "narrow bracketing" (see, e.g., Read et al., 1999) and fails to take the implications of collections of decisions into account. Thus, although the expected utility model with risk-averse individuals leads to inconsistencies, this does not rule out the possibility that it does a good job of accounting for actual decision-making (indeed individual decision-making can often appear inconsistent). Finally, there can be reasons to favor expected utility models, including tractability and comparability to existing theories (e.g., to the theory of risk aversion in auctions discussed in the previous section).

Our second main argument is that the key predictions of our theory, and the key forces described above, can survive also for formal models of consumer loss aversion. To see this, we study a model of loss aversion that adapts ideas in work such as Kahneman and Tversky (1979) and Kőszegi and Rabin (2006). While there are various formulations of loss-averse preferences, we deliberately make choices to yield a smaller departure from the model with risk aversion already examined, and at the same time to reflect our own intuitions about how loss aversion could be relevant for consumers.

Model of loss aversion. Our starting point is to posit gain-loss utility for each type of consumer $\theta_n > 0$ that depends on a type-specific reference point $\rho_n > 0$. For simplicity, we focus on the case of two consumer types: the low type θ_1 and the high type θ_2, with $\theta_1 < \theta_2$. The probability of each type θ_n is β_n. We specify the payoff of type θ_n from purchasing at price $p \in \mathbb{R}_+$ to be

$$v\left(p;\theta_n,\rho_n\right) \equiv \theta_n - p + \mu(\rho_n - p),$$

where

$$\mu(x) = \begin{cases} \lambda\eta x & \text{if } x \leq 0 \\ \eta x & \text{if } x > 0 \end{cases}$$

and where $\eta > 0$ and $\lambda > 1$. We specify the payoff in case not purchasing to be zero. Similar to our theory with risk-averse preferences, we view $v\left(p;\theta_n,\rho_n\right)$ as the instantaneous utility of a purchase at price p, so that intertemporal payoffs from a purchase after delay t are given by $e^{-rt}v\left(p;\theta_n,\rho_n\right)$ where $r > 0$ is the discount rate.

To understand this formulation, note here that $\mu(\rho_n - p)$ is intended to capture the "gain-loss utility": a gain from purchasing in case the price p is below the reference point ρ_n for type θ_n, or a loss from purchasing in case p is above ρ_n. That losses loom larger than gains is captured by the assumption that $\lambda > 1$.[18] The parameter η can be thought of as the magnitude or importance of loss aversion. Note that the type θ_n now represents the willingness to pay in the absence of gain-loss utility. Our interpretation of this specification is that, when a consumer considers buying the good, he has a reference price ρ_n in mind, which could reflect the price he thinks he would "usually" pay. An increase in the price above this level influences more how the consumer views the purchase than an equally-sized reduction in price below the reference point.

Several comments are worth making at this juncture. First, the representative consumer's reference point is taken to be type dependent. This follows Carbajal and Ely (2016) who consider type-dependent reference points for quality in a mechanism design problem (see also Spiegler, 2012, where each consumer is viewed as having a different reference point). As Carbajal and Ely explain, the approach views consumers as having different reference points in different "states" (here states with high or low willing-

[18] The assumption that losses are experienced more intensely than gains was initially suggested by Kahneman and Tversky (1979). The assumption of linearity of payoffs above and below the reference point has been common in the subsequent literature on loss aversion.

ness to pay), and is thus comparable to the state-dependent model of Sugden (2003).[19]

A second comment is that we specify loss aversion over only one dimension of the consumption experience, the price. We do not specify loss aversion over whether buyers receive the good. This is different from a number of models where loss aversion over the consumption quantity can manifest in an "attachment" to the good (see, e.g., Kőszegi and Rabin, 2006, and Heidhues and Kőszegi, 2014). However, we are not alone in specifying loss aversion only over the price dimension (see, e.g., Spiegler, 2012, and Herweg and Mierendorff, 2013).

A third and important comment is that our specification of preferences is non-standard in that gain-loss utility is only specified in the event the buyer purchases. In a static context, this amounts to specifying a zero payoff in case the buyer does not receive the good and thus ignoring any gain-loss utility. In a dynamic setting, it amounts to reducing a buyer's concern to the discounted payoff from purchasing, namely $e^{-rt}v(p;\theta_n,\rho_n)$. This abstracts from the flow "gains" that a buyer might enjoy from not having to pay over periods when he is still yet to purchase. Our assumption not only simplifies buyers' intertemporal decisions, but it makes our specification and arguments easily comparable to the baseline model with risk-averse buyers. It also strikes us as plausible that consumers would only experience gains and losses relative to a reference price when actually purchasing the good.

Fourth, note that we will focus on deterministic reference points that for now we take as exogenous. However, each type of consumer could also entertain random reference points, evaluating preferences in light of their type but also in light of a distribution of reference points, integrating over them as in Kőszegi and Rabin (2006). Suppose that the distribution over reference points for each type θ_n is G_n, a continously differentiable function with density g_n. Then the expected instantaneous utility of a type θ_n from purchasing at price p is

$$\begin{aligned}\int_0^\infty v(p;\theta_n,\rho)\,g_n(\rho)\,d\rho = &\int_p^\infty (\theta_n - p + \eta(\rho - p))\,g_n(\rho)\,d\rho \\ &+ \int_0^p (\theta_n - p + \lambda\eta(\rho - p))\,g_n(\rho)\,d\rho.\end{aligned}$$

The derivative with respect to p is

$$-1 - \eta(1 - G_n(p)) - \eta\lambda G_n(p)$$

which is differentiable and decreasing in p. In other words, the expected instantaneous utility is concave in p and so the buyer's preferences are the same as those of a risk-averse buyer as specified in our baseline

[19]De Giorgi and Post (2011, p 1094) note that, in Sugden (2003), "the decision maker compares the prospect and the reference point only in the same state and not across states, and experiences loss ... only if the outcome of the prospect falls below the outcome of the reference point in the same state". De Giorgi and Post contrast this to the approach in Kőszegi and Rabin (2006), which does not make agents' reference points state dependent. De Giorgi and Post suggest that the state-dependent model has possible advantages in some contexts.

model. Additional conditions on the functions G_n can be made to ensure the properties A1 and A2 of the baseline model are also satisfied. In this sense, our baseline model already accommodates a possible specification with consumer loss aversion.

We now consider the seller's problem when reference points are fixed to be deterministic scalars $\rho_1 = \theta_1$ for type θ_1 and $\rho_2 \in (\theta_1, \theta_2)$ for type θ_2. For now, these reference points are exogenous, but we will subsequently justify these choices as reflecting consumer expectations about market prices. As before, we begin by studying the static mechanism, say for a single buyer. We find an optimal mechanism with random allocation to type θ_1 whenever parameters satisfy

$$\beta_2 \in \left[\frac{\theta_1 (1+\eta)}{\theta_2 + \eta \rho_2}, \frac{\theta_1 (1+\eta+\eta(\lambda-1))}{\theta_2 + \eta\rho_2 + \theta_1 \eta (\lambda - 1)} \right] \tag{5}$$

(an interval of positive length). In particular, we establish the following result.

Proposition 7. *Consider the model of loss aversion with two types and reference points θ_1 for the low type and $\rho_2 \in (\theta_1, \theta_2)$ for the high type. Then if Condition (5) is satisfied, there is an optimal static mechanism in which the low type has a random allocation while the high type receives the good for sure. The low type pays θ_1 when receiving the good, while the high type pays his reference point ρ_2.*

To understand the form of the optimal mechanism described in Proposition 7, suppose that Condition (5) holds and the seller offers an optimal direct mechanism as described in Proposition 7. Let the probability of allocation to the low type in this mechanism be $\bar{x}_1$ (with value derived in the appendix). We want to explain why $\bar{x}_1 \in (0, 1)$. Note that, given this allocation probability to the low type, the high type is just indifferent to mimicking the low type when charged a price equal to his reference point ρ_2 to receive the good for sure. Consider reducing the low type allocation probability x_1 below $\bar{x}_1$ while increasing the price to the high type to maintain indifference to mimicry. Because the price increase is above the high type's reference point, the high type experiences this as a loss. This means that the high type is particularly price sensitive, and so the price increase permitted by the reduction in x_1 is relatively small. Similar logic applies to increases in x_1 above $\bar{x}_1$. There, maintaining the high type's indifference to mimicking the low type calls for a reduction in the high type's payment, which the high type experiences as a gain. Because the high type is less sensitive to gains, the price reduction needed for the high type not to mimic the low type is relatively large. This explains why, for some parameter values (those indicated in Condition (5)), there is an interior "sweet spot" allocation probability $\bar{x}_1$ that maximizes seller profits. This is the allocation probability such that the high type, made indifferent to mimicry, experiences neither gains nor losses.

Now consider a dynamic environment where buyers arrive over time, say at a constant rate r. Proportion β_2 of arrivals are high types, with the remainder low types. Reference points are fixed at ρ_2 and θ_1 respectively, and also therefore do not depend on the arrival time to the market. Suppose that parameters satisfy the condition in Equation (5). Then the same arguments that were made for the baseline model can be applied to conclude the following. There is an optimal random price path in which the seller sets a constant price ρ_2, while there is a Poisson arrival process for instantaneous discounts to price θ_1 that target the low type. The Poisson arrival rate for discounts is $\frac{r\bar{x}_1}{1-\bar{x}_1}$. This is the rate that ensures that a buyer arriving in the market at an arbitrary (non-discount) date anticipates that the expected discounting associated with the date of the next price reduction is $\bar{x}_1$. Given this price process, high types purchase immediately on arrival at price ρ_2, while low types purchase at discounts at price θ_1. The seller earns the same profits per buyer as in the static mechanism, which as before represents an upper bound on the profits available to the seller in the dynamic format.

Reference point determination. For the parameter values where the condition in Equation (5) is satisfied, we have seen that there is an optimal mechanism in the static format and optimal price process in the dynamic format in which both types of buyer make payments for purchase equal to their reference points. This seems to suggest the possibility that reference points could be formed through rational expectations about play. In fact we suggest below that this can be formalized in a similar way to Heidhues and Kőszegi's (2008) concept of "market equilibrium" whereby consumer reference points are determined by lagged rational expectations about play.

To follow Heidhues and Kőszegi's description of market equilibrium as closely as possible, we first adapt the idea of "personal equilibrium" (as introduced by Kőszegi and Rabin, 2006) to our setting. Consider first a static environment. To see that there exist equilibria with reference points as introduced above, it will be enough to consider (static) *direct* mechanisms (x_1, H_1, x_2, H_2), where x_n is type θ_n's allocation probability and H_n is the price distribution specified for this type.[20] A "reporting strategy" can be defined as a probability distribution over reported types for each type θ_n; this would, say, assign probability σ_n to a report of his "own" type θ_n and complementary probability to a report of the other type. Then a "personal equilibrium" for a buyer, given mechanism (x_1, H_1, x_2, H_2), is a reporting strategy (σ_1, σ_2) where, for each type θ_n, the buyer's participating and reporting according to σ_n is optimal given a reference point that is the distribution of payments this type expects given the mechanism and his reporting strategy, *conditional on acquisition of the good*. We leave unspecified the formation of reference points in case the probability a given type receives the good is zero, but this will not present a difficulty as we focus on equilibria where the probabilities of allocation are strictly positive. Our requirement that reference points be determined only relative to payments made when acquiring the good is different to Kőszegi and Rabin (2006), but this seems to fit well with our specification that the buyer only takes reference points into account when receiving the good.

Next, consider how our adapted version of "personal equilibrium" applies to optimal direct mechanisms as characterized above. In particular, consider any reference points θ_1 and $\rho_2 \in (\theta_1, \theta_2)$, suppose the condition in Equation (5) is satisfied, and fix an optimal static mechanism as in Proposition 7. Truthful reporting is then a personal equilibrium. Given truthful reporting, low types pay θ_1 when receiving the good, and high types pay ρ_2. Also, truthful reporting is optimal taking these payments to be the reference points of each type. Note that while we have so far introduced the personal equilibrium concept

[20]That we can focus on direct mechanisms will follow because the seller will never have an incentive to deviate from the use of direct mechanisms.

for static mechanisms, the same concept can be applied to the dynamic format where buyers arrive over time. When a buyer arrives to the market, he has reference points determined by his purchase strategy (the mapping from type and price histories to purchase decisions) and given the expected evolution of prices.[21] Given the random price path, the purchase strategy is a personal equilibrium if it is optimal for the buyer given these reference points. In particular, for the price process described above, where there is a regular price of ρ_2 and stochastic discounts to θ_1, there is a personal equilibrium in which high types purchase immediately at ρ_2 and low types purchase at discounts at price θ_1. The reference points of the buyer are therefore ρ_2 for the high type and θ_1 for the low type.

Now consider our version of *market* equilibrium in the static mechanism design environment. This is a mechanism $M = (x_1, H_1, x_2, H_2)$ and a reporting strategy (σ_1, σ_2) such that (i) (σ_1, σ_2) is a personal equilibrium for the buyer given M, and (ii) M, together with reporting strategy (σ_1, σ_2), is an optimal mechanism for the seller given the reference points induced by $(M, (\sigma_1, \sigma_2))$. Let us reiterate that, under this concept, reference points should be viewed as buyers' lagged rational expectations about future play. That is, reference points are determined before the mechanism is offered, and the mechanism is chosen optimally taking these reference points as given. Again, consider any reference points θ_1 and $\rho_2 \in (\theta_1, \theta_2)$, and suppose the condition in Equation (5) is satisfied. Then the optimal mechanism we described in Proposition 7 and a truthful reporting strategy constitutes a market equilibrium. In particular, given the mechanism characterized above, and given truthful reporting, high types purchase with a payment ρ_2 and low types with a payment θ_1. These payments are exactly the reference points upon which the design of the optimal mechanism was predicated.

The "market equilibrium" concept introduced here can be extended straightforwardly to the dynamic format with dynamic arrivals. The interpretation is that buyers form their reference points before arrival to the market based on the price process they anticipate, and based on their anticipated response to the evolution of prices.[22] The seller's choice of price process is optimal *given* these reference points. It again consists of a "regular price" set equal to the reference point of the high type, at which high types purchase immediately, and a discounted price to the low type equal to the low type valuation θ_1, with these discounts arriving at the same Poisson rate determined above.

An interesting question that our approach is silent on is what predictions could be made if the seller's prices could influence the buyer's reference points over time. Also, it is important to point out that our very partial analysis of market equilibrium only indicates that reference points could be sustained in an equilibrium with rational expectations, but does not seek to characterize the entire equilibrium set. It is already clear from the above that a range of market equilibria exist corresponding to different reference

[21] We specified that reference points are determined by the distribution of payments made by each type conditional on purchasing. Note that this is also well-defined in the dynamic setting, but that it assigns equal weights to purchases in the distant and near futures. Alternative specifications of the reference point would be possible while still delivering our main observations. For instance, the probabilities of making different payments could be weighted also according to the *time* the buyer expects them to occur (for instance, weighting according to the discount factor may make sense).

[22] Heidhues and Kőszegi (2008, footnote 9) write "While the expectations that are relevant for specifying the reference point are clearly lagged, the fact that we do not specify when exactly these expectations are formed is a weakness of our approach." Our application of the market equilibrium idea is subject to the same criticism.

points for the high type. Note that multiplicity of market equilibria is observed also by Heidhues and Kőszegi (2008). They write that, in their setting, "there is typically a continuum of focal prices, so firms have a strong incentive to manage consumers' price expectations" (footnote 9). Similarly in our setting it may be interesting to view firms as influencing expectations so as to pick among market equilibria, but we leave the possible implications for future work.

8 Competition

Some of the markets for which our theory appears relevant are characterized by the presence of more than one seller. While we predict random price discounting in our model with a single seller, the basic forces that we studied appear pertinent also in markets with more sellers. Below we describe two situations where this is possible; one with collusion among firms and one with frictions in competition. We argue that such forces can be strong enough to rule out the undercutting argument in models such as Varian (1980) (see Section 9 for a further discussion).

Frictionless competition and collusive equilibria. What formal models of competition predict seems heavily dependent on the details of the environment, including what kinds of commitments sellers can make. To begin, suppose that buyers arrive over as described in our model set-up, but that they are free to purchase from any of several identical firms without restriction. Also, firms can commit to future prices. One equilibrium involves the "Bertrand outcome" in which the firms commit to charge a price equal to marginal cost (in our set-up, zero) forever irrespective of what other firms do. The reason is simply that, fixing this strategy for competitors, a given firm cannot earn positive profit, and hence is willing to set price at marginal cost. However, especially if one permits commitments to prices that are contingent on the pricing decisions of opposing firms, one anticipates many other equilibria as well. We anticipate behavior reminiscent of the "folk theorem" results that arise in "contractible contracts" problems, discussed for instance in Peters and Szentes (2012). For instance, firms can commit to high prices, but also commit to cut prices at some point if a competitor deviates. There may also be equilibria where firms commit to random price paths that maximize joint industry profits, achieving the same profits as a hypothetical monopolist in the relevant market. Punishments could ensue if either firm deviates from this prescription; this is especially true if firms' commitments can be made contingent on the pricing *commitments* of other firms, rather than merely the realized prices, since then even the slightest deviation in a firm's pricing policy could be directly detected and punished.[23] Profit sharing could be achieved with each firm taking it in turn to drop prices, selling to lower types who remain in the market. In these kinds of equilibria, random discounts would be sustained as a sort of collusive outcome among many sellers.

Switching/search costs and the Diamond Paradox. A different direction is to consider imperfections in competition that follow from restrictions on where a consumer can shop, as well as on consumers' information about prices. For instance, suppose there are two identical firms, each specified as the incumbent

[23] More generally, equilibria could be sustained with the help of a public randomization device.

seller for an equal split of buyers. That is, suppose half of arriving buyers are randomly assigned to one firm, and half to the other. Firms commit simultaneously at the outset to random price paths, and as in our baseline model cannot reneg on these commitments. Buyers can only purchase from their designated seller, unless paying a switching cost to purchase from the other.

We make some additional assumptions. First, buyers only observe the pricing commitment of their originally designated seller and must pay the switching cost to learn both the commitment and prices of the other firm. Also, buyers do not observe others' purchasing decisions and so cannot learn from realized demand. In addition, upon paying the cost, a buyer can no longer purchase from the firm to which he was originally assigned.[24] These assumptions could plausibly describe a market where consumers only have a limited capacity for attention, so can only pay attention to the prices of one seller at a time, incurring a cost to switch.[25] Finally, we assume that an optimal price process for the monopolist exists (for instance, consider our two-type model).

In the above setting, similar to Diamond's (1971) paradox, there is an equilibrium in which both firms commit to the same optimal random price process as a monopolist and no consumer switches. To see this, consider the best response of a firm whose competitor chooses the monopoly price process, and where the firm's customers correctly conjecture the competitor's process. Assuming none of the firm's assigned customers will switch away, the firm maximizes profits from these customers by committing to the monopoly price process. None of the assigned customers will in fact switch to the competitor, because they believe the competitor chooses the same price process. In addition, a deviation from monopoly pricing does not lead the firm to attract customers from its competitor, as customers assigned to the competitor are assumed unable to learn about the deviation unless paying the switching cost.[26]

The above claims hold even if switching costs are small (or zero), and the arguments extend to more than two firms. This suggests how our characterization of the monopoly price path might remain directly applicable in markets with several firms. In spite of the simplicity of the above arguments, a more difficult question that we leave open is the conditions under which the equilibrium outcomes described represent the *only* equilibrium outcomes of the model.[27]

[24]A richer model would permit buyers to switch multiple times; such a model could be chosen so as not to alter our main conclusions.

[25]We believe that switching costs of the kind described here could well be the reality in many markets. A customer may have a store that he frequents and thinks about purchasing from, and while being aware of the existence of competitors, finds there are cognitive costs of understanding their price offers. Relatedly, consumers are likely to have capacity constraints in terms of the attention they can dedicate to following the prices at different stores, perhaps leading them to focus on the prices offered by only one.

[26]To complete the characterization of equilibrium, it is necessary to specify consumer beliefs also in case they observe a deviation from the proposed equilibrium strategy of their assigned firm. Here any specification of beliefs will do, since each firm will not want to deviate to another price process *whether or not* this new process induces switching away by its assigned consumers (customers switching away to the other firm would only lower profits). We could, for instance, follow much of the literature on competing contracts and specify "passive beliefs", meaning that customers continue to have the same beliefs about the other firm's prices irrespective of the price process chosen by their assigned firm.

[27]For instance, note that Diamond's (1971) original argument for uniqueness relied on the quasi-concavity of profits in prices. Given that random random price paths are complicated objects, we have not established analogous properties for the dependence of seller profits on such price paths. A related consideration is the admissible strategy space for firms; uniqueness would seem to require at the least that firms not be able to condition prices on those of competitors.

9 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), Varian (1980) and Rosenthal (1980). Such theories assume that not all consumers in the market can access all competing price offers on the same terms. In the celebrated work of Varian (1980), for instance, some customers are informed of all offers in the market, and others are informed of the offer of only one seller. The theories of randomized pricing based on mixed strategies under competition are generally developed in a static framework. The models can be most simply extended to dynamic settings by considering non-durable goods sold in every period with sellers redrawing prices in each period.[28] Such models would then predict mixed strategies independent of historical prices. For instance, Varian (1980) sees this idea as integral to his theory. Extending the implications of his static model to dynamic environments, he writes (p. 651) "because of intentional fluctuations in price, consumers cannot learn by experience about stores that consistently have low prices". This emphasizes the idea of independence of prices from historical ones. The more often prices are drawn, the more variation in price the repetition of the static model would predict.

Theories of price discounting based on mixed strategies and competition therefore seem to have difficulty accounting for the considerable temporal price stability that is observed in many markets: consider the various papers identifying persistent "regular" prices punctuated by occasional price discounts (see the first paragraph of our Introduction). In fact, the observed temporal persistence of empirically observed prices motivates Myatt and Ronayne (2019) to provide a dynamic model of price setting where firms first set list prices and then subsequently may only charge prices that are no higher than the list prices. In their equilibrium it turns out that firms choose prices equal to the pre-specified list prices, so there is temporal price stability. However, firms do not discount their prices relative to the list price in equilibrium, so they do not explain price discounting in equilibrium. Rather, the objective is to account for price dispersion across firms that is persistent over time in the sense that firms which are initially the highest priced remain so. While our monopoly model speaks little to the question of price dispersion across firms, our theory captures well the idea of a persistent "regular" price punctuated by occasional stochastic discounts.

Part of the difficulty in reconciling models of competitive mixed-strategy pricing with empirical patterns also lies in the fact that the price distributions predicted are often atomless. An empirical prediction that we have emphasized is the idea that there is a price most commonly offered termed the "regular price". The idea of a more common regular price under mixing by the seller is, however, captured in the work of Heidhues and Kőszegi (2014). This paper specifies a static model where consumers are loss averse in both the payment dimension and the quantity dimension. Loss aversion on the quantity dimen-

[28]Note that Fershtman and Fishman (1992) do consider an explicitly dynamic model of durable goods pricing where firms re-choose prices in every period. Consumers cannot recall prices and firms are viewed as choosing prices independently in every period.

sion implies an "attachment effect" is created from the anticipation of price discounts which will surely induce a buyer to purchase (we mentioned this effect in Section 7). The theory then predicts an atomistic "regular" price and continuously distributed "sale" prices, with a gap between the regular and sale prices. As the authors point out, the gap between discounted and regular prices has an empirical counterpart as observed for instance in data on supermarket pricing. Compared to the work of Heidhues and Kőszegi, our theory is instead explicitly dynamic. An advantage of an explicitly dynamic theory is that it permits predictions on the timing of price discounts. Note that the predictions of Heidhues and Kőszegi could be taken to a dynamic setting by supposing that the seller draws prices from the characterized distribution repeatedly over time. How often sales prices are drawn would, however, depend on the frequency with which the seller draws new prices from the distribution. So the temporal frequency of price discounts appears indeterminate (i.e., it is not pinned down how many sales would occur on average in a given year). Instead we are able to provide predictions on the rate of price discounting over time that do not depend on arbitrary considerations such as the frequency of price setting.

Another difference to Heidhues and Kőszegi (2014) is our key finding that each consumer type is (nearly) perfectly insured with respect to the price paid for acquisition. This contrasts with Heidhues and Kőszegi where consumers are exposed to variation in the price paid. In a static setting, fully insuring buyers against pricing risk in our setting often requires the seller to offer a revelation mechanism. In a dynamic environment when buyers arrive over time, insurance against price risk is achieved through a regular price with random discounting episodes. An important part of our objective is to contribute to the literature on durable-goods pricing where buyers are forward-looking. As emphasized in the Introduction, the finding of sudden and random price discounts in the dynamic model contrasts with the gradual declines in prices that have often been characterized. Note that we do not, however, provide a theoretical underpinning for the presence of a gap between regular and discounted prices. This is indeed a feature of our model with a small number of discrete types, but we did not find a gap for instance in examples like that at the end of Section 6 with many types.[29]

A leading set of theories of price discounting that can be viewed as alternatives to the mixed-strategy theories mentioned above 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 and Meilijson (1985).[30] 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

[29]Our model can produce a "gap" between regular and discount prices when there are many types if there is bunching in the allocation probability to different types. While this seems a potential outcome of our model, we have not investigated the possibility in detail.

[30]Related, it could also be that buyers with high valuations are myopic as in Pesendorfer (2002) (see also Sobel, 1984, for a model of sales with myopic consumers, though in a setting with competition). 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).

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 et al. (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).[31]

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.

It is worth pointing out that there are other papers on dynamic pricing to risk-averse buyers, in particular Liu and van Ryzin (2008) and Bansal and Maglaras (2009). Liu and van Ryzin show that price discrimination can be optimal in settings with risk-averse consumers, implying the optimality of setting different prices to different types as we explained in Section 6. Moreover, for the same preferences as in Liu and van Ryzin, Bansal and Maglaras provide certain results on static mechanisms in which the buyer only pays if he gets the good, relating the findings to a dynamic pricing problem. However, note that our full characterization of optimal static mechanisms and the resulting optimal stochastic price process is new. Most importantly, while Bansal and Maglaras do capture the relevance of the static mechanism design problem, they *assume* that each customer type pays a deterministic price in equilibrium rather than deriving this result. Thus, crucially, there is no analogue of our Lemma 2 that shows the desirability of fully insuring customers against pricing risk. This result underlies the nature of (virtually) optimal dynamic pricing in our paper. In addition, note that Liu and van Ryzin and Bansal and Maglaras do not consider dynamic arrivals, and therefore do not uncover the important role of random price discounts that we explained here.

Finally, it is worth pointing out that there are many other papers where prices evolve randomly over time, but where this randomness is a response to exogenous uncertainty that is realized over time. Examples of these papers include Hörner and Samuelson (2011), Board and Skrzypacz (2016), Gershkov et al. (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 infinitesimal 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.

[31] 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.

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, 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, 1257–1273.

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

BOARD, S. AND A. SKRZYPACZ (2016): "Revenue management with forward-looking buyers," *Journal of Political Economy*, 124, 1046–1087.

BUDISH, E. B. AND L. N. TAKEYAMA (2001): "Buy prices in online auctions: irrationality on the internet?" *Economics Letters*, 72, 325–333.

CARBAJAL, J. C. AND J. C. ELY (2016): "A model of price discrimination under loss aversion and state-contingent reference points," *Theoretical Economics*, 11, 455–485.

CHEN, T., A. KALRA, AND B. SUN (2009): "Why do consumers buy extended service contracts?" *Journal of Consumer Research*, 36, 611–623.

CHEVALIER, J. AND A. GOOLSBEE (2009): "Are durable goods consumers forward-looking? Evidence from college textbooks," *The Quarterly Journal of Economics*, 124, 1853–1884.

CHEVALIER, J. A. AND A. K. KASHYAP (2017): "Best Prices: Price Discrimination and Consumer Substitution," Tech. rep.

CICCHETTI, C. J. AND J. A. DUBIN (1994): "A microeconometric analysis of risk aversion and the decision to self-insure," *Journal of Political Economy*, 102, 169–186.

CONLISK, J., E. GERSTNER, AND J. SOBEL (1984): "Cyclic pricing by a durable goods monopolist," *The Quarterly Journal of Economics*, 99, 489–505.

CORREA, J., J. F. ESCOBAR, AND A. PERLROTH (2019): "Revenue maximization with heterogeneous discounting: auctions and pricing," Tech. rep.

COX, J. C. AND V. SADIRAJ (2006): "Small-and large-stakes risk aversion: Implications of concavity calibration for decision theory," *Games and Economic Behavior*, 56, 45–60.

DE GIORGI, E. G. AND T. POST (2011): "Loss aversion with a state-dependent reference point," *Management Science*, 57, 1094–1110.

DIAMOND, P. A. (1971): "A model of price adjustment," *Journal of Economic Theory*, 3, 156–168.

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, 234–62.

FEHR, E. AND L. GOETTE (2007): "Do workers work more if wages are high? Evidence from a randomized field experiment," *American Economic Review*, 97, 298–317.

FERSHTMAN, C. AND A. FISHMAN (1992): "Price cycles and booms: dynamic search equilibrium," *American Economic Review*, 82, 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.

Gächter, S., E. J. Johnson, and A. Herrmann (2022): “Individual-level loss aversion in riskless and risky choices,” *Theory and Decision*, 92, 599–624.

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

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

Gowrisankaran, G. and M. Rysman (2012): “Dynamics of consumer demand for new durable goods,” *Journal of Political Economy*, 120, 1173–1219.

Heidhues, P. and B. Kőszegi (2008): “Competition and price variation when consumers are loss averse,” *American Economic Review*, 98, 1245–68.

——— (2014): “Regular prices and sales,” *Theoretical Economics*, 9, 217–251.

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

Herweg, F. and K. Mierendorff (2013): “Uncertain demand, consumer loss aversion, and flat-rate tariffs,” *Journal of the European Economic Association*, 11, 399–432.

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

Kahneman, D. and A. Tversky (1979): “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica: Journal of the Econometric Society*, 47, 263–291.

Kehoe, P. and V. Midrigan (2015): “Prices are sticky after all,” *Journal of Monetary Economics*, 75, 35–53.

Kőszegi, B. and M. Rabin (2006): “A model of reference-dependent preferences,” *The Quarterly Journal of Economics*, 121, 1133–1165.

Lan, H., T. Lloyd, W. Morgan, and P. W. Dobson (2022): “Are food price promotions predictable? The hazard function of supermarket discounts,” *Journal of Agricultural Economics*, 73, 64–85.

Landsberger, M. and I. Meilijson (1985): “Intertemporal price discrimination and sales strategy under incomplete information,” *The RAND Journal of Economics*, 16, 424–430.

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

Maskin, E. and J. Riley (1984): “Optimal auctions with risk averse buyers,” *Econometrica: Journal of the Econometric Society*, 52, 1473–1518.

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

Moore, J. (1984): “Global incentive constraints in auction design,” *Econometrica: Journal of the Econometric Society*, 52, 1523–1535.

Myatt, D. P. and D. Ronayne (2019): “A Theory of Stable Price Dispersion,” Tech. rep.

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

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

Öry, A. (2017): “Consumers on a Leash: Advertised Sales and Intertemporal Price Discrimination,” Tech. rep.

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

PETERS, M. AND B. SZENTES (2012): "Definable and contractible contracts," *Econometrica*, 80, 363–411.

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: Cambridge University Press, 202–208.

READ, D., G. LOEWENSTEIN, AND M. RABIN (1999): "Choice Bracketing," *Journal of Risk and Uncertainty*, 19, 171–197.

REYNOLDS, S. S. AND J. WOODERS (2009): "Auctions with a buy price," *Economic Theory*, 38, 9–39.

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

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

ROSENTHAL, R. W. (1980): "A model in which an increase in the number of sellers leads to a higher price," *Econometrica: Journal of the Econometric Society*, 1575–1579.

RUBINSTEIN, A. (2002): "Comments on the risk and time preferences in economics," *Working paper, Tel-Aviv University, Foerder Institute for Economic Research, Sackler Institute for Economic Studies*.

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

SOBEL, J. (1984): "The timing of sales," *The Review of Economic Studies*, 51, 353–368.

——— (1991): "Durable goods monopoly with entry of new consumers," *Econometrica: Journal of the Econometric Society*, 59, 1455–1485.

SPIEGLER, R. (2012): "Monopoly pricing when consumers are antagonized by unexpected price increases: a "cover version" of the Heidhues–Kőszegi–Rabin model," *Economic Theory*, 51, 695–711.

STOKEY, N. L. (1979): "Intertemporal price discrimination," *The Quarterly Journal of Economics*, 93, 355–371.

SUGDEN, R. (2003): "Reference-dependent subjective expected utility," *Journal of Economic Theory*, 111, 172–191.

VARIAN, H. R. (1980): "A model of sales," *American Economic Review*, 70, 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, 321–352.

Appendix I: Proofs of the formal results

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,

$$\begin{aligned} x''v_l(p'') &\geq x'v_l(p')\frac{v_k(p')}{v_k(p'')}\frac{v_l(p'')}{v_l(p')} \\ &= x'v_l(p')e^{\int_{p'}^{p''}\left(\left(-\frac{v_k'(p)}{v_k(p)}\right)-\left(-\frac{v_l'(p)}{v_l(p)}\right)\right)dp} \\ &> x'v_l(p'), \end{aligned} \tag{6}$$

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 $\mathscr{J}$ of cardinality $J \equiv |\mathscr{J}|$ 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 $\mathscr{J} = \{\theta_{n_j} : j = 1, ..., J\}$. 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 $\mathscr{J}$.

Consider now the reporting decision of any type θ_{n_j} in M^R, with $\theta_{n_j} \in \mathscr{J}$. 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'.[32]

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 $\mathscr{J}$. □

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.[33] 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 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

[32] 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, $x_{n_j} > x_{n_{j'}}$.

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

(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 $\{(x_1, ..., x_N) \in [0, 1]^N : x_1 \leq ... \leq x_N\}$.

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 $p^A_{\underline{n}^A} = \theta_{\underline{n}^A} > p^B_{\underline{n}^A}$, contradicting our previous assumption which implies $p^A_{\underline{n}^A} = p^B_{\underline{n}^A}$. 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

$$\frac{x^A_n}{x^A_{n-1}} = \frac{v_n(p^A_{n-1})}{v_n(p^A_n)} = \frac{v_n(p^B_{n-1})}{v_n(p^B_n)} = \frac{x^B_n}{x^B_{n-1}}.$$

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

$$\frac{x^A_n}{x^A_{\underline{n}}} = \frac{x^B_n}{x^B_{\underline{n}}}.$$

Since $x^A_N = x^B_N = 1$, we have $x^A_n = x^B_n$ for all n, but then the mechanisms $(x^A_n, p^A_n)^N_{k=1}$ and $(x^B_n, p^B_n)^N_{k=1}$ 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 $x^C_n = \frac{1}{2}x^A_n + \frac{1}{2}x^B_n$ for report θ_n. For θ_n such that $x^C_n > 0$, it specifies H^C_n to put mass $\frac{x^A_n}{x^A_n + x^B_n}$ on p^A_n and the remaining mass on p^B_n. Incentive compatibility of the new mechanism is equivalent to the requirement that, for all n, k,

$$\begin{aligned}&\left(\tfrac{1}{2}x^A_n + \tfrac{1}{2}x^B_n\right)\left(\tfrac{x^A_n}{x^A_n + x^B_n} v_n(p^A_n) + \tfrac{x^B_n}{x^A_n + x^B_n} v_n(p^B_n)\right)\\ &\quad \geq \left(\tfrac{1}{2}x^A_k + \tfrac{1}{2}x^B_k\right)\left(\tfrac{x^A_k}{x^A_k + x^B_k} v_n(p^A_k) + \tfrac{x^B_k}{x^A_k + x^B_k} v_n(p^B_k)\right)\end{aligned}$$

or

$$x^A_n v_n(p^A_n) + x^B_n v_n(p^B_n) \geq x^A_k v_n(p^A_k) + x^B_k v_n(p^B_k).$$

This inequality holds by incentive compatibility of $(x^A_n, p^A_n)^N_{n=1}$ and $(x^B_n, p^B_n)^N_{n=1}$, so $(x^C_n, H^C_n)^N_{n=1}$ is incentive compatible. Individual rationality similarly is inherited from $(x^A_n, p^A_n)^N_{n=1}$ and $(x^B_n, p^B_n)^N_{n=1}$. Moreover, it is readily checked that the new mechanism $(x^C_n, H^C_n)^N_{n=1}$ attains the same optimal profit as $(x^A_n, p^A_n)^N_{n=1}$ and $(x^B_n, p^B_n)^N_{n=1}$. However, it does not have deterministic payments, which contradicts Lemma 2. □

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 $\beta_2 \geq \overline{\beta} \equiv \frac{\theta_1}{\theta_1 - \frac{v_2(\theta_1)}{v_2'(\theta_2)}} > \frac{\theta_1}{\theta_2}$.[34] 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 $\beta_2 \leq \underline{\beta} \equiv \frac{\theta_1}{\theta_1 - \frac{v_2(\theta_1)}{v_2'(\theta_1)}} < \frac{\theta_1}{\theta_2}$.[35] 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. **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 $\lambda_{n_1} = r x_{n_1}^* / (1 - x_{n_1}^*)$. 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 $m_j^\Lambda \Lambda$ the state changes to state σ_{j-1}, where m_j^Λ is obtained below.

For each $j = 2, ..., J-1$, we choose m_j^Λ 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:

$$x_{n_{j-1}}^* = x_{n_j}^* \left(\frac{\Lambda}{\Lambda + m_j^\Lambda \Lambda + r} x_{n_{j-1}}^* + \frac{m_j^\Lambda \Lambda}{\Lambda + m_j^\Lambda \Lambda + r} \right) \Rightarrow m_j^\Lambda = \frac{((1 - x_{n_j}^*)\Lambda + r) x_{n_{j-1}}^*}{\Lambda (x_{n_j}^* - x_{n_{j-1}}^*)} > 0 \, .$$

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

[34] 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 $v_2(\theta_2) = 0$.

[35] 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 $v_2(\theta_2) = 0$.

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

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

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 4

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

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

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 $\tilde{\tau}^t_1 > s$, and let $K'_{t,s}$ be its complement. Then there is $s > t$ such that $\Pr(K_{t,s}) > 0$ and $\mathbb{E}\left[e^{-r(\tilde{\tau}^t_1-s)}\middle|K_{t,s}\right] < x^*_1$.

Since $\mathbb{E}\left[e^{-r(\tilde{\tau}^t_1-t)}\right] = x^*_1$, we have

$$x^*_1 = (1-\Pr(K_{t,s}))\mathbb{E}\left[e^{-r(\tilde{\tau}^t_1-t)}\middle|K'_{t,s}\right] + \Pr(K_{t,s})e^{-r(s-t)}\mathbb{E}\left[e^{-r(\tilde{\tau}^t_1-s)}\middle|K_{t,s}\right].$$

So, necessarily,

$$\big(1-\Pr(K_{t,s})\big)\mathbb{E}\left[e^{-r(\tilde{\tau}^t_1-t)}\middle|K'_{t,s}\right] > \big(1-\Pr(K_{t,s})e^{-r(s-t)}\big)x^*_1. \tag{7}$$

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

$$(1-\Pr(K_{t,s}))\mathbb{E}\left[e^{-r(\tilde{\tau}^t_1-t)}\middle|K'_{t,s}\right]v_2(\theta_1) + \Pr(K_{t,s})e^{-r(s-t)}v_2(p^*_2) \tag{8}$$

The first term of Equation (8) 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 (7)), 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 (8) 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. $\square$

Proof of Proposition 5

Proof. Note that for θ^*, the interior optimum price for the deterministic posted-price mechanism, we have

$$1 - F(\theta^*) - \theta^* f(\theta^*) = 0.$$

Now consider perturbing the optimal posted-price mechanism by introducing mechanisms that will induce a small interval of types to purchase with probability $\alpha \in (0,1)$. For $\varepsilon \in (0, \theta^* - \underline{\theta})$, we consider mechanisms in which types $\theta \geq \theta^*$ obtain the good with certainty, types $\theta \in [\theta^* - \varepsilon, \theta^*)$ obtain it with probability α, and types $\theta < \theta^* - \varepsilon$ do not obtain the good at all. To represent the original deterministic mechanism, we will set $\varepsilon = 0$.

Such a mechanism can be obtained by setting the payment for obtaining the good with probability α to $\theta^* - \varepsilon$ and setting the payment for purchasing with certainty so that type θ^* is indifferent between purchasing with probability α or 1. That is, the payment in case purchasing with certainty is $p(\varepsilon) \in (\theta^* - \varepsilon, \theta^*)$ satisfying

$$v(p(\varepsilon); \theta^*) = \alpha v(\theta^* - \varepsilon; \theta^*). \tag{9}$$

Note that, when $\varepsilon = 0$, we have $p^*(\varepsilon) = \theta^*$.

Let us verify that these payments induce the purchasing strategy of the buyer as specified above. All types above $\theta^* - \varepsilon$ prefer one of the options that involves receiving the good with positive probability to receiving it with probability zero, while all lower types prefer not receiving the good. Types in $[\theta^* - \varepsilon, p(\varepsilon)]$ prefer the option of acquiring the good with probability α, using that the payoff from acquiring with certainty is non-positive. That types in $(p(\varepsilon), \theta^*)$ prefer to acquire with probability α follows from Lemma 1. In particular, fix $\theta_k \in (p(\varepsilon), \theta^*)$ and $\theta_l = \theta^*$. By contraposition of the claim in Lemma 1, because $v(p(\varepsilon); \theta_l) \leq \alpha v(\theta^* - \varepsilon; \theta_l)$, we have $v(p(\varepsilon); \theta_k) < \alpha v(\theta^* - \varepsilon; \theta_k)$, establishing the result. That types $\theta > \theta^*$ strictly prefer to obtain the good with certainty follows a direct application of Lemma 1.

Now let us write profits from the new mechanism as

$$\Pi(\varepsilon) = \alpha(F(\theta^*) - F(\theta^* - \varepsilon))(\theta^* - \varepsilon) + (1 - F(\theta^*))p(\varepsilon).$$

We are interested in determining whether $\Pi(\varepsilon) > \Pi(0) = \theta^*(1 - F(\theta^*))$ for some $\varepsilon > 0$; i.e., whether a small perturbation in our class can deliver higher profits than the optimal deterministic mechanism. For this, it is useful to determine the derivatives of $p(\varepsilon)$ at $\varepsilon = 0$. Differentiating (9) with respect to ε yields

$$p'(\varepsilon) v'(p(\varepsilon); \theta^*) = -\alpha v'(\theta^* - \varepsilon; \theta^*).$$

Differentiating again with respect to ε yields

$$p''(\varepsilon)v'(p(\varepsilon);\theta^*) + p'(\varepsilon)^2 v''(p(\varepsilon);\theta^*) = \alpha v''(\theta^* - \varepsilon;\theta^*).$$

Substituting the previous equation, we have

$$p''(\varepsilon)v'(p(\varepsilon);\theta^*) + \left(-\alpha\frac{v'(\theta^* - \varepsilon;\theta^*)}{v'(p(\varepsilon);\theta^*)}\right)^2 v''(p(\varepsilon);\theta^*) = \alpha v''(\theta^* - \varepsilon;\theta^*)$$

or

$$p''(\varepsilon) = \frac{\alpha v''(\theta^* - \varepsilon;\theta^*) - \left(\alpha\frac{v'(\theta^* - \varepsilon;\theta^*)}{v'(p(\varepsilon);\theta^*)}\right)^2 v''(p(\varepsilon);\theta^*)}{v'(p(\varepsilon);\theta^*)}.$$

Now consider the derivative of profits with respect to ε. This is

$$\Pi'(\varepsilon) = \alpha f(\theta^* - \varepsilon)(\theta^* - \varepsilon) - \alpha(F(\theta^*) - F(\theta^* - \varepsilon)) + (1 - F(\theta^*))p'(\varepsilon).$$

Note therefore that

$$\begin{aligned}\Pi'(0) &= \alpha(\theta^*)\theta^* + (1 - F(\theta^*))p'(0)\\ &= \alpha f(\theta^*)\theta^* - (1 - F(\theta^*))\alpha.\end{aligned}$$

This is equal to zero by the optimality condition for θ^*.

Next, consider the second derivative:

$$\begin{aligned}\Pi''(\varepsilon) =& -\alpha f'(\theta^* - \varepsilon)(\theta^* - \varepsilon) - \alpha f(\theta^* - \varepsilon) - \alpha f(\theta^* - \varepsilon) + (1 - F(\theta^*))p''(\varepsilon)\\ =& -\alpha f'(\theta^* - \varepsilon)(\theta^* - \varepsilon) - 2\alpha f(\theta^* - \varepsilon)\\ &+ (1 - F(\theta^*))\frac{\alpha v''(\theta^* - \varepsilon;\theta^*) - \left(\alpha\frac{v'(\theta^* - \varepsilon;\theta^*)}{v'(p(\varepsilon);\theta^*)}\right)^2 v''(p(\varepsilon);\theta^*)}{v'(p(\varepsilon);\theta^*)}.\end{aligned}$$

Therefore,

$$\Pi''(0) = -\alpha f'(\theta^*)\theta^* - 2\alpha f(\theta^*) + (1 - F(\theta^*))\frac{(\alpha - \alpha^2)v''(\theta^*;\theta^*)}{v'(\theta^*;\theta^*)}.$$

We then observe that, if $\Pi''(0) > 0$, then $\Pi(\varepsilon) > \Pi(0)$ for $\varepsilon > 0$ sufficiently small. This condition can be written as

$$\alpha\left(-f'(\theta^*)\theta^* - 2f(\theta^*) + (1 - F(\theta^*))\frac{(1 - \alpha)v''(\theta^*;\theta^*)}{v'(\theta^*;\theta^*)}\right) > 0$$

for some $\alpha \in (0,1)$. There exists such an α if and only if

$$-f'(\theta^*)\theta^* - 2f(\theta^*) + (1 - F(\theta^*))\frac{v''(\theta^*;\theta^*)}{v'(\theta^*;\theta^*)} > 0.$$

This can be written as Equation (4). □

Proof of Proposition 6

Proof. First, note that we can use the allocation to define a function $L_m(\cdot)$ on $[\underline{\theta}, \bar{\theta})$ given for each θ by $L_m(\theta) = x_n^m$ if $\theta \in [\theta_n^m, \theta_{n+1}^m)$, where recall that $\theta_{N_m+1}^m = \bar{\theta}$. Note that since each $L_m(\cdot)$ is monotone and bounded, there is a pointwise convergent subsequence $(L_{m_k}(\cdot))$ by Helly's Selection Theorem. The following lemma concerns such a subsequence.

Lemma 3. *Suppose there is $\hat{\theta} \in [\underline{\theta}, \bar{\theta}]$ such that: (i) for all $\theta < \hat{\theta}$, $L_{m_k}(\theta) \to 0$, while (ii) for all $\theta > \hat{\theta}$, $L_{m_k}(\theta) \to 1$. Then profits in the optimal mechanisms along the sequence of environments (E^{m_k}) converge to $\hat{\theta}(1 - F(\hat{\theta}))$.*

Proof. **Proof of Lemma 3. Step 1.** Suppose the sequence $(L_{m_k}(\theta))$ satisfies the assumption of the lemma. We want to show that for any $\varepsilon > 0$, we can find K_ε large enough that, for all $k > K_\varepsilon$, $x_n^{m_k} < \varepsilon$ for all $\theta_n^{m_k} < \hat{\theta} - \varepsilon$ and $x_n^{m_k} > 1 - \varepsilon$ for all $\theta_n^{m_k} > \hat{\theta} + \varepsilon$. Otherwise, there is some $\varepsilon > 0$ such that this is not the case. Then, using monotonicity of each $L_{m_k}(\theta)$, either there is not convergence of $L_{m_k}(\hat{\theta} - \varepsilon)$ to zero where $\hat{\theta} - \varepsilon \geq \underline{\theta}$, or there is not convergence of $L_{m_k}(\hat{\theta} + \varepsilon)$ to 1 where $\hat{\theta} + \varepsilon \leq \bar{\theta}$, contradicting the assumption of the lemma.

Step 2. We now show that for any $\eta > 0$ we can find $Q_\eta > 0$ large enough that, for all $k > Q_\eta$, if n is such that $\theta_n^{m_k} > \hat{\theta} + \eta$ then $p_n^{m_k} > \hat{\theta} - \eta$.[36] Note first that, by Proposition 1, the claim must hold if $\hat{\theta} = \underline{\theta}$ so suppose that $\hat{\theta} > \underline{\theta}$ and suppose for a contradiction that the claim is not true. Then there is an $\eta' \in (0, \hat{\theta} - \bar{\theta})$ and a further subsequence of environments (E^{l_k}) (i.e., a subsequence of (E^{m_k})) such that, for each k, there is a type $\theta_{n_{l_k}}^{l_k} > \hat{\theta} + \eta'$ with $p_{n_{l_k}}^{l_k} \leq \hat{\theta} - \eta'$. Then, pick $\bar{\varepsilon} > 0$ but small enough that

$$(1 - \bar{\varepsilon})v(\hat{\theta} - \eta'; \hat{\theta} - \eta'/2) > \bar{\varepsilon}v(\underline{\theta}; \hat{\theta} - \eta'/2). \tag{10}$$

That such a value of $\bar{\varepsilon}$ exists follows because $v(\hat{\theta} - \eta'; \hat{\theta} - \eta'/2) > 0$. From Step 1, we have that there is a value k' large enough that we are assured of the existence of an $\hat{n}$ such that $\theta_{\hat{n}}^{l_{k'}} \in (\hat{\theta} - \eta'/2, \hat{\theta})$ and such type is assigned the good under the optimal mechanism for environment $E^{l_{k'}}$ with a probability no greater than $\bar{\varepsilon}$, while type $\theta_{n_{l_{k'}}}^{l_{k'}}$ receives the good with probability at least $1 - \bar{\varepsilon}$.

First note that, by the inequality (10) and Lemma 1, we have

$$(1 - \bar{\varepsilon})v(\hat{\theta} - \eta'; \theta_{\hat{n}}^{l_{k'}}) > \bar{\varepsilon}v(\underline{\theta}; \theta_{\hat{n}}^{l_{k'}}).$$

We therefore have

$$x_{n_{l_{k'}}}^{l_{k'}} v(p_{n_{l_{k'}}}^{l_{k'}}; \theta_{\hat{n}}^{l_{k'}}) > x_{\hat{n}}^{l_{k'}} v(p_{\hat{n}}^{l_{k'}}; \theta_{\hat{n}}^{l_{k'}})$$

after using that $x_{n_{l_{k'}}}^{l_{k'}} \geq 1 - \bar{\varepsilon}$, $p_{n_{l_{k'}}}^{l_{k'}} \leq \hat{\theta} - \eta'$, $x_{\hat{n}}^{l_{k'}} \leq \bar{\varepsilon}$, and $p_{\hat{n}}^{l_{k'}} \geq \underline{\theta}$. Therefore, the assumed optimal mechanism in environment $E^{l_{k'}}$ is not incentive compatible, as type $\theta_{\hat{n}}^{l_{k'}}$ strictly prefers the option designed

[36] Recall that, as stated in the main text, p_n^m represents the price paid by type θ_n^m in the optimal mechanism for environment E^m.

for type $\theta^{l_{k'}}_{n_{l_{k'}}}$. This is a contradiction.

Step 3. We show that for any $\eta > 0$, there exists Q_η large enough that, for all $k > Q_\eta$, all types in the optimal mechanism for environment E^{m_k} pay no more than $\hat{\theta}+\eta$. Note that this is clearly true when $\hat{\theta} = \bar{\theta}$, by individual rationality of optimal mechanisms. So suppose that $\hat{\theta} < \bar{\theta}$ and suppose for a contradiction that the claim is not true. Then there is an $\eta' > 0$ and a further subsequence of environments (E^{l_k}) along which there is some type paying more than $\hat{\theta} + \eta'$; without loss of generality let this be the highest type $\theta^{l_k}_{N_{l_k}}$ (recall that payments in the optimal mechanism are increasing in the buyer's type). By Step 1, there is a choice of types $(\theta^{l_k}_{n_{l_k}})$ such that $x^{l_k}_{n_{l_k}} \to 1$ and $\theta^{l_k}_{n_{l_k}} \to \hat{\theta}$. Note then that types $(\theta^{l_k}_{n_{l_k}})$ are assigned by the optimal mechanism a probability of allocation approaching one, and these types pay no more than $\theta^{l_k}_{n_{l_k}}$.

Now note that

$$v(\hat{\theta};\bar{\theta}) > v(\hat{\theta}+\eta';\bar{\theta}).$$

Using continuity of the function $xv(y;z)$ in (x,y,z), for all k sufficiently large, we have

$$\begin{aligned} x^{l_k}_{n_{l_k}} v(\theta^{l_k}_{n_{l_k}};\theta^{l_k}_{N_{l_k}}) &> x^{l_k}_{N_{l_k}} v(\hat{\theta}+\eta';\theta^{l_k}_{N_{l_k}}) \\ &\geq x^{l_k}_{N_{l_k}} v(p^{l_k}_{N_{l_k}};\theta^{l_k}_{N_{l_k}}). \end{aligned}$$

We conclude that, for all k sufficiently large, type $\theta^{l_k}_{N_{l_k}}$ strictly prefers to mimic type $\theta^{l_k}_{n_{l_k}}$ than to report truthfully, implying a violation of incentive compatibility of the optimal mechanism.

Step 4. We have established that, for any $\eta > 0$, there is K large enough that the following hold for all $k > K$: (i) types $\theta^{m_k}_n$ above $\hat{\theta}+\eta$ make payments in the optimal mechanism for model E^{m_k} that are within η of $\hat{\theta}$, (ii) types $\theta^{m_k}_n$ above $\hat{\theta}+\eta$ acquire the good with probability at least $1-\eta$, (iii) types $\theta^{m_k}_n$ below $\hat{\theta}-\eta$ acquire the good with a probability no greater than η. This permits us to conclude that an upper bound on revenues for the optimal mechanism in environment E^{m_k} with $k > K$ is given by

$$\eta F(\hat{\theta}-\eta)(\hat{\theta}-\eta) + (\hat{\theta}+\eta)(1-F(\hat{\theta}-\eta)).$$

Now, take K large enough so that, in addition to points (i)-(iii), the probability of types above $\hat{\theta}+\eta$ is at least $1-F(\hat{\theta}+2\eta)$ for all $k > K$. Then a lower bound on revenue in the optimal mechanism in environment E^{m_k} with $k > K$ is given by

$$(1-F(\hat{\theta}+2\eta))(1-\eta)(\hat{\theta}-\eta).$$

This follows because, for $k > K$, types above $\hat{\theta}+\eta$, which have probability at least $1-F(\hat{\theta}+2\eta)$, acquire the good with probability at least $1-\eta$ and when doing so pay at least $\hat{\theta}-\eta$.

Finally, using continuity of F, both lower and upper bounds converge to $\hat{\theta}(1-F(\hat{\theta}))$ as $\eta \to 0$, which shows that profits converge to $\hat{\theta}(1-F(\hat{\theta}))$ considering optimal mechanisms along the sequence of environments (E^{m_k}). $\square$

We now show that the convergence hypothesized in the previous lemma cannot occur.

Lemma 4. *Suppose that Condition (4) is satisfied. Consider any subsequence (E^{m_k}) of (E^m) such that the corresponding sequence of functions $(L_{m_k}(\theta))$ converges pointwise. Then there is no $\hat{\theta} \in [\underline{\theta}, \bar{\theta}]$ such that $L_{m_k}(\theta) \to 0$ for $\theta < \hat{\theta}$ and $L_{m_k}(\theta) \to 1$ for $\theta > \hat{\theta}$.*

Proof. **Proof of Lemma 4.** Suppose for a contradiction that there is such a $\hat{\theta}$ with $L_{m_k}(\theta) \to 0$ for $\theta < \hat{\theta}$ and $L_{m_k}(\theta) \to 1$ for $\theta > \hat{\theta}$. Then, by Lemma 3, optimal profits along the sequence of environments (E^{m_k}) converge to a value no greater than $\theta^*(1-F(\theta^*))$, where recall that θ^* maximizes $\theta(1-F(\theta))$. However, we argue that there exists $\eta > 0$ such that, for all large enough k, profits are at least $\theta^*(1-F(\theta^*)) + \eta$.

This can be seen by adapting the proof of Proposition 5. In particular, pick $\alpha > 0$ and $\varepsilon > 0$ small enough that the perturbed mechanism in that proof generates strictly higher profits than the optimal deterministic mechanism which is a take-it-or-leave-it offer with payment θ^*. This mechanism offers a probability of awarding the good α with a payment $\theta^* - \varepsilon$ upon award, and a payment $p(\varepsilon) \in (\theta^* - \varepsilon, \theta^*)$ to receive the good with certainty. We saw that only types above $\theta^* - \varepsilon$ pick one of these options, with all types in $(\theta^* - \varepsilon, \theta^*)$ picking probabilistic award, and types above θ^* choosing award with certainty. Now consider these mechanisms in the environments E^{m_k}. As $k \to \infty$ the probability that the probabilistic option is chosen converges to $F(\theta^*) - F(\theta^* - \varepsilon)$. The probability that the sure option is chosen converges to $1 - F(\theta^*)$. Therefore, using the same calculations as in the proof of Proposition 5, profits converge to a level strictly greater than $\theta^*(1 - F(\theta^*))$. In particular, for any k sufficiently large, we have that the specified mechanism generates profits that are above $\theta^*(1 - F(\theta^*)) + \eta$ for some fixed $\eta > 0$. This contradicts the supposed optimality of mechanisms in environments E^{m_k} for large k. □

Proof of the proposition. We now conclude the proof of Proposition 6. In environment E^m, attribute a property to a "sequence of types of length at least ε" if the property is satisfied for some adjacent types $\theta^m_{n'}, \theta^m_{n'+1}, \ldots, \theta^m_{n''}$ with $\theta^m_{n''} - \theta^m_{n'} \geq \varepsilon$. Suppose for a contradiction that the result in the Proposition is not true. Then, for any $\varepsilon > 0$ and any $Z \in \mathbb{N}$, we can find $m > Z$ such that, in the optimal mechanism for environment E^m, there is no sequence of types of length at least ε for which the allocation probability is in $[\varepsilon, 1-\varepsilon]$. This means that there are three mutually exclusive possibilities: (i) there is a smallest type $\theta^m_{n'_m}$ for which $x^m_{n'_m} \geq \varepsilon$ and a largest type $\theta^m_{n''_m}$ for which $x^m_{n''_m} \leq 1-\varepsilon$, and $\theta^m_{n''_m} - \theta^m_{n'_m} < \varepsilon$, (ii) the allocation for all types is strictly below ε, and (iii) the allocation for all types is strictly above $1-\varepsilon$.

We can pick a subsequence (E^{m_k}) where, for each k, in the optimal mechanism of environment E^{m_k}, there is no sequence of types of length at least $1/k$ for which the allocation probability is in $[1/k, 1-1/k]$. Along this subsequence, one of (i)-(iii) occurs infinitely often, taking ε to equal $1/k$. That is, one of the following occur infinitely often: (i) there is a smallest type $\theta^{m_k}_{n'_{m_k}}$ for which $x^{m_k}_{n'_{m_k}} \geq 1/k$ and a largest type $\theta^{m_k}_{n''_{m_k}}$ for which $x^{m_k}_{n''_{m_k}} \leq 1 - 1/k$, and $\theta^{m_k}_{n''_{m_k}} - \theta^{m_k}_{n'_{m_k}} < 1/k$, (ii) the allocation for all types in the optimal mechanism of environment E^{m_k} is strictly below $1/k$, and (iii) the allocation for all types in the optimal mechanism of environment E^{m_k} is strictly above $1 - 1/k$.

If (ii) occurs infinitely often, then pick a subsequence that we now denote $(E^{m_{k_j}})$ along which it always occurs. Then $(L_{m_{k_j}}(\theta))$ converges pointwise to a constant 0. If (iii) occurs infinitely often, then pick a subsequence $(E^{m_{k_j}})$ along which it always occurs. Then $(L_{m_{k_j}}(\theta))$ converges pointwise to a constant 1. In either case, we have a violation of Lemma 4. So suppose (i) occurs infinitely often and pick a subsequence $(E^{m_{k_j}})$ along which it always occurs. Recall that $\theta^{m_{k_j}}_{n'_{m_{k_j}}}$ is the smallest type for which the probability of allocation is at least $1/k_j$. We can pick a further subsequence of $(E^{m_{k_j}})$, call it (E^{q_l}), such that $\theta^{q_l}_{n'_{q_l}} \to \hat{\theta}$ for some $\hat{\theta} \in [\underline{\theta}, \bar{\theta}]$. That such a subsequence exists follows from the Bolzano-Weierstrass Theorem. By Helly's Selection Theorem, we may suppose this subsequence is such that $(L_{q_l}(\theta))$ is convergent pointwise. For any $\theta < \hat{\theta}$, we have $L_{q_l}(\theta) \to 0$ and for any $\theta > \hat{\theta}$,we have $L_{q_l}(\theta) \to 1$. Again we have a violation of Lemma 4. This completes the proof of Proposition 6. $\square$

Proof of Proposition 7

Proof. To characterize an optimal static mechanism, first the usual replication argument permits us to conclude that it is without loss of generality to consider direct mechanisms. As in the baseline model, we may assume that the direct mechanism specifies, for each type θ_n, an allocation probability x_n and a distribution over payments conditional on allocation, H_n. We first show that it is optimal for the seller to specify deterministic prices.

Lemma 5. *It is optimal to set a payment equal to θ_1 for the low type and a deterministic payment $p_2 \in [\theta_1, \theta_2)$ for the high type.*

Proof. **Proof of Lemma 5.** The utility of type θ_n when reporting θ_k is

$$U_{n,k} = x_k \int \left(\theta_n - p + \mu\left(\rho_n - p\right)\right) dH_k\left(p\right).$$

Incentive compatibility is the requirement that, for all θ_n and θ_k, $U_{n,k} \leq U_{n,n}$. Individual rationality is the requirement that, for all θ_n, $U_{n,n} \geq 0$.

Consider an incentive-compatible and individually-rational mechanism. We want to show that there is a weakly more profitable mechanism with the properties in the lemma.

First, note that we can focus on the case where the original mechanism allocates with positive probability to the low type. This is because, if the probability of allocation to the low type is zero, the seller can (without sacrificing profits) specify a payment for the high type equal to the high type's willingness to pay.[37] The payment specified to the low type is irrelevant since it is never charged (so we can let this equal θ_1).

Next note individual rationality implies that, in any mechanism, the low type does not make an expected payment higher than θ_1 for receiving the good. We may then assume that the high type makes

[37] This is the price p that sets $\theta_2 - p + \lambda\eta\left(\rho_2 - p\right) = 0$. That is, the high type's willingness to pay is $\frac{\theta_2+\lambda\eta\rho_2}{1+\lambda\eta} \in (\rho_2, \theta_2)$.

an expected payment at least θ_1 for receiving the good, otherwise the seller could instead offer the weakly more profitable mechanism that has both types purchase at price θ_1 (say a take-it-or-leave-it offer for the good at price θ_1). This mechanism satisfies the properties in the lemma.

We can take as given then that the low type is awarded the good with positive probability and the price distribution conditional on purchase is given by H_1. We can replace the random payment with a deterministic payment which is the certainty-equivalent payment for the low type, (weakly) raising the seller's profits because the low type's payoff is concave in the payment. Because the low type earns a non-negative payoff in the original mechanism (by individual rationality), this certainty-equivalent payment is no greater than θ_1. Denote this new payment $\bar{p}_1$. It is given by

$$(\theta_1 - \bar{p}_1)(1+\eta) = \bar{U}_1$$

where

$$\bar{U}_1 = \int_{[0,\theta_1)} (\theta_1 - p + \eta(\theta_1 - p))\, dH_1(p) + \int_{[\theta_1,\infty)} (\theta_1 - p + \lambda\eta(\theta_1 - p))\, dH_1(p)$$

is θ_1's payoff in the original mechanism. We therefore have

$$\bar{p}_1 = \theta_1 - \bar{U}_1/(1+\eta).$$

We now show that θ_2's incentive compatibility constraint is satisfied in the new mechanism (θ_1's incentive compatibility constraint is unaffected by the adjustment to the mechanism, and both individual-rationality constraints are also unaffected). Note that θ_2's payoff from mimicking θ_1 in the original mechanism is

$$x_1 \left(\int_{[0,\rho_2)} (\theta_2 - p + \eta(\rho_2 - p))\, dH_1(p) + \int_{[\rho_2,\infty)} (\theta_2 - p + \lambda\eta(\rho_2 - p))\, dH_1(p) \right). \tag{11}$$

After the adjustment to the mechanism (where we replaced θ_1's payment by $\bar{p}_1$), θ_2's payoff from mimicry is

$$x_1(\theta_2 - \bar{p}_1 + \eta(\rho_2 - \bar{p}_1)). \tag{12}$$

The decrease in payoff for type θ_2 when mimicking θ_1 due to the change in the mechanism is the expression in Equation (11) less that in Equation (12). This is equal to

$$x_1 \begin{pmatrix} \int_{[\theta_1,\rho_2]} \eta(\lambda-1)(p-\theta_1)\, dH_1(p) \\ +\eta(1-H_1(\rho_2))(\lambda-1)(\rho_2-\theta_1) \end{pmatrix} \geq 0.$$

We conclude that, given the original mechanism satisfied the high type's incentive constraint, the new mechanism does too. In addition, by the concavity of type θ_1's payoffs, the new mechanism is weakly more profitable for the seller. So we have found a weakly more profitable mechanism that is incentive compatible and individually rational, and where the low type makes a deterministic payment $\bar{p}_1$.

Next, noting $\bar{p}_1 \le \theta_1$, consider further (weakly) raising the payment of the low type to θ_1. This relaxes the high type's incentive constraint, keeps individual rationality constraints intact, and raises the profits of the mechanism. To see that the low type does not prefer to mimic the high type, recall we could assume that the high type makes an expected payment upon receiving the good of at least θ_1. The low type therefore earns a non-positive payoff from mimicking the high type.

Now make a final adjustment to the mechanism by replacing the high type's payment with its certainty equivalent. Because the high type has concave payoffs, this payment is again at least θ_1. Again, the low type's incentive constraint, and hence all incentive and individual rationality constraints remain intact.

We have constructed, then, our weakly more profitable mechanism where the low type pays θ_1 and the high type pays at least θ_1. Because the high type obtains a non-negative payoff, the high type's payment as determined above is no greater than his willingness to pay, i.e. $\frac{\theta_2+\lambda\eta\rho_2}{1+\lambda\eta}$ (see footnote 37). This is strictly less than θ_2 as we assumed $\rho_2 < \theta_2$. □

Now, to establish Proposition 7, consider maximizing seller profits among mechanisms where the payment distribution H_n is degenerate for each n. Among mechanisms in the class considered in Lemma 5, we only need to impose the high type's incentive constraint.[38] This constraint can be written as:

$$x_2\left(\theta_2 - p_2 + \mu\left(\rho_2 - p_2\right)\right) \ge x_1\left(\theta_2 - \theta_1 + \eta\left(\rho_2 - \theta_1\right)\right) \tag{13}$$

where $p_2 \in [\theta_1, \theta_2)$ is the price charged to the high type θ_2. Here, the left-hand side is the high type's payoff from truthful reporting, while the right-hand side is the payoff from mimicking θ_1, in which case the buyer makes a payment θ_1 when receiving the good. When the constraint (13) is satisfied, the high type earns a non-negative rent. Setting $x_2 = 1$ therefore relaxes the constraint. It does so without introducing any violation of the low type's incentive constraint or in either of the individual rationality constraints. Therefore, it is indeed profit maximizing to set $x_2 = 1$. Moreover, raising p_2, provided it does not lead to a violation of the constraint (13), increases profits. We can therefore assume that the constraint binds. That is, p_2 is given by

$$\theta_2 - p_2 + \mu\left(\rho_2 - p_2\right) = x_1\left(\theta_2 - \theta_1 + \eta\left(\rho_2 - \theta_1\right)\right).$$

Viewing p_2 now as a (decreasing) function of x_1, there is a value of x_1, call it $\bar{x}_1$, at which the high type's payment equals the reference point. This value is given by

$$\bar{x}_1 = \frac{\theta_2 - \rho_2}{\theta_2 - \theta_1 + \eta\left(\rho_2 - \theta_1\right)}.$$

For $x_1 > \bar{x}_1$ the high type's payment is below the reference point and the high type experiences gains. For $x_1 < \bar{x}_1$, the high type's payment is above the reference point and so the high type experiences losses.

[38]The low type's incentive constraint is satisfied whenever the high type pays at least θ_1. Because the low type pays his valuation θ_1 for the good, the high type's individual rationality constraint is then satisfied whenever his incentive constraint is satisfied.

We can observe that, for $x_1 \geq \bar{x}_1$, we have

$$p_2 = \frac{\theta_2 + \eta\rho_2 - x_1(\theta_2 - \theta_1 + \eta(\rho_2 - \theta_1))}{1+\eta}.$$

For $x_1 \leq \bar{x}_1$, we have

$$p_2 = \frac{\theta_2 + \eta\lambda\rho_2 - x_1(\theta_2 - \theta_1 + \eta(\rho_2 - \theta_1))}{1+\eta\lambda}.$$

We conclude that p_2 decreases in x_1 at rate

$$\frac{\theta_2 - \theta_1 + \eta(\rho_2 - \theta_1)}{1+\eta\lambda}$$

for x_1 below $\bar{x}_1$, and at rate

$$\frac{\theta_2 - \theta_1 + \eta(\rho_2 - \theta_1)}{1+\eta}$$

above $\bar{x}_1$. Since the former is smaller than the latter, this shows that profits, as a function of x_1, have a concave kink at $\bar{x}_1$.

It is then readily checked that $x_1 = \bar{x}_1$ in the optimal mechanism if and only if

$$-\beta_2 \frac{\theta_2 - \theta_1 + \eta(\rho_2 - \theta_1)}{1+\eta} + \theta_1(1-\beta_2) \leq 0 \leq -\beta_2 \frac{\theta_2 - \theta_1 + \eta(\rho_2 - \theta_1)}{1+\eta\lambda} + \theta_1(1-\beta_2).$$

This is equivalent to Condition (5). □

Appendix II: Conditions for attaining Π^*

This Appendix proves the statement at the end of Section 4: there exists a price process attaining Π^* if and only if there are no two values $x^*_{n'}, x^*_{n''} \in (0,1)$ with $x^*_{n'} \neq x^*_{n''}$. The "if" part follows from the arguments in the proof of Proposition 3. To show the "only if" part, we first introduce incentive-compatible price processes and their optimality. Then we state and show the result.

Let $\langle \Omega, (\mathscr{F}_t)_{t\in[0,\infty)}, \mathscr{P} \rangle$ 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 $\langle \Omega, (\mathscr{F}_t)_{t\in[0,\infty)}, \mathscr{P} \rangle$.*
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 $\omega \in \Omega$.*[39]

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.

[39]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.

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 8. *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 $(\mathscr{F}_t)_t$ is the filtration generated by P. We divide the proof into three steps:

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

$$\mathbb{E}[e^{-r(\tau_{n,t}-t)}|A_t] \neq x_n^* \tag{14}$$

or (ii) there is $A_t \in \mathscr{F}_t$ with $\mathscr{P}(A_t) > 0$ such that $\mathscr{P}(P_{\tau_{n,t}} = p_n^*|A_t) < 1$. The result to be established is that the set $\mathscr{T}$ 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 $t \in \mathscr{T}$ we can find $A_t \in \mathscr{F}_t$ (chosen so that $\mathbb{E}[e^{-r(\tau_{n,t}-t)}|A_t] \neq x_n^*$ or $\mathscr{P}(P_{\tau_{n,t}} = p_n^*|A_t) < 1$ for some n) for which $\mathbb{E}[\Pi_t|A_t] < \Pi^*$. In addition, we know that, for all t, and any $B \in \mathscr{F}_t$, $\mathbb{E}[\Pi_t|B] \leq \Pi^*$. Therefore, taking $B = \Omega \backslash A_t$, we obtain

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

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

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

$$\mathscr{P}\left(\tau_{n,t} < \tau_{m,t} \text{ or } \tau_{n,t} = \tau_{m,t} = \infty \,\middle|\, A \right) = 1$$

for all n and all m such that $x_n^* > x_m^*$. This is immediate when $x_m^* = 0$, since then $\mathscr{P}\left(\tau_{m,t} = \infty\right) = 1$. If instead $x_m^* > 0$, we argue that $\mathscr{P}(\tau_{n,t} < \tau_{m,t} \text{ or } \tau_{n,t} = \tau_{m,t} = \infty \,|\, A) < 1$ 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 6. *Fix some n and m, with $n > m$ and with $x_n^* \in (0,1)$. Then, for $t, t', t'' \in \mathbb{R}_+ \backslash \mathscr{T}$ with $t \leq t' < t'' \leq \infty$, we have $\mathscr{P}(K_{n,t}^{t',t''}) > 0$ and*

$$\frac{x_m^*}{x_n^*} = \frac{\mathbb{E}\left[e^{-r(\tau_{m,t}-t')} \middle| K_{n,t}^{t',t''}\right]}{\mathbb{E}\left[e^{-r(\tau_{n,t}-t')} \middle| K_{n,t}^{t',t''}\right]}\,. \tag{15}$$

Proof. We first want to show that, for $t, t', t'' \in \mathbb{R}_+ \backslash \mathscr{T}$ with $t \leq t' < t'' \leq \infty$, $\mathscr{P}(K_{n,t}^{t',t''}) > 0$. This will be a result of what we call Claim A: For $t, t_1, t_2 \in \mathbb{R}_+ \backslash \mathscr{T}$ with $t \leq t_1 < t_2 < t_1 - \log(x_n^*)/r$ and with $\mathscr{P}(K_{n,t}^{t_1,\infty}) > 0$, we have $\mathscr{P}(K_{n,t}^{t_1,t_2}), \mathscr{P}(K_{n,t}^{t_2,\infty}) > 0$. Given $t, t', t'' \in \mathbb{R}_+ \backslash \mathscr{T}$ with $t \leq t' < t'' \leq \infty$, we then arrive at $\mathscr{P}(K_{n,t}^{t',t''}) > 0$ 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 $t_1, t_2 \in \mathbb{R}_+ \backslash \mathscr{T}$ with $t \leq t_1 < t_2 < t_1 - \log(x_n^*)/r$ and with $\mathscr{P}(K_{n,t}^{t_1,\infty}) > 0$. Applying Step 1, after noting $K_{n,t}^{t_1,\infty} \in \mathscr{F}_{t_1}$, we have:[40]

$$\begin{aligned} x_n^* &= \mathbb{E}[e^{-r(\tau_{n,t}-t_1)} | K_{n,t}^{t_1,\infty}] \\ &= \frac{\mathscr{P}(K_{n,t}^{t_1,t_2})}{\mathscr{P}(K_{n,t}^{t_1,\infty})} \underbrace{\mathbb{E}\left[e^{-r(\tau_{n,t}-t_1)} \middle| K_{n,t}^{t_1,t_2}\right]}_{(*)} + \frac{\mathscr{P}(K_{n,t}^{t_2,\infty})}{\mathscr{P}(K_{n,t}^{t_1,\infty})} \underbrace{e^{-r(t_2-t_1)}}_{(**)} \underbrace{\mathbb{E}\left[e^{-r(\tau_{n,t}-t_2)} \middle| K_{n,t}^{t_2,\infty}\right]}_{(***)}\,. \end{aligned} \tag{16}$$

[40]Note that, in expressions such as equation (16), if the conditioning event for a conditional expectation has probability zero, we take the conditional expectation to equal zero.

Note first that, if $\mathscr{P}(K_{n,t}^{t_1,t_2})>0$, then the term $(*)$ is no smaller than $e^{-r(t_2-t_1)}$. It then follows that, because $t_2<t_1-\log(x_n^*)/r$, $\mathscr{P}(K_{n,t}^{t_1,t_2})/\mathscr{P}(K_{n,t}^{t_1,\infty})<1$ and hence $\mathscr{P}(K_{n,t}^{t_2,\infty})>0$. Note also that, since $K_{n,t}^{t_2,\infty}\in\mathscr{F}_{t_2}$ 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 $\mathscr{P}(K_{n,t}^{t_2,\infty})>0$ (from Step 1). Furthermore, given $t_2>t_1$, we have that $({*}{*})$ is strictly smaller than 1; hence it must be that $\mathscr{P}(K_{n,t}^{t_1,t_2})>0$. 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 $t,t',t''\in\mathbb{R}_+\backslash\mathscr{T}$ with $t\le t'<t''\le\infty$. Observe that

$$\begin{aligned}
x_m^* &= \mathbb{E}[e^{-r(\tau_{m,t'}-t')}|K_{n,t}^{t',\infty}]\\
&= \mathbb{E}[e^{-r(\tau_{m,t}-t')}|K_{n,t}^{t',\infty}]\\
&= \frac{\mathscr{P}(K_{n,t}^{t',t''})}{\mathscr{P}(K_{n,t}^{t',\infty})}\mathbb{E}\big[e^{-r(\tau_{m,t}-t')}\big|K_{n,t}^{t',t''}\big]+\frac{\mathscr{P}(K_{n,t}^{t'',\infty})}{\mathscr{P}(K_{n,t}^{t',\infty})}e^{-r(t''-t')}\underbrace{\mathbb{E}\big[e^{-r(\tau_{m,t}-t'')}\big|K_{n,t}^{t'',\infty}\big]}_{(**)}\,.
\end{aligned}\tag{17}$$

The second equality holds because $\tau_{m,t}(\omega)\ge t'$ for almost all $\omega\in K_{n,t}^{t',\infty}$ (since $\mathscr{P}(\tau_{n,t}<\tau_{m,t}$ 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 $K_{n,t}^{t'',\infty}\in\mathscr{F}_{t''}$, and since $\tau_{m,t}=\tau_{m,t''}$ on $K_{n,t}^{t'',\infty}$, we have that $({*}{*})$ is equal to x_m^*. Considering Equation (17) for distinct m and n as in the lemma, as well as m taken equal to n, generates two equations which together imply Equation (15). □

(End of proof of Lemma 6, proof of Proposition 8 continues.)

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

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

Since $K_{n,t}^{t',t''}\cap\overline{K}_{m,t}^{t',t''}\in\mathscr{F}_{t''}$, we have that $({*}{*})$ is equal to x_m^*. Now, let $\bar\delta=\frac{-1}{r}\log\big(\frac{1}{2}x_n^*+\frac{1}{2}\big)$, 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 (15) (to replace $(*)$) and that $\mathscr{P}\big(\overline{K}_{m,t}^{t',t''}\big|K_{n,t}^{t',t''}\big)=1-\mathscr{P}\big(K_{m,t}^{t',t''}\big|K_{n,t}^{t',t''}\big)$, we obtain

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

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

$$\mathscr{P}\left(K_{m,t}^{t',t''} \middle| K_{n,t}^{t',t''}\right) \geq \frac{x_m^*(1-x_n^*)}{x_n^*(2-(1+x_n^*)x_m^*)} > 0. \tag{18}$$

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

$$\begin{aligned} \mathscr{P}\left(K_{n,m,t}^{\delta} | \tau_{n,t}(\omega) < \infty\right) &\geq \sum_{k=0}^{\infty} \mathscr{P}\left(K_{n,t}^{t_k,t_{k+1}} | \tau_{n,t}(\omega) < \infty\right) \mathscr{P}\left(K_{m,t}^{t_k,t_{k+1}} \middle| K_{n,t}^{t_k,t_{k+1}}\right) \\ &\geq \frac{x_m^*(1-x_n^*)}{x_n^*(2-(1+x_n^*)x_m^*)} \,. \end{aligned} \tag{19}$$

The first inequality holds because $K_{m,t}^{t_k,t_{k+1}} \cap K_{n,t}^{t_k,t_{k+1}} \subset K_{n,m,t}^{\delta} \cap K_{n,t}^{t_k,t_{k+1}}$. The second inequality follows from Equation (18). Considering Equation (19) as $\delta \to 0$, we obtain a contradiction to the previous observation that $\lim_{\delta\to 0} \mathscr{P}\left(K_{n,m,t}^{\delta}\right)=0$.. □