

Discussion Paper Series – CRC TR 224

Discussion Paper No. 688

Project B 02

A Dynamic Theory of Random Price Discounts

Francesc Dilmé¹

Daniel F. Garrett²

May 2025

¹University of Bonn, Email : fdilme@uni-bonn.de

²University of Essex, Email : d.garrett@essex.ac.uk

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

A Dynamic Theory of Random Price Discounts

Francesc Dilmé[†]

Daniel F. Garrett[‡]

22nd April 2025

Abstract

A seller with commitment power sets prices over time. Risk-averse buyers arrive to the market and decide when to purchase. We show that it is optimal for the seller to choose a constant high price punctuated by occasional episodes of sequential discounts that occur at random times. This 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

For helpful discussions and comments, we would like especially to thank Pierre Dubois, Drew Fudenberg, Matthew Gentry, Holger Gerhardt, Paul Heidhues, Stephan Lauerermann, Niccolò Lomys, Volker Nocke, Alessandro Pavan, Martin Pesendorfer, Andrew Rhodes, Anton Sobolev, Philipp Strack, Curtis Taylor, and Yang Yang. We would also like to thank seminar participants at numerous conferences and invited seminars at which this paper was presented. This project was initiated when Garrett was at Toulouse School of Economics, University of Toulouse Capitole. 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). Garrett thanks support from the French National Research Agency (ANR) under the Investments for the Future (Investissements d’Avenir) program (grant ANR-17-EURE-0010) and UKRI Frontier Research Grant EP/Z001528/1. 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

[‡]University of Essex. d.garrett@essex.ac.uk

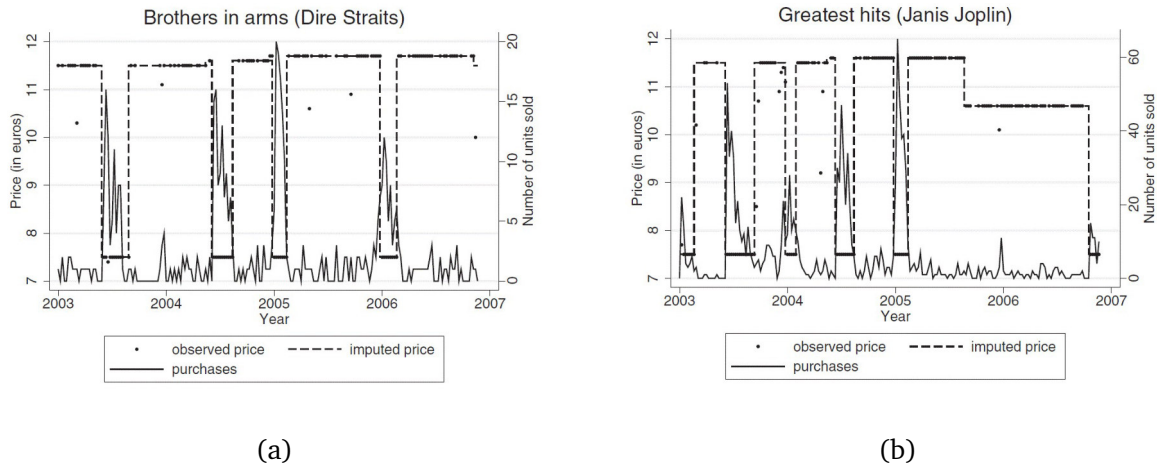


Figure 1: Illustration of typical price and quantity patterns (Figure 1 from Février and Wilner, 2016) for two albums – ‘Brothers in arms’ by Dire Straits, and ‘Greatest hits’ by Janis Joplin – featuring two focal prices (a high price and a discounted or sale price).

as the date of a price discount approaches. In a range of models with flexible prices, this means that the seller gradually reduces prices as the date with the steepest discount draws near. Stokey (1979), Conlisk 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). The latter observation might be interpreted as indicating that buyers are unable to foresee the timing of discounts. Février and Wilner 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

¹While 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.

²Some of these papers are reviewed in Section 7 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 that

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. Aversion to small risks has often been observed. Among the evidence on aversion to small-scale risks is that relating 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). A prominent interpretation is that small-stakes risk aversion is reflective of agent loss aversion and Section 6 discusses how we can adapt our theory when viewing buyers as loss averse.

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 stochastic process for prices we construct is a stationary Markov process in which future prices depend 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 independently of their arrival time because the distribution of the arrival of the next discounting episode 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 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.

commitment “is reasonable with applications such as retailing, online ads, and concerts in which the seller automates the pricing scheme and uses it repeatedly.”

³The 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.

As we explain below, these stochastic price paths 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 show that when buyers only observe prices after their arrival to the market, there is a range of optimal stochastic processes for prices, 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, so a price discount becomes more likely the longer since the last discount. Thus, under an optimal stochastic process for prices, price discounts may be somewhat predictable to an econometrician with access to historical price data, but not to an arriving buyer with no access to previous prices. We argue that this permits the theory to better accommodate observed discount patterns, as several empirical studies find that the arrival of price discounts has an increasing hazard rate. 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 ideas above connect to a range of work on price discounts discussed further in Section 7.

2 Set-up

In this section, we introduce the environment of our static model. Section 3 will study seller-optimal mechanisms in this environment. The dynamic setting is introduced in Section 4.

There is a seller and a buyer. The seller can produce one unit of a good at cost zero and the buyer has unit demand. The buyer's enjoyment of the good depends on his private "type", labeled $\{\theta_n | n=1, \dots, N\}$, with $\theta_N > \dots > \theta_1 > 0$. The probability that the buyer's type is θ_n is $\beta_n > 0$.

An outcome of our model is a production decision and a price in case the good is produced. If the good is not produced, both the buyer and the seller obtain 0. Alternatively, if the buyer's type is θ_n , the good is produced, and the price is $p \in \mathbb{R}_+$, then the seller obtains p and the buyer obtains $v(p; \theta_n) \in \mathbb{R}$. We generally use the alternative notation $v_n(p)$ to denote $v(p; \theta_n)$.

We assume that, for each n , $v_n(\cdot)$ is a strictly decreasing, strictly concave, and twice continuously-differentiable function. We normalize $v_n(\theta_n) = 0$ for each θ_n , which means that types have the interpretation of maximum willingness to pay. We also make the following additional assumptions.

Condition A.

A1 *Condition on price sensitivity:* For any $n=1, \dots, N-1$ 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 $n=1, \dots, N-1$ and $p \in \mathbb{R}_+$, $\frac{v''_{n+1}(p)}{v'_{n+1}(p)} \leq \frac{v''_n(p)}{v'_n(p)}$.

Assumption A1 is a single-crossing condition: it ensures that at any given price, a higher-type buyer has a relatively lower sensitivity to price than a lower-type buyer. Assumption A2 says that higher types are less risk averse. 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.

A natural special case is where we 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 $u: \mathbb{R} \rightarrow \mathbb{R}$. If $v_n(p) = u(\theta_n - p)$, the above restrictions on preferences require that u is strictly increasing, strictly concave, twice continuously differentiable, and satisfies $u(0) = 0$. Assumption A1 is then met because $\frac{u'(y)}{u(y)}$ is strictly decreasing over $y > 0$.⁴ 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 given by specifying CARA preferences, where $u(y) = 1 - e^{-Ry}$ for a coefficient of absolute risk aversion $R > 0$.

3 Analysis of the static model

This section considers static mechanisms for the environment presented in Section 2, where the buyer pays only if he gets the good. This anticipates the relevance for dynamic pricing in Section 4.

By the revelation principle, it will be without loss of generality to consider direct mechanisms. These mechanisms allocate a unit to each type θ_n with probability denoted x_n . In addition, they stipulate a potentially random and non-negative price conditional on assignment, with cumulative distribution function H_n for each type θ_n . To ensure finiteness of expected payoffs, we assume that H_n has a bounded support. 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 $\mathcal{M} = (x_n, H_n)_{n=1}^N$.

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 $\mathcal{M}^D = (x_n, p_n)_{n=1}^N$. The following result implies monotonicity of the allocation in mechanisms with deterministic payments.

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 (shown in the Appendix together with the other proofs) follows from a combination of the fact that higher types have higher maximum willingness to pay and Assumption A1, which limits the price

⁴Assumption A1 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 $\theta > p$.

sensitivity of higher types. The result implies that if some type θ_k would not accept a specified reduction in the probability of getting the good in return for a discount on the price, then any higher type θ_l would not accept this reduction either. Note that the result uses deterministic payments. Our next result is that this is the relevant case for seller-optimal mechanisms, where we use now both Assumptions A1 and A2.

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

Note that Matthews (1983), Maskin and Riley (1984), and Moore (1984) find similar results. However, their settings feature a payment also in case the buyer is not awarded the good, so their results do not apply (especially consider Theorem 1 in Moore, 1984). Another point of comparison is provided by Bansal and Maglaras (2009) in a model where buyers have CRRA utility and only pay if they get the good. However, their work *assumes* each type pays a sure price, rather than deriving the implication.

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 $n = 1, \dots, N$, $x_n^* v_n(p_n^*) = x_{n-1}^* v_n(p_{n-1}^*)$, where we put $x_0^* = p_0^* = 0$.*

While Proposition 1 follows closely from the above lemmas, it is worth drawing attention to the proof of uniqueness, which appears new to the literature. 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 it is important to demonstrate 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 first establish this in the case where $N = 2$ before discussing the case with many types.

In the two-type model, 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 expected profits as:

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

The optimal mechanism is then determined by maximizing the expression in Equation (1) with respect to x_1 .

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.⁵ 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.

Many types. We now show that random allocations can be a robust feature of the optimal mechanism also with a large number of types. We consider sequences of discrete-type models that, in the limit, approximate settings with a continuum of types. We can provide conditions under which the probability of types receiving random allocations (i.e., the probability of types θ_n with $x_n \in (0, 1)$) remains bounded away from zero in this limit.

In the main text, we specialise to the case of CARA utility mentioned above, where $v_n(p) = 1 - e^{-R(\theta_n - p)}$ with $R > 0$ the coefficient of absolute risk aversion. More general preferences are considered in the proofs in the Appendix. We consider type distributions that approximate a twice continuously differentiable distribution F with density f and support on an interval $[\underline{\theta}, \bar{\theta}]$ with $0 \leq \underline{\theta} < \bar{\theta}$.

We impose a regularity condition on F , namely that $p(1 - F(p))$ has a unique maximiser $p^* \in (\underline{\theta}, \bar{\theta})$.⁶ Note that $p(1 - F(p))$ represents the seller's expected profit from posting price p to get the good for sure in a model with types continuously distributed according to F (using that types are equal to the maximum willingness to pay). We then consider a sequence of models with evenly spaced types, where the number of types is given by the increasing sequence $(N_m)_{m=1}^\infty$. For each $m \in \mathbb{N}$, and each $n \in \{1, 2, \dots, N_m\}$, we let $\theta_n^m = \underline{\theta} + (n-1)\frac{\bar{\theta} - \underline{\theta}}{N_m}$ and let $\beta_n^m = F(\theta_{n+1}^m) - F(\theta_n^m)$, where we set $\theta_{N_m+1}^m = \bar{\theta}$. Here, β_n^m is the probability of type θ_n^m .

For each $m \in \mathbb{N}$, we let E^m represent the environment with the above specified types, payoffs, and distribution over types. Taking m large closely approximates a model with a continuum of types distributed according to F . 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.

Proposition 3. *Suppose the buyer is sufficiently risk averse that the following inequality holds:*

$$R > \frac{f'(p^*)p^* + 2f(p^*)}{1 - F(p^*)}. \quad (2)$$

Then there exists $\varepsilon > 0$ and K sufficiently large that, for all $m > K$, the following is true: There exist adjacent

⁵That 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).

⁶The 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) $\lim_{p \uparrow \bar{\theta}} f(p)^{-1}(1 - F(p)) < \bar{\theta}$.




Figure 2: Numerically determined optimal mechanism 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 $n=1, \dots, 100$.

types $\theta_{n'}^m, \theta_{n'+1}^m, \dots, \theta_{n''}^m$ with $x_{n'}^m, x_{n'+1}^m, \dots, x_{n''}^m \in [\varepsilon, 1-\varepsilon]$ and $\theta_{n''}^m - \theta_{n'}^m \geq \varepsilon$.

We thus find that, when the buyer is sufficiently risk averse, the fraction of types that are assigned a random allocation does not vanish as the number of types N_m becomes large. Alternatively stated, the probability of these types (as given by β_n^m) does not vanish. The sufficient condition on the coefficient of absolute risk aversion R depends on the distribution of types. For instance, if F is the uniform distribution on $[0, \bar{\theta}]$, so that each environment E^m has N_m equally likely types, it is enough that $R > 4/\bar{\theta}$. Figure 2 depicts the price and probability of allocation under the optimal mechanism for $N_m=100$, $\underline{\theta}=0$, $\bar{\theta}=10$, and $R=1$.

The key idea of the proof is to show that, when Condition (2) is satisfied, the optimal deterministic mechanism can be slightly perturbed to increase expected profits. The perturbation involves augmenting the deterministic mechanism by including a single additional option to get the good probabilistically. This option is made attractive only to an arbitrarily small set of types.

It is worth noting there are some antecedents to our results. In Liu and van Ryzin (2008), Proposition 9 implies the optimality of random allocations in a model with CRRA utility and a continuum of types. The literature on auctions with risk aversion also finds random allocation; for instance, see Theorem 1 of Matthews (1983). However, we require a separate analysis because we restrict attention to settings where the buyer pays only when receiving the good.

4 Pricing with dynamic arrivals

This section aims at understanding stochastic price paths in the dynamic environment of interest where buyers arrive over time. Time is continuous and the horizon infinite, with time indexed by $t \in [0, \infty)$.

There is a seller and a continuum of buyers. The seller faces no capacity constraints and zero production costs, while each buyer has a unit demand. Both the seller and buyers have a common discount rate

$r > 0$. Buyers arrive to the market at a fixed rate $\gamma > 0$ and we therefore refer to this set-up as the *model with dynamic arrivals*. The parameter γ simply scales demand and hence will scale the realised profits. It is convenient to normalize by setting $\gamma = r$, implying that $\int_0^\infty \gamma e^{-rt} dt = 1$. We can then conveniently understand the seller's expected profits as a per-buyer (weighted) average over the profits from all arrival times $t \geq 0$.⁷ We assume that buyers observe the entire history of past prices upon arrival (Section 5 relaxes this assumption).

As in Section 2, buyers' enjoyment of the good depends on their types, labeled $\{\theta_n | n = 1, \dots, N\}$, with $\theta_N > \dots > \theta_1 > 0$. For each arrival time, a fixed proportion $\beta_n > 0$ has type θ_n , where $\sum_{n=1}^N \beta_n = 1$. Each buyer can transact at most once with the seller; that is, allocation of the good and payment must occur on the same date. If a buyer of type θ_n obtains the good at time t at price p , his payoff in date-zero terms is $e^{-rt} v_n(p)$, where v_n satisfies the same properties as in Section 2. The seller's profit from selling to this buyer (again in date-zero terms) is $e^{-rt} p$.

There is naturally a close connection between the dynamic model specified here and the static model of Section 3. In considering the dynamic model, we sometimes refer to the corresponding static model. This is simply the static model with the same type distribution (i.e., the same set of types and type probabilities β_n), and with the same utility functions $v_n(p)$. We denote the seller's optimal profits in the static model by Π^* . These static profits will be shown to be an upper bound on the expected discounted profits attainable in the dynamic environment. Our main question will be whether and how random price paths can generate expected discounted profits equal or close to this bound.

Price-path mechanisms: We take a mechanism design approach to defining the seller's choice. This means that we specify not only the random price path to which the seller commits, but also view the acceptance decision of buyers as being determined by the seller subject to incentive compatibility. We introduce the notation that \mathcal{B}_A represents the Borel sigma algebra on a set $A \subset \mathbb{R}$, while $\mathcal{A}_1 \otimes \mathcal{A}_2$ is the product sigma algebra associated with two sigma algebras \mathcal{A}_1 and \mathcal{A}_2 . Then we have:

Definition 1. A price-path mechanism is a pair (P, τ) and a filtered probability space $\langle \Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{R}_+}, \mathcal{P} \rangle$ satisfying that:

1. $P = (P_t)_{t \in \mathbb{R}_+}$ is a non-negative valued, bounded, and progressively measurable stochastic process defined on $\langle \Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{R}_+}, \mathcal{P} \rangle$, where the filtration $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ is the one generated by P .
2. $\tau = (\tau_{n,t})_{n \in \{1, 2, \dots, N\}, t \in \mathbb{R}_+}$ is a collection of stopping times. For each pair (n, t) , $\tau_{n,t}$ is a stopping time with respect to the given filtration and satisfying that $\tau_{n,t}(\omega) \geq t$ for all $\omega \in \Omega$. A stopping time may be finite or infinite valued. We impose the additional restriction that, for each θ_n , $\tau_{n,t}(\omega)$ is $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ measurable as a function from $\mathbb{R}_+ \times \Omega$ into the non-negative extended real numbers.

A price-path mechanism prescribes a random price path P , and a buyer acceptance strategy τ , with the acceptance strategy specifying the stopping times $\tau_{n,t}$. The filtered probability space is generally left implicit. A realization of P consists of the path of prices posted by the seller. For each n and t , the realization

⁷While a constant arrival rate is a convenient simplification, all our arguments and results extend also to settings with time-varying arrival rates.

of $\tau_{n,t}$ is the time at which a buyer of type θ_n arriving at time t purchases the good. A realization equal to $+\infty$ indicates that he never purchases the good.

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

$$\mathbb{E}[e^{-r\tau_{n,t}} v_n(P_{\tau_{n,t}})] \geq \mathbb{E}[e^{-r\hat{\tau}_{n,t}} v_n(P_{\hat{\tau}_{n,t}})] \quad (3)$$

for all stopping times $\hat{\tau}_{n,t}$ with respect to the given filtration satisfying that $\hat{\tau}_{n,t}(\omega) \geq t$ for all $\omega \in \Omega$.⁸ The left side of Equation (3) corresponds to the expected payoff of a buyer of type θ_n who arrives at time t and purchases according to the prescribed stopping time $\tau_{n,t}$, while the right side corresponds to his expected payoff if he deviates and purchases according to the stopping time $\hat{\tau}_{n,t}$. The assumption that buyers observe the past prices is encoded by the requirement that the stopping times are with respect to the filtration $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$, which is the filtration generated by P .

We say that the price-path mechanism (P^*, τ^*) is *optimal* if it maximizes the seller's expected discounted profits

$$\mathbb{E} \left[\int_0^\infty \sum_{n=1}^N \beta_n e^{-r\tau_{n,t}} P_{\tau_{n,t}} r dt \right] \quad (4)$$

among all incentive compatible price-path mechanisms (P, τ) .⁹

Analysis of Price-Path Mechanisms. We now present one of the paper's central results, which pertains to our model with dynamic arrivals as described above.

Proposition 4. *Consider the model with dynamic arrivals. The seller's expected discounted profits are no greater than Π^* . For any $\varepsilon > 0$, there is a price-path mechanism such that the seller's expected discounted profits are at least $\Pi^* - \varepsilon$.*

The first part of the result bounds expected discounted profits by the optimal profits in the static model and thus establishes that the dynamic environment does not provide any screening advantage to the seller. Our finding here follows from a mathematical equivalence between allocation probabilities and discount factors. Dynamics introduces the possibility of delayed allocation, but delay has an analogous effect on expected discounted payoffs as a reduction in the allocation probability. This observation is familiar from the literature. For instance, Rochet and Thanassoulis (2019) write (p 953), “the discount factor in the intertemporal pricing problem can be interpreted as a probability of delivery [in a static problem]”. They attribute this idea to Salant (1989).

Now consider the second part of Proposition 4: price-path mechanisms can at least approach, or possibly exactly attain, the seller expected discounted profits Π^* . In fact, whether expected discounted profits of Π^* can be attained turns on the number of distinct prices that occur in the optimal mechanism of the corresponding static problem. Let J denote the number of distinct prices at which purchase occurs with

⁸Throughout, we define $e^{-\infty}$, $e^{-\infty} P_{+\infty}(\omega)$, and $e^{-\infty} v_n(P_{+\infty}(\omega))$ to be zero for all $\omega \in \Omega$. Measurability of the integrands on each side of Equation (3) is established in the proof of Proposition 4.

⁹Measurability of the integrand in Equation (4) is established in the proof of Proposition 4.

positive probability in the optimal static mechanism; equivalently, this is the number of distinct positive values of x_n^* . For instance, $J=1$ indicates that the optimal static mechanism is deterministic. For the two-type model of Proposition 2, for parameters such that allocation to the low type is random, we have $J=2$. Expected discounted profits of Π^* can be exactly attained when J is equal to 1 or 2, but – as we show in the Online Appendix – not exactly attained when $J \geq 3$.

To start understanding the argument behind the second part of Proposition 4, 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 $x_{n_{j-1}}^* < x_{n_j}^*$. This means that, for each $j=1, \dots, J$, type θ_{n_j} is the smallest type receiving the good with the positive probability $x_{n_j}^*$ (the price paid when purchasing with this probability is $p_{n_j}^*$). Hence, for all $j \in \{1, \dots, J\}$ and $n \in \{n_j, \dots, n_{j+1}-1\}$, we have $p_n^* = p_{n_j}^*$ and $x_n^* = x_{n_j}^*$, where $n_{J+1} := N+1$.

Now, consider the case of $J=1$. The seller can attain expected discounted profits Π^* in the model with dynamic arrivals through a constant price path with price equal to $p_{n_1}^* = \theta_{n_1}$. All buyers with type weakly above θ_{n_1} buy upon arrival, and all buyers with type strictly below θ_{n_1} never buy. Considering the date- t value of the payoffs associated with buyers arriving at date t , this replicates the outcome in the static problem. For instance, buyers with types $\theta_n \geq \theta_{n_1}$ earn payoff $v_n(p_{n_1}^*)$ while generating profit of $p_{n_1}^*$ for the seller. We can note that the absence of delayed purchase in the dynamic model corresponds to the sure allocation in the static model. The optimality of a constant price path here is familiar, for instance, from Conlisk et al. (1984) who show in a setting with risk-neutral buyers arriving over time that the seller optimally commits to a constant price path.

If $J=2$, the seller can attain expected discounted profits Π^* with a stochastic process for prices (or price process) consisting of a constant “high” price $p_{n_2}^*$, punctuated by random reductions to price $p_{n_1}^* = \theta_{n_1}$ arriving at a constant Poisson rate $\lambda_{n_1} = r x_{n_1}^* / (1 - x_{n_1}^*)$ and which last only an instant before the resumption of the high price. In this case, all buyers with type weakly above θ_{n_2} buy upon arrival at the high price, all buyers with type strictly below θ_{n_2} and weakly above θ_{n_1} wait and buy at the reduced price, and all buyers with type strictly below θ_{n_1} never buy. Note that the arrival rate of price reductions λ_{n_1} is calculated so that, when the price is high, the expected discount factor associated with the time to the next price reduction is $x_{n_1}^*$ (i.e., we have $\lambda_{n_1} / (r + \lambda_{n_1}) = x_{n_1}^*$). Incentive compatibility of the stopping times described above follows from incentive compatibility of the static mechanism. For instance, for any type $\theta_n \geq \theta_{n_2}$, we have from the static incentive compatibility constraints that $v_n(p_{n_2}^*) \geq x_{n_1}^* v_n(p_{n_1}^*)$, which ensures θ_n prefers to purchase immediately at the high price rather than wait for the arrival of the reduced price.

Now, consider the case where $J \geq 3$. Our approach to show the approximation of expected discounted profits Π^* is, in essence, to spread episodes of sequential discounts over intervals of short duration. We construct price-path mechanisms where the price process is a stationary Markov process with J states. In each state j , the price is $p_{n_j}^*$. In state J , the state transitions to state $J-1$ at rate $\lambda_{n_{J-1}} = r x_{n_{J-1}}^* / (1 - x_{n_{J-1}}^*)$ (note the similarity with the case $J=2$). In each state j , with $2 \leq j \leq J-1$, the state transitions at a high rate to either state $j-1$ or to state J . Thus, when the state is in state j with $2 \leq j \leq J-1$, the state quickly transitions to either state J (and the sequential discount episode ends) or to state $j-1$ with a more discounted price. In state $j=1$, the state transitions at a high rate to state J .

The proof of Proposition 4 shows that the transition rates can be constructed so that, conditional on state J , the expected discount factor associated with stopping only once reaching state j is $x_{n_j}^*$. We can then show that, as a consequence of incentive compatibility in the static mechanism, each type θ_n finds it optimal to follow an acceptance strategy that specifies a highest purchase price equal to p_n^* . Expected discounted profits from buyers of type θ_n arriving in state J are then $x_n^* p_n^*$, the same as the expected profits in the optimal static mechanism. Then, by choosing high enough transition rates, the seller's expected discounted profits are close to the static profits Π^* . Expected discounted profits are less than Π^* only because, for some values of n , buyers of type θ_n who arrive during a discounting episode (i.e., when the state is not J) may purchase at a price lower than p_n^* .

It is worth commenting on the price patterns generated by the price processes that approximate profits Π^* as the transition rates between states increase. As we increase the transition rates out of state $J-1$ and below, the transition rate from J to $J-1$ is held fixed at $r x_{n_{J-1}}^* / (1 - x_{n_{J-1}}^*)$. Prices therefore hold in state J at $p_{n_j}^*$ except for occasionally being discounted, at first instance, to $p_{n_{J-1}}^*$. But once a price reduction has occurred, further price reductions follow in quick succession over a short period, with the episode of price discounting terminating randomly as the process returns to state J . Note that a buyer of type θ_{n_j} , $1 < j < J$, finds it optimal to purchase when the price reaches $p_{n_j}^*$ rather than waiting for further discounts even though these may occur very soon. The reason the buyer optimally chooses not to wait is the threat of the price returning to the non-discounted price $p_{n_j}^*$.

The price patterns exhibited by the process can be described in lay terminology. It is common to refer to a retailer's "regular price" as the price before any discount has been applied. Our price processes can then naturally be thought of as consisting of a regular price $p_{n_j}^*$, interrupted by short-lived episodes of price discounting, where the depth of the price discount reached in each episode is uncertain. Figure 3(a) provides an example of a sample price path with $N=3$ types and $0 < x_1^* < x_2^* < x_3^* = 1$, where the regular price is p_3^* . Figure 3(b) provides an example with $N=100$ types and parameters as in Figure 2. Note that in the price paths in the figures, there are several discounting episodes, each of them involving a price decline that stops at a random time.

We have demonstrated that price-path mechanisms can either attain or approach the static profits Π^* in the model with dynamic arrivals. We can show that randomness in the price path is necessary for this end whenever $J \geq 2$. Towards this goal, we consider *deterministic price-path mechanisms*, with the natural interpretation that the price path $(P_t)_{t \in \mathbb{R}_+}$ that the seller commits to is deterministic. Formally, a pair (P, τ) is a deterministic price-path mechanism if there is some bounded function $p: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $P_t(\omega) = p_t$ for all $\omega \in \Omega$.

Proposition 5. *Consider the model with dynamic arrivals and suppose the optimal static mechanism involves random allocation, i.e. $J \geq 2$. There exists $\varepsilon > 0$ such that no deterministic price-path mechanism achieves seller expected discounted profits greater than $\Pi^* - \varepsilon$.*

The reason for this result can be understood by comparing outcomes in the dynamic setting with those in the static model. For a deterministic price-path mechanism to achieve discounted profits Π^* for some cohorts, the seller must drop the price to $p_{n_1}^* = \theta_{n_1}$, say at some date $t^* > 0$. But then higher buyer types

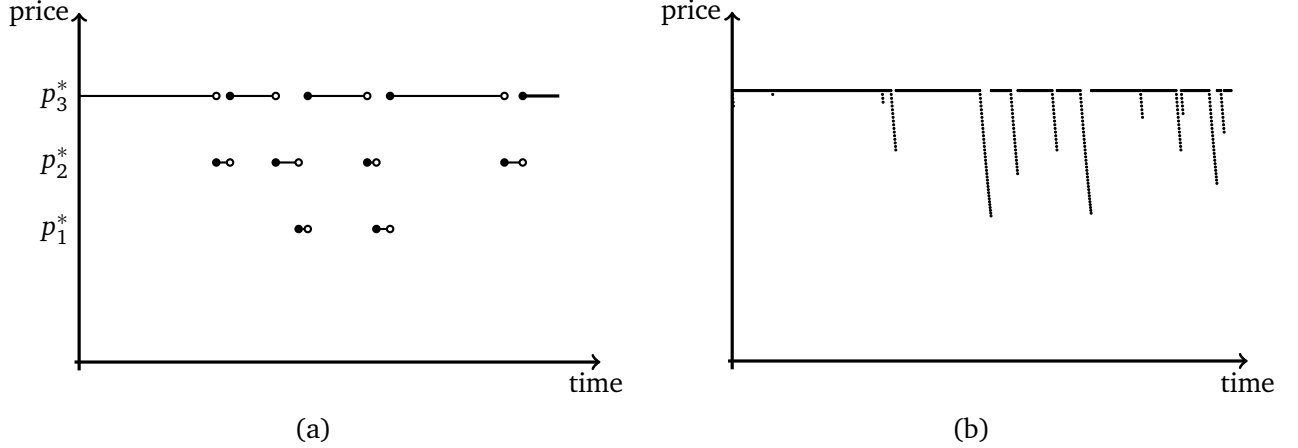


Figure 3: (a) Example of a realization of a price path for $N=3$ when $0 < x_1^* < x_2^* < x_3^* = 1$. In the example, there are four discount episodes from the regular price p_3^* . In the first and the last, only p_2^* is offered before p_3^* is offered again. In the other two discount episodes a second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

entering the market shortly before have the opportunity to purchase at a low price with little delay. The outcome for these buyers is a lower purchase price than in the unique optimal static mechanism, and this represents a loss in expected discounted profits.

We have concluded that, with $J \geq 2$, there is a loss in discounted profits from a restriction to deterministic price-path mechanisms. This loss turns out to be the result of dynamic arrivals. To see this, suppose a unit mass of buyers arrive to the market at a fixed date, say zero, and no buyers arrive thereafter. For each j with $1 \leq j \leq J$, the seller can set a price $p_{n_j}^*$ at a date $t_{n_j}^*$ satisfying $e^{-rt_{n_j}^*} = x_{n_j}^*$. Prices at other times can be set at θ_N (so no buyer has a strict preference to purchase at these prices). Then a buyer who purchases at price $p_{n_j}^*$ in the optimal static mechanism can be asked to purchase at the same price $p_{n_j}^*$ under the price path, and the discount factor associated with this purchase is $e^{-rt_{n_j}^*} = x_{n_j}^*$. This is enough to conclude that the seller's discounted profits are Π^* .

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 this section, we examine what our theory implies for the timing of discounts. We do so by focusing on the case where $N=2$ and where $\beta_2 \in (\underline{\beta}, \bar{\beta})$, the interval identified in Proposition 2 (hence, the probability of allocation to the θ_1 -buyer in the optimal static mechanism, x_1^* , is interior). Our findings extend readily to the case of $N > 2$ but $J = 2$.

As explained in Section 4, an optimal price-path mechanism exists when $N=2$, generating expected discounted profits Π^* . As demonstrated in the proof of Proposition 4, attaining expected profits of Π^*

requires that for almost all cohorts t (i) high types θ_2 purchase immediately and at a price p_2^* almost surely, and (ii) low types θ_1 purchase almost surely at a price $p_1^* = \theta_1$, and with delay to purchase such that the associated expected discount factor is $\mathbb{E}[e^{-r(\tau_{1,t}-t)}] = x_1^*$.

Motivated by these observations and to avoid technicalities, we focus throughout this section on price-path mechanisms with the following *Regularity Conditions*. First, sample paths of the price process $(P_t)_{t \in \mathbb{R}_+}$ are constant at p_2^* except at isolated points (“sales”) when the price is $p_1^* = \theta_1$. Also, the stopping times for high types are given by $\tau_{2,t}(\omega) = t$ for all t and all ω . To specify stopping times for low types, we assume for each ω and each t that either (i) there is a *minimum* date $s \geq t$ such that the price $P_s(\omega)$ is equal to θ_1 , or (ii) the price is never θ_1 after t . In the first case $\tau_{n,t}(\omega)$ is equal to the minimum date, and in the second $\tau_{n,t}(\omega) = +\infty$. It then makes sense to describe the low type θ_1 as purchasing on the *next* sales date if there are any future sales.

Optimal price-path mechanisms satisfying the Regularity Conditions exist: the case where sales arrive according to a Poisson process with parameter $\lambda_1 = rx_1^*/(1-x_1^*)$ is considered above and in the proof of Proposition 4. When sales are optimally determined according to a Poisson process, at any t and irrespective of the history of past prices, the time until the next sale at date $s > t$ is exponentially distributed with parameter λ_1 . Thus, buyers are not able to use the price history to better predict future sales. We can show that this property is in fact shared by any optimal price-path mechanism satisfying the Regularity Conditions in the following sense.

Proposition 6. *Consider the model with dynamic arrivals and number of types $N = 2$, and with $x_1^* \in (0, 1)$. Fix any optimal price-path mechanism satisfying the Regularity Conditions. For any date $t > 0$, any positive probability event $A \in \mathcal{F}_t$, the distribution of the time to the next sale is exponential with parameter λ_1 conditional on A . That is, for any $s \geq t$, $\mathcal{P}(\tau_{1,t}(\omega) \leq s | A) = 1 - e^{-\lambda_1(s-t)}$.*

An explanation for this result is that maximising profit from (almost) every cohort of buyers requires treating buyers who arrive at different times, or at different price histories, symmetrically. The exponential distribution achieves this as the only distribution with the memoryless property ensuring that the distribution of the waiting time to the next sale is the same irrespective of the price history.

The analysis to this point views buyers as observing the entire history of past prices: purchase times are stopping times with respect to the filtration generated by the entire price history. We now argue that the findings are quite different to the one in Proposition 6 if buyers are unable to observe prices before their arrival to the market. This is arguably realistic and allows us to better reconcile empirical price patterns as we argue below.¹⁰

To capture the limited observability of prices, price-path mechanisms (P, τ) are now modified by changing the filtration for the stopping times τ . In particular, each $\tau_{n,t}$ is now a stopping time with respect to

¹⁰It is worth comparing our analysis to Öry (2017), who considers a model where buyers do not observe prices before or after their arrival and also are uncertain about their arrival time.

the filtration generated by $P^{\geq t} = (P_s^{\geq t})_{s \in \mathbb{R}_+}$ defined by

$$P_s^{\geq t} = \begin{cases} 0 & \text{if } s < t, \\ P_s & \text{if } s \geq t. \end{cases}$$

(We still require that $\tau_{n,t}(\omega) \geq t$ for all $\omega \in \Omega$ as well as $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ measurability of $\tau_{n,t}(\omega)$.) We refer to these mechanisms as *price-path mechanisms with unobserved pre-arrival prices*. Note that now, (P, τ) is *incentive compatible* if, for all n and all t , condition (3) holds for all stopping times $\hat{\tau}_{n,t}$ with respect to the filtration generated by $P^{\geq t}$ and satisfying that $\hat{\tau}_{n,t}(\omega) \geq t$ for all $\omega \in \Omega$.

An initial observation is that it remains optimal for the seller to determine sales according to a Poisson process with rate λ_1 . Specifying that high types stop immediately and low types stop at the next sale date, as described above, satisfies the new restrictions on stopping times. The seller obtains what is still an upper bound on expected discounted profits, Π^* .

A range of patterns of sales is now consistent with optimality. The reason is that, because buyers cannot condition their deviations on prices before arrival, the incentive compatibility constraints are easier to satisfy. To understand these patterns, we consider price-path mechanisms with unobserved pre-arrival prices satisfying the Regularity Conditions as well as the requirement that $\mathbb{E}[e^{-r(\tau_{1,t}-t)}] = x_1^*$ for all t . This last requirement is necessary for obtaining expected discounted profits of Π^* from the cohort arriving at t . We show the following result that provides a condition for incentive compatibility and thus defines a class of price-path mechanisms with unobserved pre-arrival prices that achieve optimality.

Proposition 7. *Consider the model with dynamic arrivals and number of types $N=2$, and with $x_1^* \in (0, 1)$. Consider a price-path mechanism with unobserved pre-arrival prices subject to the Regularity Conditions, and satisfying that $\mathbb{E}[e^{-r(\tau_{1,t}-t)}] = x_1^*$ for all t . Such a price-path mechanism is incentive compatible if and only if, for all t and all $s > t$ such that $\tau_{1,t} > s$ with positive probability,*

$$\mathbb{E}[e^{-r(\tau_{1,t}-s)} | \tau_{1,t} > s] \geq x_1^*. \quad (5)$$

To understand Condition (5), consider the incentive compatibility condition for high types (incentive compatibility for low types is straightforward). Note that the assumption $\mathbb{E}[e^{-r(\tau_{1,t}-t)}] = x_1^*$ implies that a high type arriving at date t expects the same payoff either from purchasing immediately or waiting and purchasing at the next sale. That is, there is no profitable deviation for high types to simply waiting and purchasing at the next sale. Next, because high types always find it optimal to purchase at sales, only one other kind of deviation needs to be checked. This is where the buyer waits and purchases after a fixed amount of delay, or at the next sale if one comes earlier. Such strategies, however, do not generate higher

payoffs under Condition (5). Indeed, for any t and $s > t$ such that $\Pr(\tau_{1,t} > s) > 0$,

$$\begin{aligned} v_2(p_2^*) &= \Pr(\tau_{1,t} \in [t, s]) \mathbb{E}[e^{-r(\tau_{1,t}-t)} | \tau_{1,t} \in [t, s]] v_2(\theta_1) \\ &\quad + \Pr(\tau_{1,t} > s) e^{-r(s-t)} \mathbb{E}[e^{-r(\tau_{1,t}-s)} | \tau_{1,t} > s] v_2(\theta_1) \\ &\geq \Pr(\tau_{1,t} \in [t, s]) \mathbb{E}[e^{-r(\tau_{1,t}-t)} | \tau_{1,t} \in [t, s]] v_2(\theta_1) + e^{-r(s-t)} \Pr(\tau_{1,t} > s) v_2(p_2^*). \end{aligned}$$

The equality follows from the condition that $v_2(p_2^*) = x_1^* v_2(\theta_1) = \mathbb{E}[e^{-r(\tau_{1,t}-t)}] v_2(\theta_1)$. The inequality follows from the fact that Condition (5) holds. Because the expression on the right side of the inequality is the payoff the high type obtains by deviating by purchasing at the first sales date in $[t, s)$ or at s if there is no discount in $[t, s)$, such deviation is not profitable for the buyer. If, instead, Condition (5) does not hold, the inequality is reversed, and then the price-path mechanism with unobserved pre-arrival prices is not incentive compatible.

To further interpret Condition (5), suppose it is satisfied for a buyer arriving at a given date t . If he delays purchase until $s > t$ and observes no sale (i.e., in case of the event $\{\tau_{1,t} > s\}$), then this absence of a sale is effectively “good news” in that he expects the next sale relatively sooner at date s than at date t (as measured in terms of expected discounting). Hence, the condition accommodates discount patterns where the hazard rates for sales are increasing with time since the last sale, as found in the empirical work mentioned above. That is, we can anticipate a rich multiplicity of optimal discount patterns, some of which may coincide with the ones found in practice.

To further illustrate, consider the following example where the time between sales is constant. Suppose that the first sale is uniformly distributed on $[0, \Delta]$ and then sales occur a fixed time Δ apart, where $\Delta > 0$ is determined below. Buyers arriving at any date t with no information on past prices then believe the next sale is uniformly distributed on $[t, t + \Delta]$. The appropriate choice of Δ satisfies

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

It is easy to check that Condition (5) is satisfied: i.e., a buyer who delays purchase becomes more optimistic about the wait time to the next sale in the relevant sense. This price process is then part of an optimal price-path mechanism with unobserved pre-arrival prices.¹¹

6 Loss aversion

We have used expected utility theory to capture the trade-offs faced by consumers averse to risk and therefore the problem of the seller. In doing so, we have been able to demonstrate random discounts in a dynamic pricing model while remaining close to existing literature on auctions (e.g., Matthews, 1983, Maskin and Riley, 1984, and Moore, 1984).

In this section, we show that our results are robust to replacing expected utility by loss aversion.

¹¹Stopping times are as described above. The measurability conditions can be verified by similar arguments as for the Poisson price-path mechanism in the proof of Proposition 4.

Our motivation for doing so is that the stakes involved in individual markets are often small. Rabin (2000) shows how 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. This observation, and the fact that aversion to small risks is often observed, is sometimes viewed as favoring loss aversion instead of expected utility.¹² While we concur with the view that expected utility models with risk-averse agents are illuminative even in settings with small risk (as these models capture economic forces in a simple and canonical manner; e.g., see Rubinstein, 2002), that our theory can be couched in a model of loss aversion is reassuring.

Model of loss aversion. We now present a model that adapts ideas in work such as Kahneman and Tversky (1979) and Kőszegi and Rabin (2006) to our economic environment. We first consider the static model with one seller and one buyer. Our starting point is to posit gain-loss utility for each type of buyer $\theta_n > 0$ 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

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

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 $\lambda > 1$. That $\lambda > 1$ captures the idea that losses, i.e. paying a price about the reference point, are experienced more intensely than gains. We specify the payoff in case not purchasing to be zero; as in the earlier models, the buyer makes no payment if not receiving the good. The seller has no cost of providing the good.

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; we later argue that these choices can reflect consumer expectations about market prices. We find an optimal static mechanism with random allocation to type θ_1 whenever parameters satisfy

$$\beta_2 \in \left[\frac{\theta_1(1+\eta)}{\theta_2 + \eta\rho_2}, \frac{\theta_1(1+\eta+\eta(\lambda-1))}{\theta_2 + \eta\rho_2 + \theta_1\eta(\lambda-1)} \right] \quad (6)$$

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

Proposition 8. *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 (6) 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 prices paid by each type are equal to their reference points.*

To understand the form of the optimal mechanism described in Proposition 8, suppose that Condition (6) holds and the seller offers an optimal direct mechanism as described in Proposition 8. Let the probability

¹²Evidence from experimental settings includes Read et al. (1999), Fehr and Goette (2007), and Gächter et al. (2022).

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 (6)), 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.

Dynamic analysis with an inflow of buyers. Now consider a dynamic environment where buyers arrive over time. Here, we view $v(p; \theta_n, \rho_n)$ as type θ_n 's instantaneous utility of a purchase at price p , so that intertemporal payoffs from a purchase after delay t are given by $e^{-rt} v(p; \theta_n, \rho_n)$ where $r > 0$ is the common discount rate. We assume a constant inflow of buyers at rate r . Proportion β_2 of arrivals are high types, with the remainder low types. Reference points are fixed at $\rho_2 \in (\theta_1, \theta_2)$ and θ_1 respectively, and also therefore do not depend on the arrival time to the market. The seller has no costs and no capacity constraints.

Suppose that parameters satisfy the condition in Equation (6). We can define price-path mechanisms analogously to in Section 4 to be a price process and stopping times subject to incentive compatibility constraints. There is then an optimal price-path mechanism in which the seller sets a constant price ρ_2 , while a Poisson process determines 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 the expected discounting until the price is θ_1 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 .

Reference point determination. For the parameter values where the condition in Equation (6) is satisfied, we have seen that there is an optimal static mechanism and also an optimal price-path mechanism in the dynamic environment where both types of buyer purchase at prices equal to their reference points. If reference points are formed by anticipating (as "lagged rational expectations" in the terminology of Heidhues and Kőszegi (2008)) the prices to be paid when purchasing in equilibrium, then there is a natural sense in which reference points can be self-fulfilling. Suppose that high types anticipate paying $\rho_2 \in (\theta_1, \theta_2)$ in case of a purchase, while low types anticipate paying θ_1 in case of a purchase, and that these expectations determine buyers' reference points.¹³ Then the seller, taking these reference points as given, determines an optimal mechanism (in either of the static or dynamic environments described above) that induces each type to obtain the good paying prices equal to the reference points as shown above. This idea is closely related to the concept of "market equilibrium" introduced by Heidhues and Kőszegi (2008). In the

¹³In the static model above, the low type purchases with probability less than one. Assume the reference point is based on the price to be paid *conditional on purchase*.

Online Appendix, we explain in more detail the sense in which reference points are determined in a market equilibrium of the static model.

7 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 own. 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 are generally developed in a static framework, and rely on the assumption that not all consumers in the market can access all competing price offers on the same terms (especially, due to search frictions). The models can be extended to dynamic settings by considering non-durable goods sold in every period with sellers redrawing prices in each period. Varian seems to see this idea as integral to his theory, writing for example (p. 651) “because of intentional fluctuations in price, consumers cannot learn by experience about stores that consistently have low prices”.¹⁴

One reason it can be difficult to reconcile these theories with price data is that equilibrium prices are drawn from a continuous distribution. In practice, it is often the case that retail prices spend most of the time at their highest level, often termed the “regular price”, and are then only occasionally discounted (see the references in the first paragraph of the Introduction). In this sense, the models of price fluctuations based on competition and search frictions fail to explain the “mass point” in prices at the regular price.

An important related point is that extending the static models with search frictions to dynamic settings by assuming prices are redrawn in each period (as seems to be suggested by Varian, 1980) means that the amount of price variation over a given interval depends crucially on the *frequency* with which prices are redrawn. This approach does not provide a way to determine how often prices are redrawn and therefore has little to say about the number of occasions on which the lowest prices are offered in a given time interval.¹⁵ In contrast, we have provided clear predictions on the frequency of price discounting. For instance, see Section 5.

The popularity of theories of price discounts based on competition and search frictions can, however, be partly explained by the observation that many of the relevant markets have several sellers. What the ingredients of our theory imply when introduced to settings with more than one seller is likely to depend on the details of the model. For instance, we would expect the forces examined in our paper to remain relevant with many sellers in the presence of collusion. Another possibility is that search frictions are sufficiently strong that most consumers consider only one among many sellers, or perhaps one seller at a time subject to infrequent switching. The predictions of such a model could well be similar to the ones in our paper. The fact that our theory only requires a monopolist seller could also perhaps be viewed as broadening its

¹⁴Note 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.

¹⁵A similar point is made by Myatt and Ronayne (2019) who argue that the usual models of price discounts based on competition and search produce more price variation than encountered in real-world markets. They propose a variation of the model that addresses this.

applicability (e.g., to explain price discounts that occur even in markets with little competition).

Another theory of random price discounts that does not involve competition is Heidhues and Kőszegi (2014). This considers a monopoly model with loss averse consumers, where loss aversion is specified in both the payment dimension and the quantity dimension. This theory captures the empirical regularity of a mass point in the distribution at the highest price, and also a gap below the highest price. Again this model is static and extending it to make dynamic predictions about sales seems to require a view on how often prices are redrawn.

As described in the Introduction, our theory is closest to work on intertemporal price discrimination which has presented a range of reasons why prices may fluctuate. 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). None of these papers, however, point to the seller profiting from random price discounting. Some of the work contains arguments why randomisation is not a feature of the optimal price path. In Garrett and Board, the optimality of deterministic pricing can be seen from the maximisation of virtual surpluses which are linear in the probability of allocation.

As explained in the Introduction, because the dynamic pricing models with intertemporal price discrimination do not predict random discounting, prices fall gradually when prices are lowered to sell to low value buyers. This is not the case, however, if high-valuation buyers are simply myopic. In this case, sellers occasionally reduce prices to sell to low-value consumers and can do so without affecting the willingness of high-value consumers to purchase at other times. While this idea can be expressed in a monopoly setting, its first expression appears to be Sobel (1984) in a model with competition. A difficulty with this theory is that the myopia assumption can seem too strong in many markets; for evidence that buyers are forward-looking see Chevalier and Goolsbee (2009) and Février and Wilner (2016).

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, while Bansal and Maglaras consider a related problem. Neither work considers models with dynamic arrival of buyers.

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.
- BEICHELT, F. (2016): *Applied probability and stochastic processes*, CRC Press.
- 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.
- BILLINGSLEY, P. (1995): *Probability and measure*, John Wiley & Sons.
- 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.
- 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.
- CONLISK, J., E. GERSTNER, AND J. SOBEL (1984): “Cyclic pricing by a durable goods monopolist,” *The Quarterly Journal of Economics*, 99, 489–505.
- 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.
- 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. KÓSZEGI (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.
- 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.
- KLENKE, A. (2013): *Probability theory: a comprehensive course*, Springer Science & Business Media.
- KÓSZEGI, 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.
- 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.
- 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.
- ROCHET, J.-C. AND J. THANASSOULIS (2019): “Intertemporal price discrimination with two products,” *The RAND Journal of Economics*, 50, 951–973.
- 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*.
- SALANT, S. W. (1989): "When is inducing self-selection suboptimal for a monopolist?" *The Quarterly Journal of Economics*, 104, 391–397.
- SEPPÄLÄINEN, T. (2012): "Basics of Stochastic Analysis," .
- 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.
- STOKEY, N. L. (1979): "Intertemporal price discrimination," *The Quarterly Journal of Economics*, 93, 355–371.
- 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.

Proofs of the formal results

Proof of Lemma 1

Proof. Let k and l satisfy the conditions in the statement. Suppose that $x''v_k(p'') \geq x'v_k(p')$. Note that, using that higher types have higher willingness to pay, $v_l(p'') > 0$. Hence, $x''v_l(p'') > x'v_l(p')$ is immediate if $p'' \leq p'$ or if $x' = 0$. Hence, we may assume $x' > 0$ and $p' < p'' < \theta_k$. Then,

$$\begin{aligned} x''v_l(p'') &\geq x'v_l(p') \frac{v_k(p')}{v_k(p'')} \frac{v_l(p'')}{v_l(p')} \\ &= x'v_l(p') e^{\int_{p'}^{p''} \left(-\frac{v'_k(p)}{v_k(p)} - \left(-\frac{v'_l(p)}{v_l(p)} \right) \right) dp} \\ &> x'v_l(p'), \end{aligned} \tag{7}$$

where the strict inequality follows from Assumption A1, and because $v_l(p') > 0$ (since $v_l(p'') > 0$ and $p' < p''$). \square

Proof of Lemma 2

Proof. Consider an arbitrary (incentive compatible and individually rational) mechanism $\mathcal{M} = (x_n, H_n)_{n=1}^N$ 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 $\mathcal{M}^D = (x_n, p_n)_{n=1}^N$. Note that (by Jensen's inequality and strict concavity of each v_n) this is strictly more profitable than the original mechanism \mathcal{M} if the buyer reports the truth. 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 \mathcal{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 \mathcal{M}^D . We can construct a set of types \mathcal{J} of cardinality $J \equiv |\mathcal{J}|$ along which expected profit strictly increases. We begin by letting θ_{n_1} be the lowest type assigned the good with strictly positive probability in \mathcal{M}^D . Then, having determined θ_{n_j} , we let $\theta_{n_{j+1}}$ be the next smallest type such that profits exceed those for θ_{n_j} . That is, for each $j \geq 1$, we let $n_{j+1} = \min\{n : n > n_j, \pi_n > \pi_{n_j}\}$ if the set is non-empty, and stop otherwise so that $J = j$. This determines $\mathcal{J} = \{\theta_{n_j} : j = 1, \dots, J\}$. We then denote \mathcal{M}^R the "restricted" indirect mechanism which is the same as \mathcal{M}^D except that the buyer is permitted to choose only among messages in \mathcal{J} .

Consider now the reporting decision of any type θ_{n_j} in \mathcal{M}^R , with $\theta_{n_j} \in \mathcal{J}$. Because the original mechanism \mathcal{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 \mathcal{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' .¹⁶

¹⁶Note 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, $x_{n_j} > x_{n_{j'}}$.

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

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.¹⁷ 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 (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 \tilde{p} where type θ_k is indifferent between (x, \tilde{p}) and (x', p') , type θ_k then prefers (x, \tilde{p}) to any smaller allocation probability. Again by Lemma 1, any type higher than θ_k then strictly prefers (x, \tilde{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 \tilde{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

¹⁷The buyer can also choose not to participate, which means zero allocation probability and zero payment.

to the previous case, we consider raising this price to $\tilde{p} = \theta_k$. 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 $n=1, \dots, N$, $x_n^* v_n(p_n^*) = x_{n-1}^* v_n(p_{n-1}^*)$, 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 $\{(x_1, \dots, x_N) \in [0, 1]^N : x_1 \leq \dots \leq x_N\}$.

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 $p_{\underline{n}^A}^A = \theta_{\underline{n}^A} > p_{\underline{n}^A}^B$, contradicting our previous assumption which implies $p_{\underline{n}^A}^A = p_{\underline{n}^A}^B$. Given that a contradiction can also be reached for the case $\underline{n}^A < \underline{n}^B$, it must be that $\underline{n}^A = \underline{n}^B = \underline{n}$. Since downward constraints bind in both mechanisms, for all $n > \underline{n}$, we have

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

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

$$\frac{x_n^A}{x_n^B} = \frac{x_{n-1}^A}{x_{n-1}^B}.$$

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

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

p_n^B . Incentive compatibility of the new mechanism is equivalent to the requirement that, for all n, k ,

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

or

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

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

Proof of Proposition 2

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

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

Hence, strict concavity of the expression implies that $x_1^* = 0$ if and only if $\beta_2 \geq \bar{\beta} \equiv \frac{\theta_1}{\theta_1 - \frac{v_2(\theta_1)}{v_2'(\theta_2)}} > \frac{\theta_1}{\theta_2}$.¹⁸ 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 $\beta_2 \leq \underline{\beta} \equiv \frac{\theta_1}{\theta_1 - \frac{v_2(\theta_1)}{v_2'(\theta_1)}} < \frac{\theta_1}{\theta_2}$.¹⁹ It is then necessarily the case that $x_1^* \in (0, 1)$ if and only if $\beta_2 \in (\underline{\beta}, \bar{\beta})$. \square

Proof of Proposition 3

Proof. The proof plan is the following. We first construct a model with a continuum of types, referred to as the ‘‘continuum model’’. Such a model can be seen as the limit of a sequence the models defined in the main text as $N_m \rightarrow \infty$. We first establish a sufficient condition for a deterministic mechanism not to be optimal in the continuum model (Proposition 9). We then establish that, under the sufficient condition, not only the optimal mechanism is stochastic for N_m large enough, but that such stochasticity does not vanish as $N_m \rightarrow \infty$ (Proposition 10).

Continuum model. Our proof approach is to first provide conditions under which deterministic mechanisms are less profitable than mechanisms with random allocations in environments with a continuum of types, and then extend the finding to the nearby discrete-type models.

¹⁸The 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 $v_2(\theta_2) = 0$.

¹⁹The 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 $v_2(\theta_2) = 0$.

It comes at little cost to generalise preferences in the main text: here we denote the payoff of type $\theta \in [\underline{\theta}, \bar{\theta}]$ who acquires the good paying $p \geq 0$ by $v(p; \theta)$. As mentioned in the main text, type θ is the buyer's willingness to pay (i.e., $v(\theta; \theta) = 0$). 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. We denote by $v'(p; \theta)$ and $v''(p; \theta)$ the first and second derivatives with respect to p . We then 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. This additional continuity will be used to extend our finding for a continuum-type model to the discrete-type models of interest. The CARA utility specification introduced in the main text (where $v(p; \theta) = 1 - e^{-R(\theta - p)}$ with $R > 0$) satisfies all these restrictions.

The distribution F together with buyer preferences $v(p; \theta)$ define a continuum-of-types model. Under the regularity condition on F in the main text, expected profits in the deterministic mechanism are uniquely optimised by posting a price (for sure acquisition of the good) interior to the type space. We now refer to this price as θ^* (equivalently, θ^* is the ‘‘marginal type’’ in the optimal deterministic mechanism for a continuum of types). We show the following.

Proposition 9. *Consider the continuum-type static model and a unique optimal posted price $\theta^* \in (\underline{\theta}, \bar{\theta})$. Then there is a mechanism with a random allocation that is more profitable than the optimal deterministic mechanism provided that*

$$\frac{v''(\theta^*; \theta^*)}{v'(\theta^*; \theta^*)} > \frac{f'(\theta^*)\theta^* + 2f(\theta^*)}{1 - F(\theta^*)}. \quad (8)$$

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

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

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 $\varepsilon = 0$. Considering Equation (9) and the implicit function theorem, we see that $p(\cdot)$ is differentiable and

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

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

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

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

Next, consider the second derivative:

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

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

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

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 (8). □

(End of the proof of Proposition 9. The proof of Proposition 3 continues.)

Discrete approximation. We now extend the finding in Proposition 9 to discrete-type approximations of the continuum model. These approximations feature discrete type distributions as introduced in the main text. Paired with the utility functions $v(p; \theta)$, these sequences of type distributions define a sequence of environments $(E^m)_{m=1}^\infty$ as in the main text. The (unique) optimal mechanism in environment E^m is denoted $(x_n^m, p_n^m)_{n=1}^{N_m}$. We now state and prove our generalisation of the result in the main text.²⁰

Proposition 10. *Suppose Condition (8) holds. Then there exists $\varepsilon > 0$ and K sufficiently large that, for all $m > K$, the following is true: There exist adjacent types $\theta_{n'}^m, \theta_{n'+1}^m, \dots, \theta_{n''}^m$ with $x_{n'}^m, x_{n'+1}^m, \dots, x_{n''}^m \in [\varepsilon, 1 - \varepsilon]$ and $\theta_{n''}^m - \theta_{n'}^m \geq \varepsilon$.*

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 $\theta < \hat{\theta}$, $L_{m_k}(\theta) \rightarrow 0$, while (ii) for all $\theta > \hat{\theta}$, $L_{m_k}(\theta) \rightarrow 1$. Then profits in the optimal mechanisms along the sequence of environments (E^{m_k}) converge to $\hat{\theta}(1 - F(\hat{\theta}))$.*

Proof. 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 $p_n^{m_k} > \hat{\theta} - \eta$.²¹ 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 $\eta' \in (0, \bar{\theta} - \hat{\theta})$ 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

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

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

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

²⁰Proposition 10 is more general than Proposition 3 because $v(p; \theta)$ not restricted to be CARA and Condition (8) is a generalization of Condition (2).

²¹Recall 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 .

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 $p_{\hat{n}}^{l_{k'}} \geq \underline{\theta}$ if $x_{\hat{n}}^{l_{k'}} > 0$. 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 for type $\theta_{n_{l_{k'}}}^{l_{k'}}$. 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 $\theta_{N_{l_k}}^{l_k}$ (recall that payments in the optimal mechanism are increasing in the buyer's type). By Step 1, there is a choice of types ($\theta_{n_{l_k}}^{l_k}$) such that $x_{n_{l_k}}^{l_k} \rightarrow 1$ and $\theta_{n_{l_k}}^{l_k} \rightarrow \hat{\theta}$. Note then that types ($\theta_{n_{l_k}}^{l_k}$) are assigned by the optimal mechanism a probability of allocation approaching one, and these types pay no more than $\theta_{n_{l_k}}^{l_k}$.

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

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

We conclude that, for all k sufficiently large, type $\theta_{N_{l_k}}^{l_k}$ strictly prefers to mimic type $\theta_{n_{l_k}}^{l_k}$ 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 $\theta_n^{m_k}$ above $\hat{\theta} + \eta$ make payments in the optimal mechanism for model E^{m_k} that are within η of $\hat{\theta}$, (ii) types $\theta_n^{m_k}$ above $\hat{\theta} + \eta$ acquire the good with probability at least $1 - \eta$, (iii) types $\theta_n^{m_k}$ 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 $\eta \rightarrow 0$, which shows that profits converge to $\hat{\theta}(1 - F(\hat{\theta}))$ considering optimal mechanisms along the sequence of environments (E^{m_k}). \square

(End of the proof of Lemma 3. The proof of Proposition 10 continues.)

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

Lemma 4. *Suppose that Condition (8) 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 $L_{m_k}(\theta) \rightarrow 0$ for $\theta < \hat{\theta}$ and $L_{m_k}(\theta) \rightarrow 1$ for $\theta > \hat{\theta}$.*

Proof. Suppose for a contradiction that there is such a $\hat{\theta}$ with $L_{m_k}(\theta) \rightarrow 0$ for $\theta < \hat{\theta}$ and $L_{m_k}(\theta) \rightarrow 1$ 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 9. 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 $k \rightarrow \infty$ 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 9, 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 . \square

(End of the proof of Lemma 4. The proof of Proposition 10 continues.)

We now conclude the proof of Proposition 10. 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 $\theta_{n'}^m, \theta_{n'+1}^m, \dots, \theta_{n''}^m$ with $\theta_{n''}^m - \theta_{n'}^m \geq \varepsilon$. 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 $\theta_{n'_m}^m$ for which $x_{n'_m}^m \geq \varepsilon$ and a largest type $\theta_{n''_m}^m$ for which $x_{n''_m}^m \leq 1 - \varepsilon$, and $\theta_{n''_m}^m - \theta_{n'_m}^m < \varepsilon$, (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 $\theta_{n'_{m_k}}^{m_k}$ for which $x_{n'_{m_k}}^{m_k} \geq 1/k$ and a largest type $\theta_{n''_{m_k}}^{m_k}$ for which $x_{n''_{m_k}}^{m_k} \leq 1 - 1/k$, and $\theta_{n''_{m_k}}^{m_k} - \theta_{n'_{m_k}}^{m_k} < 1/k$, (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 $\theta_{n'_{m_{k_j}}}^{m_{k_j}}$ 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 $\theta_{n'_{q_l}}^{q_l} \rightarrow \hat{\theta}$ 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 $L_{q_l}(\theta) \rightarrow 0$ and for any $\theta > \hat{\theta}$, we have $L_{q_l}(\theta) \rightarrow 1$. Again we have a violation of Lemma 4. This completes the proof of Proposition 10. \square

(End of the proof of Proposition 10. The proof of Proposition 3 continues.)

As we explain above (recall Footnote 20), Proposition 10 is a generalization of Proposition 3, hence the proof of Proposition 3 is complete. \square

Proof of Proposition 4

Proof. Determining the upper bound on seller profits. We begin by showing the first part of the proposition: that the seller's expected discounted total profits in any incentive-compatible price-path mechanism are no greater than Π^* .

Measurability of payoffs. We begin by showing the measurability of the players' payoffs so that the expressions for expected payoffs are well-defined. First, we show that $r \sum_{n=1}^N \beta_n e^{-r\tau_{n,t}} P_{\tau_{n,t}}$, the integrand in Equation (4), is $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ measurable. To do so, we extend the function $P_t(\omega)$ by setting $P_\infty(\omega) = 0$ for all $\omega \in \Omega$ and recall $e^{-\infty} \equiv 0$, leaving the value of the integrand in Equation (4) unchanged. We now note that, for each θ_n , $e^{-r\tau_{n,t}(\omega)}$ is a $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ measurable function into the unit interval, which follows because (i) e^{-rz} is measurable as a function of $z \in \bar{\mathbb{R}}_+$ (where $\bar{\mathbb{R}}_+$ denotes the non-negative extended reals) by continuity of the exponential function and by our definition of $e^{-\infty}$, (ii) because $\tau_{n,t}(\omega)$ is assumed $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ measurable, and (iii) because the composition of measurable functions is measurable. The result then follows (using the measurability of the product of measurable functions, as well as of sums of measurable functions) if we can show that, for each θ_n , $P_{\tau_{n,t}(\omega)}(\omega)$ is also $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ measurable.

We show the measurability of $P_{\tau_{n,t}(\omega)}(\omega)$ by viewing it as the composition $(t, \omega) \mapsto (\tau_t(\omega), \omega) \mapsto P_{\tau_{n,t}(\omega)}(\omega)$. The first mapping is measurable from $(\mathbb{R}_+ \times \Omega, \mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F})$ into $(\bar{\mathbb{R}}_+ \times \Omega, \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F})$ because the component functions are measurable.²² Then, we want to show that $(t, \omega) \mapsto P_t(\omega)$ is measurable from $(\bar{\mathbb{R}}_+ \times \Omega, \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F})$ to $(\mathbb{R}_+, \mathcal{B}_{\mathbb{R}_+})$. Considering the preimage A of any Borel subset B of \mathbb{R}_+ , either $A \in \mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ if B excludes zero (by the measurability of P_t as originally defined), or otherwise $A = A' \cup (\{\infty\} \times \Omega)$ where $A' \in \mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ and $\{\infty\} \times \Omega \in \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F}$. The result then follows because $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F} \subset \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F}$.²³

Because $r \sum_{n=1}^N \beta_n e^{-r\tau_{n,t}} P_{\tau_{n,t}}$ is non-negative, and by the established measurability, we may apply Tonelli's Theorem to conclude the seller's expected discounted total profits satisfy

$$\mathbb{E} \left[\int_0^\infty \sum_{n=1}^N \beta_n e^{-r\tau_{n,t}} P_{\tau_{n,t}} r dt \right] = \int_0^\infty \mathbb{E} \left[\sum_{n=1}^N \beta_n e^{-r\tau_{n,t}} P_{\tau_{n,t}} \right] r dt.$$

We view $\mathbb{E} \left[\sum_{n=1}^N \beta_n e^{-r\tau_{n,t}} P_{\tau_{n,t}} \right]$ as (the date-zero value of) expected per-buyer profits from arrivals at date t . We will show that these profits are no greater than $e^{-rt} \Pi^*$ for any incentive-compatible price-path mechanism (P, τ) , thus establishing that expected discounted total profits over all buyers are bounded by Π^* .

Now, consider briefly the measurability of the integrands in Equation (3); take that on the left side for instance. Similar to above, for given t and θ_n , we have that $e^{-r\tau_{n,t}(\omega)} \nu_n(P_{\tau_{n,t}(\omega)}(\omega))$ is measurable

²²This guarantees, in particular, that the preimages of "rectangular" sets $C \times D$ with $C \in \mathcal{B}_{\bar{\mathbb{R}}_+}$ and $D \in \mathcal{F}$ are measurable; and this is sufficient because these rectangular sets generate $\mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F}$.

²³To see this, define the sigma algebra $\mathcal{S} \equiv \{E \cap (\mathbb{R}_+ \times \Omega) : E \in \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F}\}$. It contains any set $C \times D$ with $C \in \mathcal{B}_{\mathbb{R}_+}$ and $D \in \mathcal{F}$ as any such set $C \times D \in \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F}$ (since $\mathcal{B}_{\mathbb{R}_+} \subset \mathcal{B}_{\bar{\mathbb{R}}_+}$). Therefore, $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F} \subset \mathcal{S}$. Also, $\mathbb{R}_+ \times \Omega \in \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F}$, so (using that $\mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F}$ is closed under intersections) $\mathcal{S} \subset \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F}$.

with respect to \mathcal{F} . For instance, measurability of $P_{\tau_{n,t}(\omega)}(\omega)$ can be seen by considering the composite mapping $\omega \mapsto (\tau_{n,t}(\omega), \omega) \mapsto P_{\tau_{n,t}(\omega)}(\omega)$. The first map is measurable from (Ω, \mathcal{F}) to $(\bar{\mathbb{R}}_+ \times \Omega, \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F})$. For the second map, we note $P_s(\omega)$ is measurable (as a function of s and ω) from $(\bar{\mathbb{R}}_+ \times \Omega, \mathcal{B}_{\bar{\mathbb{R}}_+} \otimes \mathcal{F})$ to $(\mathbb{R}_+, \mathcal{B}_{\mathbb{R}_+})$ as established above.

Necessary conditions for incentive compatibility. For a candidate price-path mechanism, incentive compatibility requires that no type θ_n arriving at t strictly prefers to purchase according to the stopping time $\tau_{m,t}$, $m \neq n$, and also that no such type θ_n strictly prefers not to purchase. These requirements represent a subset of the incentive constraints, allowing us to formulate a relaxed program. The optimum in this program represents an upper bound on the principal's expected profits.

Relaxed program. The relaxed program can be formally stated as maximising by random price path P and stopping times $\{\tau_{1,t}, \tau_{2,t}, \dots, \tau_{N,t}\}$,

$$e^{-rt} \sum_{n=1}^N \beta_n \mathbb{E}[e^{-r(\tau_{n,t}-t)} P_{\tau_{n,t}}] \quad (11)$$

subject to, for all θ_n and θ_m ,

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

and to

$$\mathbb{E}[e^{-r(\tau_{n,t}-t)} v_n(P_{\tau_{n,t}})] \geq 0,$$

where we have removed the factor e^{-rt} from inside the expectations.

For any random price path P and stopping times $\{\tau_{1,t}, \tau_{2,t}, \dots, \tau_{N,t}\}$, we specify, for each $n \in \{1, \dots, N\}$, $\check{x}_n = \mathbb{E}[e^{-r(\tau_{n,t}-t)}]$. There are then CDFs \check{H}_n , $n = 1, \dots, N$, for bounded random variables, such that, for any θ_n and θ_m ,

$$\mathbb{E}[e^{-r(\tau_{n,t}-t)} P_{\tau_{n,t}}] = \check{x}_n \int p d\check{H}_n(p), \quad \text{and} \quad (12)$$

$$\mathbb{E}[e^{-r(\tau_{m,t}-t)} v_n(P_{\tau_{m,t}})] = \check{x}_m \int v_n(p) d\check{H}_m(p). \quad (13)$$

Hence the objective and constraints are exactly the same as in the static mechanism design problem of Section 3. This shows that, given satisfaction of the constraints in the relaxed program stated above, (11) is no greater than $e^{-rt} \Pi^*$ as required.

Verifying expected payoffs. We now show the equality in Equation (13) and derive $\check{H}_m(p)$. The proof for Equation (12) follows the same reasoning and is omitted. We may assume the event $\{\tau_{m,t} < \infty\}$ has positive probability, as otherwise the claim is immediate.

Note that

$$\begin{aligned} \mathbb{E}[e^{-r(\tau_{m,t}-t)} v_n(P_{\tau_{m,t}})] &= \Pr(\{\tau_{m,t} < \infty\}) \mathbb{E}[e^{-r(\tau_{m,t}-t)} v_n(P_{\tau_{m,t}}) | \tau_{m,t} < \infty] \\ &= \Pr(\{\tau_{m,t} < \infty\}) \mathbb{E}[\mathbb{E}[e^{-r(\tau_{m,t}-t)} v_n(P_{\tau_{m,t}}) | P_{\tau_{m,t}}] | \tau_{m,t} < \infty]. \end{aligned}$$

The first conditional expectation conditions on the positive probability event $\{\tau_{m,t} < \infty\}$. This allows us to define the random variable $e^{-r(\tau_{m,t}-t)} v_n(P_{\tau_{m,t}})$ on a restricted probability space $(\{\tau_{m,t} < \infty\}, \mathcal{F}_{\{\tau_{m,t} < \infty\}}, \mathcal{P}_{\{\tau_{m,t} < \infty\}})$ such that $\mathcal{F}_{\{\tau_{m,t} < \infty\}} = \{A \in \mathcal{F} : A \subset \{\tau_{m,t} < \infty\}\}$ and, for any

$A \in \mathcal{F}_{\{\tau_{m,t} < \infty\}}$, $\mathcal{P}_{\{\tau_{m,t} < \infty\}}(A) = \frac{\mathcal{P}(A)}{\mathcal{P}(\{\tau_{m,t} < \infty\})}$. Moreover, given the boundedness of the stopped price $P_{\tau_{m,t}}$, the random variable $e^{-r(\tau_{m,t}-t)}v_n(P_{\tau_{m,t}})$ is bounded and hence integrable on $\{\tau_{m,t} < \infty\}$. This allows us to define (see, e.g., p. 445 of Billingsley, 1995) an integrable random variable $\mathbb{E}[e^{-r(\tau_{m,t}-t)}v_n(P_{\tau_{m,t}})|P_{\tau_{m,t}}]$ which is the conditional expectation on $\{\tau_{m,t} < \infty\}$.

By Theorem 34.3 of Billingsley (1995), we have that $\mathbb{E}[e^{-r(\tau_{m,t}-t)}v_n(P_{\tau_{m,t}})|P_{\tau_{m,t}}] = v_n(P_{\tau_{m,t}})\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}}]$. Since the conditional expectation $\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}}]$ is a measurable random variable on the sigma field generated by $P_{\tau_{m,t}}$ on $\{\tau_{m,t} < \infty\}$, we may assume that there is a measurable function $\varphi_m: \mathbb{R}_+ \rightarrow [0, 1]$ such that $\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}}] = \varphi_m(P_{\tau_{m,t}})$ and we will write $\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}} = p] = \varphi_m(p)$ (see Definition 8.24 of Klenke, 2013). We have that

$$\begin{aligned} \mathbb{E}[v_n(P_{\tau_{m,t}})\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}}]|\tau_{m,t} < \infty] &= \int_{\{\tau_{m,t} < \infty\}} v_n(P_{\tau_{m,t}}(\omega))\varphi_m(P_{\tau_{m,t}}(\omega))d\mathcal{P}_{\{\tau_{m,t} < \infty\}}(\omega) \\ &= \int_{\mathbb{R}_+} v_n(p)\varphi_m(p)dG_m(p), \end{aligned}$$

where $G_m(p)$ is the conditional CDF given by $G_m(p) = \mathcal{P}_{\{\tau_{m,t} < \infty\}}(\{P_{\tau_{m,t}} \leq p\})$.

After dividing and multiplying by $\mathbb{E}[e^{-r(\tau_{m,t}-t)}|\tau_{m,t} < \infty]$, we have shown that

$$\begin{aligned} &\mathbb{E}[e^{-r(\tau_{m,t}-t)}v_n(P_{\tau_{m,t}})] \\ &= \mathbb{E}[e^{-r(\tau_{m,t}-t)}|\tau_{m,t} < \infty]\mathcal{P}(\{\tau_{m,t} < \infty\}) \int_{\mathbb{R}_+} v_n(p) \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}} = p]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|\tau_{m,t} < \infty]} dG_m(p) \\ &= \mathbb{E}[e^{-r(\tau_{m,t}-t)}] \int_{\mathbb{R}_+} v_n(p) \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}} = p]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|\tau_{m,t} < \infty]} dG_m(p). \end{aligned}$$

We are now in a position to specify \check{H}_m . First, define $\check{H}_m: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ by

$$\check{H}_m(p) = \int_{[0,p]} \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}} = q]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|\tau_{m,t} < \infty]} dG_m(q).$$

Note that $\lim_{p \rightarrow \infty} \check{H}_m(p) = 1$ by the law of total expectation and we can extend \check{H}_m by letting $\check{H}_m(p)$ equal zero for $p < 0$. \check{H}_m is non-decreasing as the integrand is non-negative and G_m is non-decreasing. We can show that \check{H}_m is right continuous by considering any sequence $(p_k)_{k=1}^\infty$ with p_k approaching p from above. Then

$$\begin{aligned} \check{H}_m(p_k) - \check{H}_m(p) &= \int_{\mathbb{R}_+} \mathbf{1}_{(p,p_k]}(q) \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}} = q]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|\tau_{m,t} < \infty]} dG_m(q) \\ &\rightarrow 0 \end{aligned}$$

by the dominated convergence theorem, as $\mathbf{1}_{(p,p_k]}(q) \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|P_{\tau_{m,t}} = q]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)}|\tau_{m,t} < \infty]} \rightarrow 0$ almost everywhere.²⁴ Hence

²⁴This argument, and indeed the remainder of the proof, follows closely the proof of Lemma 1.16 in Seppäläinen (2012).

\check{H}_m is a CDF. To conclude the proof, it will then be enough to show that

$$\int_{\mathbb{R}_+} v_n(p) \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | P_{\tau_{m,t}} = p]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | \tau_{m,t} < \infty]} dG_m(p) = \int_{\mathbb{R}_+} v_n(p) d\check{H}_m(p). \quad (14)$$

For sets $A \in \mathcal{B}_{\mathbb{R}_+}$, define

$$\mu_m(A) = \int_A \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | P_{\tau_{m,t}} = q]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | \tau_{m,t} < \infty]} dG_m(q).$$

Note that μ_m is a measure (Billingsley, 1995, Theorem 16.9). Using that $v_n(p) \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | P_{\tau_{m,t}} = p]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | \tau_{m,t} < \infty]}$ is bounded on the support of G_m , the integral on the left side of Equation (14) is equal to $\int_{\mathbb{R}_+} v_n(p) d\mu_m(p)$ (see Billingsley, 1995, Theorem 16.11).

We now show that Equation (14) holds if the domain of integration on each side were restricted to $(0, +\infty)$. To see this, consider the Lebesgue-Stieltjes measure ν_m obtained from \check{H}_m ; it is enough to show that ν_m and μ_m agree on the Borel subsets of $(0, +\infty)$. We can note that, for any interval $(a, b]$ contained in $(0, +\infty)$,

$$\nu_m((a, b]) = \check{H}_m(b) - \check{H}_m(a) = \int_{(a,b]} \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | P_{\tau_{m,t}} = q]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | \tau_{m,t} < \infty]} dG_m(q) = \mu_m((a, b]).$$

We can then assume that $\nu_m((0, +\infty)) = \mu_m((0, +\infty)) > 0$ as otherwise the result is immediate. Then $\frac{\nu_m}{\nu_m((0, +\infty))}$ and $\frac{\mu_m}{\mu_m((0, +\infty))}$ are probability measures on $\mathcal{B}_{(0, +\infty)}$. Also, the collection of subsets $(a, b]$ contained in $(0, +\infty)$, together with the empty set, form a π -system. Theorem 10.3 of Billingsley (1995) then states that $\frac{\nu_m}{\nu_m((0, +\infty))}$ and $\frac{\mu_m}{\mu_m((0, +\infty))}$ are equal on $\mathcal{B}_{(0, +\infty)}$, and so ν_m and μ_m are also equal.

We can then note additionally that

$$\mu_m(\{0\}) = \frac{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | P_{\tau_{m,t}} = 0]}{\mathbb{E}[e^{-r(\tau_{m,t}-t)} | \tau_{m,t} < \infty]} G_m(0) = \check{H}_m(0) = \nu_m(\{0\}),$$

and hence conclude (again using Theorem 16.9 of Billingsley, 1995) that Equation (14) holds for the stated domain of integration \mathbb{R}_+ .

Attaining or approaching expected profits Π^* . We now show that there exist price-path mechanisms for which seller expected discounted total profits at least approach the upper bound Π^* in the sense of the proposition. Recall the assignment of the J indices in the main text. In particular, we let $x_0^* \equiv 0$ and let $(n_j)_{j=1}^J$ be the (unique) increasing sequence containing *all* indices satisfying $x_{n_{j-1}}^* < x_{n_j}^*$, where each x_n^* is the probability of allocation in the optimal static mechanism. Then, $(p_{n_j}^*)_{1 \leq j \leq J}$ gives the prices in the static mechanism.

The main text describes price-path mechanisms achieving profits Π^* for $J=1$ and $J=2$, however we have not yet verified the measurability conditions. We do this now for $J=2$ (they are immediate for $J=1$). To formally define the price process, we let $(N_t)_{t \in \mathbb{R}_+}$ be a Poisson counting process defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{R}_+}, \mathcal{P})$, and with arrival rate $\lambda_{n_1} = r x_{n_1}^* / (1 - x_{n_1}^*)$. We follow Billingsley (1995, p 298) and associate each outcome ω with a sample path satisfying a certain Condition 0°. This states: “For each ω , $N_t(\omega)$ is a [finite] non-negative integer for $t \geq 0$, $N_0(\omega) = 0$, and $\lim_{t \rightarrow \infty} N_t(\omega) = \infty$; further, for each ω , $N_t(\omega)$ as a function of t is nondecreasing and right-continuous, and at the points of discontinuity

the saltus $N_t(\omega) - \sup_{s < t} N_s(\omega)$ is exactly 1." We take $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ to be the natural filtration.

We then define the price process $(P_t)_{t \in \mathbb{R}_+}$ by

$$P_t(\omega) = \begin{cases} \theta_{n_1} & \text{if } t = \inf\{s : N_s(\omega) \geq m\} \text{ for some } m \in \mathbb{N} \setminus \{0\}, \\ p_{n_2}^* & \text{otherwise.} \end{cases}$$

The underlying probability space is the same, and note that the filtration $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ introduced above is still the natural filtration.²⁵ This process is progressively measurable. To see this, we can define, for each $\eta > 0$, a process $(P_t^\eta)_{t \in \mathbb{R}_+}$ by

$$P_t^\eta(\omega) = \begin{cases} \theta_{n_1} & \text{if } t \in [\inf\{s : N_s(\omega) \geq m\}, \inf\{s : N_s(\omega) \geq m\} + \eta) \text{ for some } m \in \mathbb{N} \setminus \{0\}, \\ p_{n_2}^* & \text{otherwise.} \end{cases}$$

The underlying probability space is the same. The process is adapted to the same filtration $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ and has right-continuous sample paths, implying it is progressively measurable (see, e.g., Theorem 25.8 of Klenke, 2013). That is, for any $s \geq 0$, $P_t^\eta(\omega)$ restricted to a function on $[0, s] \times \Omega$ is measurable with respect to $\mathcal{B}_{[0, s]} \otimes \mathcal{F}_s$. But, $P_t(\omega) = \lim_{\eta \rightarrow 0} P_t^\eta(\omega)$ is similarly measurable as a pointwise limit (see, e.g., Theorem 13.4 (ii) of Billingsley, 1995), establishing the result.

Now, consider the stopping times. We can write, for all t and ω ,

$$\tau_{n,t}(\omega) = \begin{cases} t & \text{if } n \geq n_2, \\ \inf\{s \in [t, \infty) : P_s(\omega) = \theta_{n_1}\} & \text{if } n_1 \leq n < n_2, \\ +\infty & \text{if } n < n_1. \end{cases}$$

For any type θ_n and positive integer m , we can define $\tau_{n,t}^m(\omega) = \tau_{n, \lfloor \frac{mt}{m} \rfloor}(\omega)$. Then, for any Borel set B of \mathbb{R}_+ , the preimage under $\tau_{n,t}^m$ is

$$\bigcup_{\{a : a = \frac{q}{m} \text{ for } q \in \mathbb{N} \cup \{0\}\}} [a, a + \frac{1}{m}) \times \tau_{n,a}^{-1}(B),$$

where $\tau_{n,a}^{-1}(B) = \{\omega \in \Omega : \tau_{n,a}(\omega) \in B\}$. Each $[a, a + \frac{1}{m}) \times \tau_{n,a}^{-1}(B)$, with $a = q/m$ for $q \in \mathbb{N} \cup \{0\}$, is in $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$, and so the above countable union is also in $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ showing the measurability of $\tau_{n,t}^m$. If $n < n_1$, then $\tau_{n,t}(\omega) = \tau_{n,t}^m(\omega) = +\infty$. For $n \geq n_1$, noting the left continuity of $\tau_{n,t}(\omega)$ in t for given ω , we have that $\lim_{m \rightarrow \infty} \tau_{n,t}^m(\omega) = \tau_{n,t}(\omega)$ and so $\tau_{n,t}(\omega)$ is measurable with respect to $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ as a pointwise limit (again, see Theorem 13.4 (ii) of Billingsley, 1995).

Assume from now on that $J \geq 3$. We consider price processes with J states $\{\sigma_j\}_{j=1}^J$ and parameterised by $\Lambda > 0$, anticipating that our result will be shown by taking $\Lambda \rightarrow \infty$. The price when the state is j is $p_{n_j}^*$. Transitions of the price process for each Λ are described as follows:

1. In state σ_J 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 $j = 1, \dots, J-1$, the state changes to state σ_j at rate Λ and, if $j > 1$, the state changes to state σ_{j-1} at rate $m_j^\Lambda \Lambda$, where m_j^Λ is determined below.

²⁵For any ω , the dates s at which $P_s(\omega) = \theta_{n_1}$ for $s \leq t$ contains the same information as the path $(N_s(\omega))_{s \in [0, t]}$. See Billingsley (1995, p 299).

For each $j=2, \dots, J-1$, we choose m_j^Λ so that the following equation is satisfied:

$$x_{n_{j-1}}^* = x_{n_j}^* \left(\frac{\Lambda}{\Lambda + m_j^\Lambda \Lambda + r} x_{n_{j-1}}^* + \frac{m_j^\Lambda \Lambda}{\Lambda + m_j^\Lambda \Lambda + r} \right) \Rightarrow m_j^\Lambda = \frac{((1-x_{n_j}^*)\Lambda + r)x_{n_{j-1}}^*}{\Lambda(x_{n_j}^* - x_{n_{j-1}}^*)} > 0. \quad (15)$$

For a given initial distribution of states (q_1, q_2, \dots, q_J) , with $\sum_{j=1}^J q_j = 1$, we have defined a finite-state continuous-time Markov process that is irreducible and positive recurrent. There is then, for each Λ , a unique stationary distribution $(q_1^\Lambda, q_2^\Lambda, \dots, q_J^\Lambda)$ (e.g., see p 396 of Beichelt, 2016). We take the date-zero distribution over states to be this stationary distribution.

Consider now the acceptance strategy τ of the buyers. Because the price process is time homogeneous, it is optimal for buyers to follow a stationary strategy, described simply in terms of the set of prices (or states) at which purchase occurs. We can then restrict attention to acceptance strategies that depend only on type and state (and not, for instance, arrival date, nor calendar time).

Below we only partially determine the strategy by specifying the highest price/state at which a buyer of each type purchases; this is enough to determine buyer discounted payoffs in state J . Buyers' decisions of whether to purchase at lower prices must be chosen to be buyer-optimal so as to guarantee incentive compatibility. However, the effect on the seller's discounted expected profits will be negligible as $\Lambda \rightarrow \infty$ and so the choice can be ignored in the analysis. This follows, in particular, because the probability of state J approaches one as $\Lambda \rightarrow \infty$. To see this, note the balance equation for state J is given by

$$0 = \Lambda \sum_{i=1}^{J-1} q_i^\Lambda - \lambda_{n_{J-1}} q_J^\Lambda,$$

where the first $J-1$ terms represent flows into state J and the final term represents flows out of this state. As $\Lambda \rightarrow \infty$, we must have $\sum_{i=1}^{J-1} q_i^\Lambda \rightarrow 0$, implying that $q_J^\Lambda \rightarrow 1$.

We now show that, if j' is the highest state in which a buyer of type θ_n purchases, then his expected discounted payoff in state J is $x_{n_{j'}}^* v_n(p_{n_{j'}}^*)$. To see this we can write a system of Bellman equations for states $j \geq j'$, with V_j giving the expected discounted continuation payoff in that state. We have $V_{j'} = v_n(p_{n_{j'}}^*)$. For $j = j'+1, j'+2, \dots, J-1$,

$$rV_j = \Lambda(V_j - V_j) + m_j \Lambda(V_{j-1} - V_j),$$

and

$$rV_j = \lambda_{n_{j-1}}(V_{j-1} - V_j).$$

Using Equations (15), the system is uniquely solved by $V_j = \frac{x_{n_{j'}}^*}{x_{n_j}^*} v_n(p_{n_{j'}}^*)$.²⁶ In particular, because $x_{n_J}^* = 1$ by Proposition 1, we have shown that $V_J = x_{n_{j'}}^* v_n(p_{n_{j'}}^*)$.

We now note that it is optimal for a type θ_n to set the highest price at which he purchases to p_n^* . This follows from the above expression for buyer expected discounted payoffs in state J together with incentive compatibility of the static mechanism. We specify our price-path mechanism by taking a buyer-optimal acceptance strategy with p_n^* being the highest price at which purchase occurs for each type θ_n .²⁷ Then, analogous to the argument for buyer expected discounted payoffs, expected discounted profit from type

²⁶To see uniqueness, initially let V_j to be determined and consider solving for V_j recursively, starting with $V_{j'+1}$. By an inductive argument, for $j \geq j'+1$ we have $V_j = A_j + B_j V_j$ for constants $A_j \in \mathbb{R}$ and $B_j \in (0, 1)$ that can be determined in the recursion. In particular, we have $V_j = A_j + B_j V_j$ with $B_j \in (0, 1)$ and so V_j is given uniquely by $\frac{A_j}{1-B_j}$.

²⁷This strategy can be determined as follows. First, fix the highest price at which purchase occurs to be p_n^* . Then let the decision whether to stop in state $j=1$ be optimal and recursively determine an optimal purchase decision for all higher states, up to the state immediately below the one with price p_n^* .

θ_n , conditional on state J , is $x_n^* p_n^*$. Hence, the seller's expected discounted profit from a buyer arriving to the market in state J (with probability of each type θ_n given by β_n) is Π^* .

Now consider per-buyer expected profits at any date t whose date- t value (given stationarity of the price process and buyer acceptance strategy) is independent of t . Recalling that the probability of state J approaches one as $\Lambda \rightarrow \infty$, per-buyer profits also approach Π^* . It is then easily seen that expected discounted *total* profits also approach Π^* , as required.

It remains to consider the admissibility requirements. We can consider a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{R}_+}, \mathcal{P})$ that associates outcomes with right-continuous sample paths taking values equal to the prices in the optimal static mechanism, and the filtration can be taken as the natural one. As the process is adapted and sample paths are right-continuous, the price process is progressively measurable.

For each θ_n , let S_n be the set of prices at which the buyer stops according to the above stationary strategy. We can let, for each θ_n , each $t \in \mathbb{R}_+$, and each $\omega \in \Omega$, $\tau_{n,t}(\omega) = \inf\{s \in [t, \infty) : P_s(\omega) \in S_n\}$. Note that this stopping time is right continuous in t .²⁸ We can define $\tau_{n,t}^m(\omega) = \tau_{n, \lfloor mt \rfloor / m}(\omega)$ which, as in the case of $J=2$, is easily shown to be $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ measurable. Then, the right continuity of $\tau_{n,t}(\omega)$ implies $\lim_{m \rightarrow \infty} \tau_{n,t}^m(\omega) = \tau_{n,t}(\omega)$, showing that $\tau_{n,t}(\omega)$ is $\mathcal{B}_{\mathbb{R}_+} \otimes \mathcal{F}$ measurable as the pointwise limit of measurable functions (again, see Theorem 13.4(ii) of Billingsley, 1995). \square

Proof of Proposition 5

Proof. Suppose $J \geq 2$. Fix a sequence of deterministic price-path mechanisms determined by $p^m : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ and τ^m . Assume for a contradiction that the payoff the seller achieves in the m -th mechanism, denoted Π^m , approaches the static optimal payoff Π^* as $m \rightarrow \infty$. For each $\varepsilon > 0$ and m , let T_ε^m be the set of times satisfying that

$$\max_n \{ |e^{-r(\tau_{n,t}^m - t)} - x_n^*|, |p_{\tau_{n,t}^m}^m - p_n^*| \} < \varepsilon.$$

That is, T_ε^m is the set of times for which the price and discounting under (p^m, τ^m) is similar to the price and probability of allocation under the optimal static mechanism for all types. By the uniqueness of the optimal static mechanism (Proposition 1), we have that, for all $\varepsilon > 0$,

$$\int_{T_\varepsilon^m} e^{-rt} r dt > 1 - \varepsilon \tag{16}$$

if m is large enough.

Fix some $m \in \mathbb{N}$, a small $\varepsilon > 0$, and some $t \in T_\varepsilon^m$. Define

$$t_1 := t - \log(x_{n_{J-1}}^* + \varepsilon)/r \quad \text{and} \quad t_2 := t - \log(x_{n_{J-1}}^* - \varepsilon)/r.$$

Note that $t < t_1 < t_2$ and, since $t \in T_\varepsilon^m$, we have that $\tau_{n_{J-1},t}^m \in (t_1, t_2)$ and $p_{\tau_{n_{J-1},t}^m}^m < p_{n_{J-1}}^* + \varepsilon$. The incentive compatibility of the type θ_{n_J} implies that, for all $t' \in [t, t_1]$, we have

$$e^{-r(\tau_{n_J,t'}^m - t')} v_{n_J}(p_{\tau_{n_J,t'}^m}^m) \geq e^{-r(\tau_{n_{J-1},t}^m - t')} v_{n_J}(p_{\tau_{n_{J-1},t}^m}^m) > e^{-r(t_2 - t')} v_{n_J}(p_{n_{J-1}}^* + \varepsilon).$$

The second inequality holds because $\tau_{n_{J-1},t}^m - t' < t_2 - t'$ and $p_{\tau_{n_{J-1},t}^m}^m < p_{n_{J-1}}^* + \varepsilon$.

²⁸To define the topology on the extended real numbers, we take the base for open sets to be the open intervals in \mathbb{R} plus sets of the form $[-\infty, a)$ and $(b, \infty]$, $a, b \in \mathbb{R}$.

Let \check{t} satisfy

$$e^{-r(t_2-\check{t})}v_{n_j}(p_{n_{j-1}}^* + \varepsilon) = v_{n_j}(p_{n_j}^* - \varepsilon).$$

Note that $t_2 - \check{t}$ is independent of m , and it is close to $-\log(x_{n_{j-1}}^*)/r$ if ε is small enough (recall that, from Proposition 1, we have that $v_{n_j}(p_{n_j}^*) = x_{n_{j-1}}^* v_{n_j}(p_{n_{j-1}}^*)$). Hence, for all $t' \in (\check{t}, t_1)$ we have

$$e^{-r(\tau_{n_j, t'}^m - t')} v_{n_j}(p_{\tau_{n_j, t'}^m}^m) > e^{-r(t_2 - t')} v_{n_j}(p_{n_{j-1}}^* + \varepsilon) > v_{n_j}(p_{n_j}^* - \varepsilon).$$

Note that, because $\tau_{n_j, t'}^m - t' \geq 0$ and $x_{n_j}^* = 1$, this implies that $p_{\tau_{n_j, t'}^m}^m < p_{n_j}^* - \varepsilon$, so $t' \notin T_\varepsilon^m$.

We have then shown that, for each $t \in T_\varepsilon^m$, there is an interval (\check{t}, t_1) satisfying that $(\check{t}, t_1) \cap T_\varepsilon^m = \emptyset$. It is easy to see that, if $\varepsilon > 0$ is small, $t_2 - \check{t}$ is close to $-\log(x_{n_{j-1}}^*)/r$ (as we argued above) and that $t_2 - t_1$ is close to 0 (note also that, when ε is small, $\check{t} - t$ is close to 0). Hence, for each small ε and $t \in T_\varepsilon^m$, there is a closeby interval of approximate size $-\log(x_{n_{j-1}}^*)/r > 0$ not intersecting with T_ε^m . This proves that equation (16) cannot hold, hence profits are bounded away from the static profits Π^* for all m , a contradiction. \square

Proof of Proposition 6

Proof. We begin by showing the following result.

Lemma 5. *For all dates $t > 0$, any positive probability event $A \in \mathcal{F}_t$, expected discounting to the next sale $\mathbb{E}[e^{-r(\tau_{1,t}-t)}]$ is equal to x_1^* .*

Proof. We have argued in the main text that $\mathbb{E}[\sum_{n=1}^2 \beta_n e^{-r(\tau_{n,t}-t)} P_{\tau_{n,t}}] = \Pi^*$ for almost all t . For any such t , this conclusion can be extended to expected profits conditional on positive probability events $A \in \mathcal{F}_t$. In particular, consider the expected discounted profits $\sum_{n=1}^2 \beta_n \mathbb{E}[e^{-r(\tau_{n,t}-t)} P_{\tau_{n,t}} | A]$ and note that the value that can be obtained is limited by the following incentive constraints: for all θ_n and θ_m , we must have

$$\mathbb{E}[e^{-r(\tau_{n,t}-t)} v_n(P_{\tau_{n,t}}) | A] \geq \mathbb{E}[e^{-r(\tau_{m,t}-t)} v_n(P_{\tau_{m,t}}) | A] \quad \text{and} \quad \mathbb{E}[e^{-r(\tau_{n,t}-t)} v_n(P_{\tau_{n,t}}) | A] \geq 0.$$

To see, for instance, that the first inequality is necessary for $\tau_{n,t}$ to be an incentive compatible stopping time, suppose that inequality fails for $\theta_m \neq \theta_n$. Then, we can define the stopping time

$$\hat{\tau}_{n,t}(\omega) = \begin{cases} \tau_{m,t}(\omega) & \text{if } \omega \in A, \\ \tau_{n,t}(\omega) & \text{otherwise,} \end{cases}$$

which represents a profitable deviation (i.e., a failure of incentive compatibility).

Letting $\check{x}_n^A = \mathbb{E}[e^{-r(\tau_{n,t}-t)} | A]$ (where recall $e^{-\infty} \equiv 0$), we can follow the steps in the proof of Proposition 4 to show that there is a collection of distributions (\check{H}_n^A) such that the objective can be rewritten as

$$\sum_{n=1}^2 \beta_n \check{x}_n^A \int p d\check{H}_n^A(p)$$

and the incentive constraints as

$$\check{x}_n^A \int v_n(p) d\check{H}_n^A(p) \geq \check{x}_n^A \int v_m(p) d\check{H}_m^A(p) \quad \text{and} \quad \check{x}_n^A \int v_n(p) d\check{H}_n^A(p) \geq 0.$$

By reference to Proposition 1, the objective is then bounded above by Π^* and attained only if, for $n=1, 2$, $\check{x}_n^A = x_n^*$ and \check{H}_n^A is degenerate at p_n^* . Therefore, if t is such that $\mathbb{E}\left[\sum_{n=1}^2 \beta_n e^{-r(\tau_{n,t}-t)} P_{\tau_{n,t}}\right] = \Pi^*$, we must have in particular $\check{x}_1^A = x_1^*$

We now extend $\mathbb{E}[e^{-r(\tau_{1,t}-t)}|A] = x_1^*$ to all $t > 0$ and positive probability events $A \in \mathcal{F}_t$. To first show that $\mathbb{E}[e^{-r(\tau_{1,t}-t)}|A] \geq x_1^*$, note that for any $s > t$, we have $\mathbb{E}[e^{-r(\tau_{1,t}-t)}|A] \geq e^{-r(s-t)}\mathbb{E}[e^{-r(\tau_{1,s}-s)}|A]$ because $\tau_{1,t} \leq \tau_{1,s}$. The claim then follows because $A \in \mathcal{F}_t \subset \mathcal{F}_s$ and we can take s arbitrarily close to t and such that $\mathbb{E}[e^{-r(\tau_{1,s}-s)}|A] = x_1^*$.

Now suppose for a contradiction that $t > 0$ and $\mathbb{E}[e^{-r(\tau_{1,t}-t)}|A] > x_1^*$ for some positive probability event $A \in \mathcal{F}_t$. In case A has probability less than one, note that $\mathbb{E}[e^{-r(\tau_{1,t}-t)}|A^c] \geq x_1^*$ where $A^c \in \mathcal{F}_t$ is the complement of A . Therefore, we can conclude using the law of total expectation that $\mathbb{E}[e^{-r(\tau_{1,t}-t)}] > x_1^*$. Then, for $s < t$, we have $\mathbb{E}[e^{-r(\tau_{1,s}-s)}] \geq e^{-r(t-s)}\mathbb{E}[e^{-r(\tau_{1,t}-t)}]$, implying $\mathbb{E}[e^{-r(\tau_{1,s}-s)}] > x_1^*$ for s sufficiently close to t . But this contradicts that $\mathbb{E}[e^{-r(\tau_{1,s}-s)}] = x_1^*$ almost everywhere. \square

(End of the proof of Lemma 5. The proof of Proposition 6 continues.)

Now, for any $t > 0$, any $A \in \mathcal{F}_t$ with positive probability, and all $s \geq t$, we can define $M_t^A(s) = \mathcal{P}(\omega : \tau_{1,t}(\omega) \leq s | A)$ to be the distribution of the time to next sale at t conditional on the event A . It is easy to see that $M_t^A(s) < 1$ for all $s \geq t$ and that $M_t^A(t) = 0$. We then have, for all $s > t$,

$$\mathbb{E}[e^{-r(\tau_{1,t}-t)}|A] = \int_{(t,s]} e^{-r(z-t)} dM_t^A(z) + (1 - M_t^A(s))\mathbb{E}[e^{-r(\tau_{1,t}-t)}|A, \tau_{1,t} > s],$$

which yields

$$x_1^* = \int_{(t,s]} e^{-r(z-t)} dM_t^A(z) + (1 - M_t^A(s))e^{-r(s-t)}x_1^*. \quad (17)$$

This equation can be used to show that $M_t^A(\cdot)$ is absolutely continuous on compact intervals and hence has a density, call it m_t^A , on the restriction to \mathbb{R}_+ . Consider s' and s'' with $t \leq s' < s''$. Then,

$$x_1^* = \int_{(t,s']} e^{-r(z-t)} dM_t^A(z) + \int_{(s',s'']} e^{-r(z-t)} dM_t^A(z) + (1 - M_t^A(s''))e^{-r(s''-t)}x_1^*$$

and indeed

$$x_1^* \geq \int_{(t,s']} e^{-r(z-t)} dM_t^A(z) + e^{-r(s''-t)}(M_t^A(s'') - M_t^A(s')) + (1 - M_t^A(s''))e^{-r(s''-t)}x_1^*.$$

Moreover,

$$x_1^* \leq \int_{(t,s']} e^{-r(z-t)} dM_t^A(z) + e^{-r(s'-t)}x_1^* - e^{-r(s''-t)}M_t^A(s')x_1^*.$$

Subtracting the latter inequality from the former we obtain

$$M_t^A(s'') - M_t^A(s') \leq \frac{x_1^*}{1 - x_1^*} (e^{r(s''-s')} - 1).$$

This establishes the claim.

We can then write Equation (17) as

$$x_1^* = \int_{(t,s]} e^{-r(z-t)} m_t^A(z) dz + (1 - M_t^A(s)) e^{-r(s-t)} x_1^*.$$

Differentiating in s yields a first-order linear non-homogeneous ODE

$$\frac{dM_t^A}{ds} = (1 - M_t^A) \frac{rx_1^*}{1 - x_1^*},$$

which is solved by $M_t^A(s) = 1 - e^{-\lambda_1(s-t)}$ with $\lambda_1 = \frac{rx_1^*}{1-x_1^*}$. Because the right side of the ODE is Lipschitz in M_t^A , the solution is globally unique by the Picard-Lindelöf Theorem. This completes the proof. \square

Proof of Proposition 7

Proof. Sufficiency. Incentive compatibility for low types is immediate so consider the stopping time of high types $\tau_{2,t}$ which prescribes stopping at t . It is always optimal for a high type to stop at the first sale, so the only deviations we need to consider are to stopping times $\hat{\tau}_{2,t}$ where either (i) the buyer simply waits for a sale (i.e., $\hat{\tau}_{2,t}(\omega) = \inf\{s \geq t : P_s(\omega) = \theta_1\}$), or (ii) the buyer purchases either at some date $s' > t$ or at the first sale, whichever is earlier (i.e., $\hat{\tau}_{2,t}(\omega) = \min\{s', \inf\{s \geq t : P_s(\omega) = \theta_1\}\}$). For the second specification of the stopping time to yield different expected payoffs than the first requires the event $\tau_{1,t} > s'$ to have positive probability; so assume this is true when considering the second case.

In the first case, the high type's expected payoff is $\mathbb{E}[e^{-r(\tau_{1,t}-t)} v_2(\theta_1)] = x_1^* v_2(\theta_1)$. By Proposition 1, this is equal to the expected payoff from immediate purchase, $v_2(p_2^*)$.

In the second case, let A be the (positive probability) event that $\tau_{1,t} > s'$, so that its complement A^c is the event $\tau_{1,t} \in [t, s']$. Then the high type's expected payoff under $\hat{\tau}_{2,t}$ is

$$\begin{aligned} & \mathbb{E}[e^{-r(\tau_{1,t}-t)} v_2(\theta_1) | A^c] \mathcal{P}(A^c) + e^{-r(s'-t)} v_2(p_2^*) \mathcal{P}(A) \\ &= \mathbb{E}[e^{-r(\tau_{1,t}-t)} v_2(\theta_1) | A^c] \mathcal{P}(A^c) + e^{-r(s'-t)} x_1^* v_2(\theta_1) \mathcal{P}(A) \\ &\leq \mathbb{E}[e^{-r(\tau_{1,t}-t)} v_2(\theta_1) | A^c] \mathcal{P}(A^c) + \mathbb{E}[e^{-r(\tau_{1,t}-t)} v_2(\theta_1) | A] \mathcal{P}(A) \\ &= x_1^* v_2(\theta_1) \\ &= v_2^*(p_2^*), \end{aligned}$$

where $v_2^*(p_2^*)$ is the expected payoff under $\tau_{2,t}$. The inequality follows by Condition (5).

Necessity. Now, consider a t and $s > t$ such that Condition (5) fails. Let A denote the event that $\tau_{1,t} > s$, and let A^c be its complement. Then we have $\mathcal{P}(A) > 0$ and $\mathbb{E}[e^{-r(\tau_{1,t}-s)} | A] < x_1^*$.

Since $\mathbb{E}[e^{-r(\tau_{1,t}-t)}] = x_1^*$, we have

$$x_1^* = \mathcal{P}(A^c) \mathbb{E}[e^{-r(\tau_{1,t}-t)} | A^c] + \mathcal{P}(A) e^{-r(s-t)} \mathbb{E}[e^{-r(\tau_{1,t}-s)} | A].$$

So, necessarily,

$$\mathcal{P}(A^c) \mathbb{E}[e^{-r(\tau_{1,t}-t)} | A^c] > (1 - \mathcal{P}(A)) e^{-r(s-t)} x_1^*. \quad (18)$$

For a high type arriving at t , consider the deviation $\hat{\tau}_{2,t}$ defined by purchasing at the next price discount

or at date s , whichever comes first. This yields an expected payoff

$$\mathcal{P}(A^c)\mathbb{E}[e^{-r(\tau_{1,t}-t)}|A^c]v_2(\theta_1)+\mathcal{P}(A)e^{-r(s-t)}v_2(p_2^*). \quad (19)$$

The first term of Equation (19) is strictly greater than $(1-\mathcal{P}(A)e^{-r(s-t)})x_1^*v_2(\theta_1)$ by Equation (18), while the second term is equal to $\mathcal{P}(A)e^{-r(s-t)}x_1^*v_2(\theta_1)$. Therefore, the expression in Equation (19) is strictly greater than $x_1^*v_2(\theta_1)$, which is equal to $v_2(p_2^*)$. This shows that purchasing immediately with probability one gives a strictly lower payoff than the defined deviation. That is, incentive compatibility fails. \square

Proof of Proposition 8

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 incentive compatible and individually rational 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 6. *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. The utility of type θ_n when reporting θ_k is

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

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.

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.²⁹ The payment specified to the low type is irrelevant since it is never charged (so we can let this equal θ_1). So, in this case, the statement of the lemma holds. We can then focus on the case where the original mechanism allocates with positive probability to the low type.

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 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

²⁹This is the price p that sets $\theta_2 - p + \lambda\eta(\rho_2 - p) = 0$. That is, the high type's willingness to pay is $\frac{\theta_2 + \lambda\eta\rho_2}{1 + \lambda\eta} \in (\rho_2, \theta_2)$.

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

$$x_1 \left(\int_{[0, \rho_2)} (\theta_2 - p + \eta(\rho_2 - p)) dH_1(p) + \int_{[\rho_2, \infty)} (\theta_2 - p + \lambda\eta(\rho_2 - p)) dH_1(p) \right). \quad (20)$$

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

$$x_1 (\theta_2 - \bar{p}_1 + \eta(\rho_2 - \bar{p}_1)). \quad (21)$$

The decrease in payoff for type θ_2 when mimicking θ_1 due to the change in the mechanism is the expression in Equation (20) less that in Equation (21). This is equal to

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

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 $\bar{p}_1 \leq \theta_1$, 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 29). This is strictly less than θ_2 as we assumed $\rho_2 < \theta_2$. \square

(End of the proof of Lemma 6. The proof of Proposition 8 continues.)

Now, to establish Proposition 8, 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 6, we

only need to impose the high type's incentive constraint.³⁰ This constraint can be written as:

$$x_2(\theta_2 - p_2 + \mu(\rho_2 - p_2)) \geq x_1(\theta_2 - \theta_1 + \eta(\rho_2 - \theta_1)), \quad (22)$$

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 (22) 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 (22), increases profits. We can therefore assume that the constraint binds. That is, p_2 is given by

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

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

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

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. 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 (6). □

³⁰The 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.