On dynamic pricing*

Ilia Krasikov Rohit Lamba

December 2020

Abstract

This paper builds a theory of dynamic pricing for the sale of timed goods such as travel tickets, hotel stays and concert seats. The main friction is private and evolving valuation of the buyer prior to the date of consumption. A combination of membership fee and continuously increasing prices induces a threshold response from the buyer, endogenously segmenting the market along timing of purchase. This pricing mechanism achieves the deterministic global optimum. Under standard assumptions, the analyst can estimate the fundamentals of the market from observables– time and price of sale. The tools developed here are shown to be useful in thinking about: refund contracts, global incentives and randomization in dynamic mechanisms, dynamic incentives beyond the one-shot deviation principle, and mapping dynamic pricing to the classic taxonomy of consumer-producer surplus and deadweight loss.

1 Introduction

Price discrimination refers to the idea of selling different goods at prices that are in different ratios to the marginal cost (Stigler [1987]). Its main goal is to segment the market of buyers into different prices based on observable or unobservable characteristics. In a survey on the topic Varian [1989] had written:

> "The easiest case is where the firm can explicitly sort consumers with respect to some exogenous category such as age. *A more complex analysis is necessary when the firm must price discriminate on the basis of some endogenous category such as time of purchase*. In this case the monopolist faces the problem of structuring his pricing so that consumers 'self-select' into appropriate categories."

Time of purchase as an endogenous determinant of price discrimination is the subject matter of this paper. Two examples to keep in mind are buying tickets for travel and booking a room in a hotel or on Airbnb. In both cases the timing of consumption is fixed, but the valuation for eventual

*Krasikov: krasikovis.main@gmail.com, Higher School of Economics Moscow; Lamba: rlamba@psu.edu; Pennsylvania State University. We are indebted to Vijay Krishna for his advice and guidance, and Karl Schurter for detailed comments. We are also also grateful to Nima Haghpanah, Berk Idem, Juuso Toikka, Rakesh Vohra, and seminar participants at University of Pennsylvania, Chinese University of Hong Kong, Indian Statistical Institute Delhi, Berlin-Bonn Theory Webinar, Canadian Economic Theory Conference at McGill University, and NSF/NBER/CEME Conference at University of Chicago for their comments.

consumption may change over time. Since the seller is not privy to the buyers' valuations, she can design a changing price path and segment buyers along *when* they self-select to trade. In particular, the seller monetizes the difference between buyers who have high initial valuations and buy early, and those that initially have a low valuation but might buy later if their valuation increases.

There are innumerable goods and services that we now use whose prices change frequently, and there are many reasons for why that may be the case. On the buyers' side, it can be fluctuating market size, changing tastes or the history of past purchases, and on the seller's side, it can be capacity consideration, changing costs, or experimentation for market research. In this paper we exclusively explore the channel of changing buyer valuations (or market size) over time.

The main conceptual message here is that the seller can use a combination of two-part tariff and second-degree price discrimination to maximize her profits. The first part of the pricing scheme is an upfront payment or membership fees that extracts total surplus modulo the rents that need to be paid to satisfy sell-selection. The second part is a time dependent sequence of prices that executes second-degree price discrimination wherein the buyers with differing valuations over time sort themselves into different timings of purchase.

The main technical novelty of the paper is that it presents a model of dynamic mechanism design that hitherto had not been solved due to reliance on the "standard local approach" that fails generically in our environment. This allows us to unpack the economic content of binding global incentive constraints in dynamic mechanism design and show its equivalence to simple pricing instruments. The pricing approach we take is shown to be useful in other problems such as identifying the gains from randomization and analyzing models beyond the one-shot deviation principle. It has potential applications to related economic environments such as optimal taxation.

The main substantive contribution to the applied literature on dynamic pricing is twofold. First, the two part-tariff we use has a classical price theoretic interpretation: Consider the "surplus triangle" generated by static pricing, divided into consumer and producer surplus and deadweight loss. Here upfront payment moves some consumer surplus to producer surplus, and sorting along timing of purchase moves some deadweight loss to producer surplus. Second, under some conditions, the analyst can back out fundamentals (distribution of values) from observables (price and time of sales) in a relatively straightforward estimation exercise afforded by our pricing approach.

The model features a buyer who wants to consume a good at a future date T. His (private) valuation evolves continuously according to a Poisson process: it is drawn from a prior F, then can change repeatedly upon the arrival of a Poisson shock with intensity λ, and in case of an arrival, the value is redrawn from the prior F. Poisson shocks here convey some new information, say a family emergency or meeting at work, that changes the value of the buyer.

A seller commits to a dynamic pricing mechanism $\langle M, \mathbf{p} \rangle$, where M is an upfront payment which grants the buyer access to a price path $\mathbf{p} = (p_t)_{t=0}^{T}$. After paying M the buyer's strategy boils down to an optimal stopping problem of when to trade. The seller takes the buyer's behavior into account and chooses the optimal pricing mechanism to maximize her profit.

This infinite dimensional pricing problem is reduced to the choice of a single variable, say α, and M and $\mathbf{p}$ are pinned down in closed form as functions of α. The seller chooses the optimal value, α^*, to maximize her profits. The optimal price path is shown to be smoothly increasing from some value p_0^* and to $p_T^* = \alpha^*$. The buyer trades at time t for price p_t^* iff $V_t \geqslant \alpha^* > \max_{s<t} V_s$. That is, the buyer waits until the first time the valuation process goes above α^* and buys exactly at that point. If it turns out that $\max_{t \geqslant 0} V_t < \alpha^*$, there is no trade.

The solution to the pricing problem is provided further credence by establishing that the allocation rule it implements is actually the global (deterministic) optimum of the dynamic mechanism design problem. To do this (following Myerson [1981]) we state the design problem in full generality, and by invoking the revelation principle, establish the optimum in the space of all possible indirect (or pricing) mechanisms.

The basic model studied here fits into the burgeoning literature on dynamic pricing, which starts at least as far back as Stokey [1979]. That paper provided a seminal benchmark of sorts– an impossibility result on the benefits of dynamic price discrimination. Qualitatively speaking, relaxing one or more assumptions from that environment generates scope for dynamic pricing. The literature on dynamic mechanism design pushes along this line of thought by focussing on, amongst other things, changing private valuations of the buyers (see Bergemann and Välimäki [2019] for the most recent survey).

The sale of a timed good by dynamically discriminating amongst buyers has been studied under the rubric of sequential screening, starting with Courty and Li [2000]. Most papers study a two-period model with different assumptions on the type space or the Markov process of valuations. Courty and Li [2000] invokes two types (business and economy class customers) in the first period to show the optimality of refund contracts. Esö and Szentes [2007] studies type dependent option contracts, where a continuum of types choose from a menu of premiums and strike prices.

A few related papers go beyond the two-period model but restrict attention to the random arrival of a single piece of new information. Deb [2014] studies a model similar to ours but restricts the type to change at most once according to a Poisson jump. It partially characterizes the optimal contract, when the local approach is satisfied, and provides an intuitive implementation in the form of of introductory pricing. Akan, Ata, and Dana Jr. [2015] re-interprets the two-period model as a continuous time setting in which the draw of the first period type in $[0, T]$ is the time at which the buyer is informed about his valuation. The value, once endowed to the buyer, is constant. The type dependent option contract observed in two-period models can then be interpreted as a price path for buyers who are informed at different times.[1]

In contrast to these papers, we consider an arbitrary finite time horizon with random strands of information modeled through arbitrary arrivals of Poisson shocks. This delivers a continuous price

[1]Other related papers include: Boleslavsky and Said [2013] generalizes Esö and Szentes [2007]'s option contracts to multiple periods through a model where the the buyer's types are i.i.d conditional on the first period information. Instead of offering a menu of option contracts, we study a single option contract of time dependent prices. Kruse and Strack [2015] considers the sale of a single good in a discrete time framework with changing private valuations of the buyer. At a high level, we add to their exercise of characterizing feasible allocations by identifying a tractable environment where the optimum is a threshold policy, and the threshold is history independent. Ely, Garrett, and Hinnosaar [2017] considers a multi player version of the two period model, and show that the optimum can be implemented through a double auction. Board and Skrzypacz [2016] and Garrett [2016] study related models of price discrimination where buyers arrive stochastically over time. Pavan, Segal, and Toikka [2014] presents a unified treatment of necessary conditions of incentive compatibility in dynamic mechanisms.

path and a sorting of dynamic valuations into time buckets for who buys when. The solution is completely characterized in that it gives the deterministic optimum with a distinct implementation strategy that is only time (and not type) dependent. In addition, it lends itself to an intuitive price-theoretic explanation of the optimum, in terms of consumer and producer surplus and deadweight loss, extending the insights of Bulow and Roberts [1989] from static to dynamic mechanism design.

The model and its solution generate qualitative predictions that could be potentially useful for an outside analyst who wants to rationalize data on dynamic pricing. First, the price path is continuously increasing, and hence the seller offers a steadily declining *advance purchase discount*. This result builds on earlier work, for example Dana Jr. [1998], that looks at static models with ex ante uncertainty to theorize such advance discounts. It is also complimentary to papers that motivate random discounts through stochastic arrivals of buyers, for example Garrett [2016].

Second, the distribution of sale times generates a simple test of the model. The probability or the fraction of completed sales is increasing in time. So, if the analyst has access to the full dataset including those goods that are unsold, this provides a quick smell test of the validity of our exercise. In a more realistic scenario where the analyst does not have access to the trades that do not materialize, she can construct a different statistic– expected time of sale, conditional on trade. This statistic is shown to be hump-shaped. which provides an altogether different way to test the basic validity of the model.

Third, beyond simple tests of validity, we execute an econometric exercise to parametrically estimate the fundamentals from the observables. Under some standard assumptions on the dataset that allow for the application of law of large numbers and identification of the thresholds of trade, we construct a consistent estimator of the parameters that define the distribution of valuations. This exercise is similar in spirit to the estimation of optimal auction models (eg. Laffont, Ossard, and Vuong [1995]), but incorporates for dynamic mechanisms. Importantly, the estimation exercise does not rely on the optimality of the price schedule. As long as the buyer is assumed to respond optimally to *any* sequence of prices posted by the seller, the fundamentals can be recovered. In addition, this exercise also produces an estimate of the deadweight loss associated with dynamic pricing.

Now, thinking through this largely normative exercise of achieving the optimal profit through a two-part tariff, the reader could invoke this positive concern: but we may not observe "upfront payments" in various dynamic pricing contexts such as travel tickets or hotel stays. An arguably positive implication of the model is that whenever possible, sellers *do* employ dual instruments in dynamic pricing, one that extracts some consumer surplus and other that materializes what would otherwise be missed opportunities for trade. To conceptualize the former, think about refund contracts, membership fees or even point-systems. We offer one such explanation in the form of tine dependent refunds, and show that it also implements the optimum.

Technically speaking, most studies in dynamic mechanism design have relied on models where local incentive constraints are sufficient for incentive compatibility.[2] The focus on such cases has

[2]The two exceptions to this are Akan, Ata, and Dana Jr. [2015] and Battaglini and Lamba [2019]. Both present two-period models with specific Markov processes where global incentives bind.

culminated in a line of thinking that since the approach is analogous to the static problem, dynamic considerations must be "irrelevant" (see Esö and Szentes [2017]). We show that going outside the box of the local approach can yield rich dynamics, and propose a model where dynamic ironing maps into a familiar price-theoretic taxonomy. So even though dynamic ironing can seem technically daunting, we argue that it can find a clear motivation in simpler pricing instruments.[3]

Finally, we illustrate the appeal of our dynamic pricing approach by (i) constructing a stochastic mechanism that improves upon the deterministic optimum, and (ii) solving the model where the information change is exogenously restricted by the number of permissible Poisson arrivals. The first problem allows for an information acquisition interpretation of refund contracts in dynamic pricing, and second problem captures situations where valuations change through a fixed number of sources. Both these problems don't seem to have a natural tack though the standard direct revelation mechanism approach. The pricing approach offers insights here that might otherwise be elusive.

Before diving into the general model, a simplified two-period version is presented in Section 2. It gets across the basic intuition of two-part tariffs as instruments of dynamic price discrimination and the economic content of binding global incentive constraints in dynamic mechanisms.

2 An example

Here we look at simple two-period example which communicates some of the key forces. A buyer wants to consume a good (or service) at the end of date 2. At date 1, he has a value for it, say v_1, that is distributed according to F on the unit interval. Come date 2, the value will remain the same with probability $1-\lambda$, that is $v_2 = v_1$, and it will be redrawn again from F with probability λ, that is $v_2 \sim F$. Here v_2 is payoff relevant, and v_1 acts a "signal" of the eventual consumption value with precision $1-\lambda$. The cost of production for the seller is zero. For further simplicity of exposition we will assume F is uniform on $[0,1]$ and $\lambda = 0.5$. The question we ask is this: What pricing mechanism(s) should the seller employ to maximize her profit?

Optimal pricing mechanism. We consider a menu of dynamic prices as depicted in Figure 1. The contract space is given by $\langle M, p, \alpha \rangle$. At the start the buyer gets to choose between paying a membership/entree fee M or ending the contract. If he pays M, he is granted access to two prices. At date 1, he can buy the good for a *forward price* p, or he can forgo p at date 1, and instead buy the good at date 2 for a *spot price* α. If he does not exercise either price, there is no trade. The key distinction between p and α is that the former is binding, it is paid no matter the realization of v_2, whereas if the buyer waits till date 2, he can decide upon observing v_2 whether to trade at α.

[3] Allocations determined by considering only local constraints feature a certain level of separation, viz. different types get different contracts. Often, when the type space in multidimensional (as in the dynamic model), these separable allocations do not satisfy relevant monotonicity properties necessary for incentive compatibility. Ironing refers to the identification of these monotonicity constraints whose shadow price is positive at the optimum, so the *virtual values* need to be smoothened and allocations pooled or bundled to make them incentive compatible. In taking the pricing approach to understand complex direct mechanism design problems which require ironing, we take a leaf from a rich literature in multidimensional mechanism design, see Armstrong [2016] for the most recent survey.

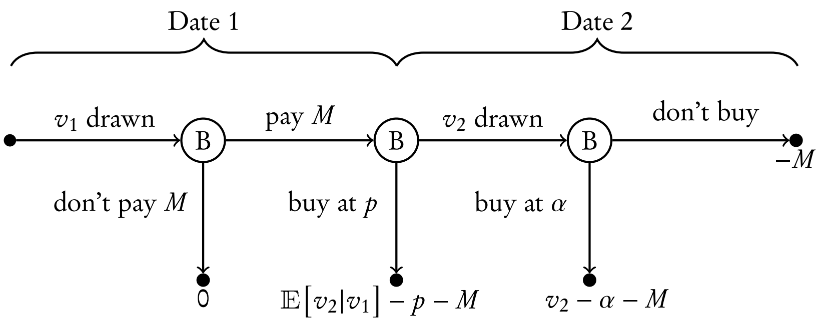

Figure 1: Timing of the contract $\langle M, p, \alpha \rangle$.

We will construct the mechanism depicted in Figure 1 in three gradual steps. Figure 2 represents the optimal trading regions for these progressively enriched pricing mechanisms. The x-axis represents valuation in the first period, y-axis valuation in the second period, and the shaded area represents the region of trade. Note that since the cost of the seller is zero, efficiency demands that the whole square in Figure 2 should be shaded.

In the first case, in Figure 2a, the seller ignores the dynamics of the problem and offers a spot price, say α. This earns the her an expected profit of $\alpha \mathbb{P}(v_2 \geqslant \alpha) = \alpha(1 - F(\alpha))$, where the equality follows from the fact that from an ex ante perspective $v_2 \sim F$. Finally, since F is uniform on $[0, 1]$, this term is $\alpha(1 - \alpha)$, and it is maximized at $\alpha = 1/2$, giving an expected profit of $1/4$.

The second contract, given by $\langle M, \alpha \rangle$, works as follows: charge M at the start of date 1, the payment of which grants the buyer the right to buy the good at date 2 for a spot price α. If the buyer does not pay M, the "game" ends with no trade. It is easy to show that the seller sets $M(\alpha) = \mathbb{E}[(v_2 - \alpha)^+ \mid v_1 = 0]$, i.e. the surplus of the lowest first period value buyer.[4] The expected profit is then given by $M(\alpha) + \alpha(1 - \alpha)$, which is maximized at $\alpha = 1/3$. This gives $M = 1/9$ and an expected profit of $1/3$. So, all types agree to pay M, and as Figure 2b shows, trade takes place when $v_2 \geq 1/3$. Finally, can the seller use the first period type to screen the buyers and is that profitable?

The contract space is now given by $\langle M, p, \alpha \rangle$, and its timing, as explained before, is depicted in Figure 1. Again, M is set to be the the expected surplus of the buyer whose valuation at date 1 is zero. The new price p is chosen by making the buyers with first period valuation above α indifferent between trading and waiting:

$$\underbrace{\mathbb{E}\left[v_2|v_1 = \alpha\right] - p(\alpha)}_{\text{net value for } v_1 = \alpha \text{ from taking the forward price}} = \underbrace{\mathbb{E}\left[(v_2 - \alpha)^+|v_1 = \alpha\right]}_{\text{net value for } v_1 = \alpha \text{ from waiting for spot price}}$$

Intuitively speaking, since p screens first period buyers, there is no advantage is using M to do that as well. So M is used only to extract surplus and not tinker with selection. Next, suppose p makes $v_1 = \alpha'$ indifferent between trading and waiting. If $\alpha' > \alpha$, then for p thus chosen, the buyer would never trade in the first period, If $\alpha < \alpha'$, then α can be reduced to increases the area of trade without disturbing selection in the first period, and then M can be increased to extract all

[4]Here a^+ denotes $\max\{0, a\}$.

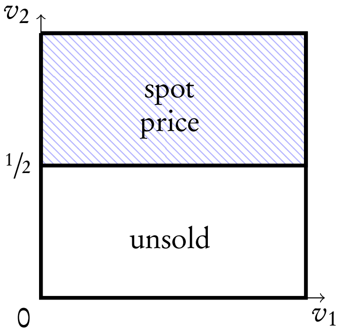

(a) Spot pricing with $\alpha = 1/2$.

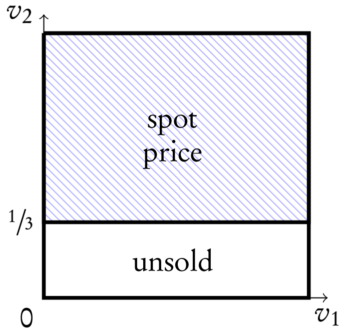

(b) Adding upfront payment: $M = 1/9, \alpha = 1/3$.

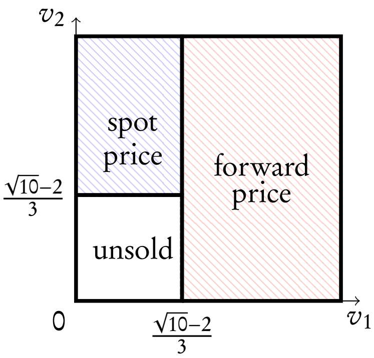

(c) Optimum:
$M \approx 0.094, p \approx 0.350, \alpha = \frac{\sqrt{10}-2}{3} \approx 0.387$.

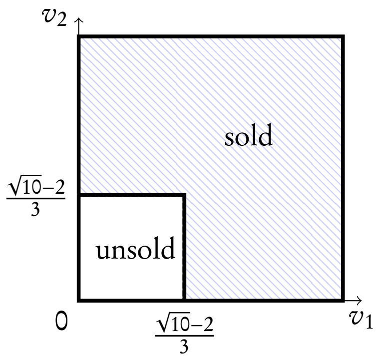

(d) Structure of trade at the optimum

Figure 2: Dynamic pricing in a two period example, F is uniform and $\lambda = 1/2$.

of the new surplus upfront.

So, M is chosen to extract the expected surplus of the lowest first period type and p is chosen to make $v_1 = \alpha$ indifferent between trading now and waiting, where α is also the final spot price. As a consequence, the unit square is split into three regions– trade in period 1, trade in period 2 and no trade, where the no trade region forms a smaller square of side α (see Figure 2c). We provide the formal arguments in the appendix for the intuition offered in the previous paragraph for the optimality of these chosen prices.

Putting it all together, the seller will choose $M = M(\alpha) := \mathbb{E}\left[(v_2 - \alpha)^+ | v_1 = 0\right]$ and $p = p(\alpha) := \mathbb{E}[v_2 | v_1 = \alpha] - M(\alpha)$. The expected profit of the seller is given by $\Pi(\alpha) = M(\alpha) + p(\alpha)\mathbb{P}(v_1 \geqslant \alpha) + \alpha\mathbb{P}(v_2 \geqslant \alpha | v_1 < \alpha)$. Substuiting and simplifying, we get:

$$\Pi(\alpha) := \alpha(1-\alpha) + M(\alpha) + 1/2 \times (1-\alpha)\int_0^\alpha v dF(v).$$

where the second and third terms represent respectively the extra profit from introducing upfront payment M and time dependent price p.

This optimization problem yields $\alpha = \frac{\sqrt{10}-2}{3}$. Thus, trade happens if and only if $\max\{v_1, v_2\} \geqslant \frac{\sqrt{10}-2}{3}$ (see Figures 2c and 2d), granting the seller an expected profit of approximately 0.354. Next we show this is the optimal dynamic contract– the seller cannot do better.

Optimal dynamic mechanism. We complete the analysis of the two-period example here by showing that our pricing strategy is the best the seller can do in terms of ex ante profits. To establish that we state the problem in a general two-period dynamic mechanism design framework, invoke the revelation principle, and then solve the dynamic problem to show that the optimal allocation is in fact given by $q(v_1, v_2) = \mathbb{1}\left(\max\{v_1, v_2\} \geqslant \frac{\sqrt{10}-2}{3}\right)$.

Now, a direct dynamic mechanism is a pair of an allocation $q \in \{0, 1\}$ and payments (p_1, p_2). The buyer reports a value in each period. Since the seller can commit, by the revelation principle there is no loss from looking only at the set of direct incentive compatible mechanisms where the buyer cannot strictly gain by misreporting his valuations.

Write $v_2 q(\hat{v}_1, \hat{v}_2) - p_2(\hat{v}_1, \hat{v}_2)$ for the buyer's second period payoff where $\hat{v}_t$ is the reported type and v_t is the true type. This payoff is independent of the true first period value, thus incentives "on-path" and "off-path" are identical. Keeping this in mind, we shall only consider $\hat{v}_1 = v_1$; then incentive compatibility at $t = 2$ reads as follows:

$$U_2(v_1, v_2) := v_2 q(v_1, v_2) - p_2(v_1, v_2) = \max_{\hat{v}_2 \in [0,1]} v_2 q(v_1, \hat{v}_2) - p_2(v_1, \hat{v}_2).$$

This constraint is quite standard and well-studied in the static mechanism design literature (Myerson [1981]). It can be shown that the above condition is equivalent to $v_2 \mapsto U_2(v_1, v_2)$ being a convex function with a derivative of $q(v_1, v_2)$.

Next, define the buyer's first period expected payoff from truth-telling as $U_1(v_1) := \mathbb{E}[U_2(v_1, v_2)] - p(v_1)$. Since the second period incentive constraint ensures that the buyer reverts to truth-telling, the incentive compatibility at $t = 1$ is the following:

$$U_1(v_1) = \max_{\hat{v}_1 \in [0,1]} {}^1\!/_2 \times U_2(\hat{v}_1, v_1) + {}^1\!/_2 \times \int_0^1 U_2(\hat{v}_1, v_2) dv_2 - p_1(\hat{v}_1)$$

Again, the function $U_1(v_1)$ is convex with a derivative of ${}^1\!/_2 \times q(v_1, v_1)$. However, unlike the second period, convexity is only necessary, but not sufficient for incentive compatibility. To produce the right condition, first subtract $U_1(\hat{v}_1)$ from each side of the incentive constraint and use the fact that the allocation pins down $U_1(v_1)$ and $U_2(v_1, v_2)$ up to a constant:

$$\int_{\hat{v}_1}^{v_1} q(v, v) dv \geqslant \int_{\hat{v}_1}^{v_1} q(\hat{v}_1, v) dv.$$

This condition, known as *integral monotonicity* (Pavan, Segal, and Toikka [2014]), characterizes implementability in the dynamic setting when information arrives gradually.

A next natural step is to write the seller's profit as a function of q. Because of linearity, this profit is the difference between surplus and buyer's expected payoff. The former is simply $\mathbb{E}\left[v_2 q(v_1, v_2)\right]$, whereas the latter is given by

$$\mathbb{E}[U_1(v_1)] = U_1(0) + {}^1\!/_2 \times \int_0^1 (1 - v) q(v, v) dv.$$

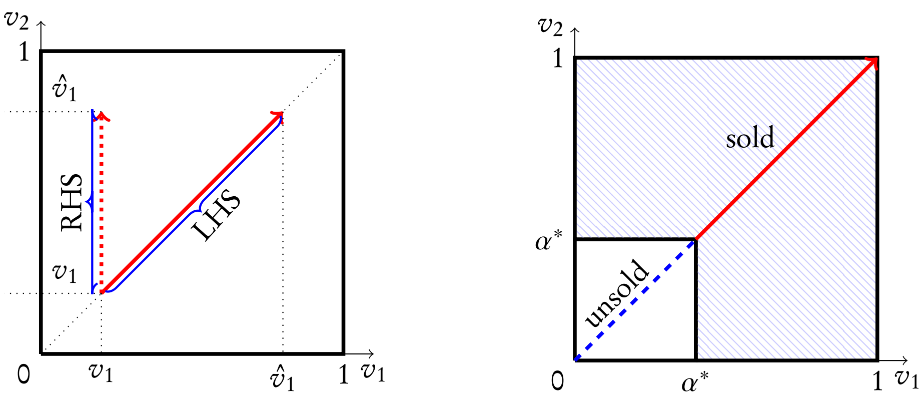

(a) Integral monotonicity (b) Construction of the optimum

Figure 3: Optimal dynamic mechanism

This expression is obtained by the dynamic envelope formula: $U'(v_1) = 1/2 \times q(v_1, v_1)$, and using integration by parts. It is important to note that locally the buyer's payoff only depends on the "diagonal" or constant allocation, that is when $v_1 = v_2$.

Now, we are in position to find the best contract for the seller. Of course, we require that the buyer cannot be forced to accept trade, in other words $U_1(v_1) \geqslant 0$. This pins downs $U_1(0) = 0$, and the seller's problem can finally be stated as

$$\max_{q \in \{0,1\}} \quad \mathbb{E}\left[v_2 q(v_1, v_2)\right] - 1/2 \times \int_0^1 (1-v) q(v,v) dv.$$

subject to (i) integral monotonicity and (ii) $v_2 \mapsto q(v_1, v_2)$ non-decreasing. A graphical illustration of both global constraints can be seen in Figure 3a. Integral monotonicity demands that average allocation along the diagonal must be greater than the allocation along its vertical projection, and second period monotonicity demands that allocation along any vertical line in the unit square should be monotonic.[5]

It is easy to see that $v \mapsto q(v,v)$ must be non-decreasing, thus there exists a number α such that $q(v,v) = 1$ if and only if $v \geqslant \alpha$. The splits the diagonal in Figure 3b into trade (red and solid) and no-trade (blue and dashed) regions. We claim that $q(v_1, v_2) = 0$ whenever $\max\{v_1, v_2\} < \alpha$.

- For $v_2 < v_1 < \alpha$, this is implied by (ii): $q(v_1, v_2) \leqslant q(v_1, v_1) = 0$.
- For $v_1 < v_2 < \alpha$, this is implied by (i): $\int_{v_1}^{v_2} q(v_1, v) dv \leqslant \int_{v_1}^{v_2} q(v,v) dv = 0$.

Graphically, (ii) forces no-trade in the lower triangle, and (i) forces no-trade in the upper triangle in the unshaded region of Figure 3b. Note also that having trade whenever $\max\{v_1, v_2\} \geqslant \alpha$ increases surplus, but keeps the buyer's payoff at the same level; thus the seller's profit is bounded above by

$$\mathbb{E}\left[v_2 \cdot \mathbb{1}\left(\max\{v_1, v_2\} \geqslant \alpha\right)\right] - 1/2 \times \int_\alpha^1 (1-v) dv = 1/2 \times \alpha[1-\alpha] + 1/2 \times \int_\alpha^1 \frac{3v^2}{2} dv.$$

[5] The standard approach in the literature is to maximize the objective while ignoring (i) and (ii). In that case the first-order optimum turns out to be $q^{foa} = \mathbb{1}\left(v_1 \neq v_2 \vee v_1 = v_2 > 1/2\right)$, which violates both (i) and (ii).

It is routine to verify that the upper bound is maximized at $\alpha^* = \frac{\sqrt{10}-2}{3}$ which yields exactly the same profit as the optimal pricing strategy described above. Thus, having trade everywhere except the lower square of dimension α is optimal, which is the shaded region of Figure 3b. The optimal allocation rule is given by $q(v_1, v_2) = \mathbb{1}\left(\max\{v_1, v_2\} \geqslant \frac{\sqrt{10}-2}{3}\right)$. It can be further shown that for the two-period model, the profit generated by this optimal deterministic contract *cannot* be improved by allowing for randomization.[6] Now, we take these insights from the two-period model to a continuous time model with Poisson arrivals.

3 Primitives

A seller (she) wants to sell one unit of a timed good (or service) to a buyer (he). The good is timed in the sense that it has a fixed date of consumption T. Time is continuous and indexed by $t \in [0, T]$. For simplicity, we assume no discounting, and zero cost of production for the seller. The buyer's valuation for the good follows a stationary Markov renewal process: $V_t = X_{N_t}$, where N_t is a Poisson process with intensity λ and $\{X_n\}_{n=1}^{N_T}$ is a sequence of iid samples from a distribution F on $[0, 1]$ with a well defined density f.

Physically, the process works as follows: a value V_0 is drawn at time zero from F. Then at each instant in time, $V_t \in [0, 1]$ is the buyer's value (or "demand") for the good that he will consume at date T. The value can either stay the same or change with the arrival of some news. In the latter case, modeled as the arrival of a Poisson shock, the value is redrawn from the distribution F. If say two shocks arrive between 0 and T, then the value is redrawn twice. We will write V^t or $V^{[0,t]}$ in place of the continuous vector $(V_s)_{s=0}^t$, and $V^{[t,T]}$ for the "future" set of values $(V_s)_{s=t}^T$. The final realized value V_T is the *actual payoff* the agent gets from consuming the good. The structure of changing values prior to the realization of the actual payoff follows the sequential screening literature.[7]

At the outset, we consider the following set of pricing instruments: At the initial date the seller asks the buyer to make an upfront payment M and offers a menu of time-dependent prices $\mathbf{p} = (p_t)_{t=0}^T$.[8] The buyer can either opt out or pay the upfront payment; the payment of M grants him the right to buy the good at any of the future prices. For example, if the buyer decides to purchase the good at time t, he will make a payment of p_t (in addition to M) to the seller in return

[6] The proof of the claim that for the two-period model the optimal contract is deterministic is available from the authors upon request. The paper restricts the allocation space to be deterministic, except in Section 7.

[7] All papers in dynamic pricing and mechanism design that study profit maximization assume some special structure: (i) time horizon, eg. Courty and Li [2000] and Esö and Szentes [2007] assume two periods, (ii) type space, eg. Courty and Li [2000] assume two-types in the first period to motivate refund contracts, Akan, Ata, and Dana Jr. [2015] assume the first period type is the "time" at which the buyer is informed of her value, (iii) number of goods, eg. all of sequential screening literature looks at the sale of one good at a fixed time of consumption, and (iv) stochastic process, eg. Deb [2014], and Garrett [2016] assume Poisson arrivals, Boleslavsky and Said [2013] assume a geometric AR(1) process, and Courty and Li [2000] and Esö and Szentes [2007] put various assumptions on the Markov process to ensure validity of the local approach. We maintain generality along time horizon and type space but restrict attention to the sale of a single good and a specific Markov process, viz. Poisson. We believe this model allows for both tractability, and non-trivial economic predictions and intuition.

[8] As with values, will denote a the continuous vector of prices till time t as $p^t = (p_s)_{s=0}^t$, and the entire menu of prices is succinctly expressed as $\mathbf{p} = p^T = (p_t)_{t=0}^T$. The future set of prices will be denoted by $p^{[t,T]} = (p_s)_{s=t}^T$.

for which he is assured the delivery of the good at time T. We assume that the physical payment and consumption both happen at T.[9] If the buyer does not buy the good till T, no trade happens and the upfront payment is lost as a sunk cost for the buyer.

As is the norm in price discrimination, this model can be seen as a single seller and single buyer interaction where the latter's valuation is drawn from a known distribution, or equivalently, behind the veil of the law of large numbers, it can also be viewed as a single seller and many buyers interaction where the distribution determines the size of market demand.

4 A dynamic pricing strategy

For a fixed menu of prices $\langle M, \mathbf{p}\rangle$, the buyer's strategy can be described as an *optimal stopping problem*: Modulo the upfront payment, the gain from stopping at time t is described by $G_t(v) := \mathbb{E}[V_T | V_t = v] - p_t$. The buyer can always refuse to trade upon stopping, thus the effective gain is $(G_t(v))^+$, where $a^+ = \max\{0, a\}$.

It is standard practice to formulate the solution to such questions as a Markov decision problem. At any point t, since we only need to keep track of the stochastic evolution of the value process $V^{[t,T]}$ and the set of remaining prices $p^{[t,T]}$, the buyer's strategy can be shown to be Markov in the current value V_t and time t. The value function of the buyer W_t at any time t is then given by

$$W_t(v) := \sup_{\tau \in [t,T]} \mathbb{E}\left[(G_\tau(V_\tau))^+ \mid V_t = v\right].$$

where "sup" is taken over all stopping times larger than t.

Internalizing the aforementioned optimal response of the buyer, the seller's optimization problem consists of the choice of upfront payment M and price path $(p_t)_{t=0}^{T}$ to maximize her expected profit. A technical point to keep in mind is that in solving the seller's problem, we will break buyer's indifference to stopping or not in favor of the seller. This is because the seller can reduce prices by a small amount and get a unique implementation with profit arbitrarily close to the optimal one.

4.1 Optimal dynamic pricing

For starters, it is intuitive that the buyer's response should be a threshold strategy: as a function of the current value V_t and future path of prices $p^{[t,T]}$, the buyer devises a threshold $\alpha_t\left(p^{[t,T]}\right) \in [0,1]$ such that if $V_t \geqslant \alpha_t$, buy, else wait. The seller can then optimize over the threshold responses. We show that these thresholds can be derived in closed form, and in fact they have a simple structure.

Theorem 1. *There exists $\alpha^* \in [0,1]$ such that an optimal pricing strategy, $\langle M^*, \mathbf{p}^*\rangle$, is as follows:*

$$M^* = \mathbb{E}\left[(V_T - \alpha^*)^+ | V_0 = 0\right], \quad p_t^* = \alpha^* - \left(1 - e^{-\lambda(T-t)}\right)\int_0^{\alpha^*} F(v)dv.$$

[9] Since there is no discounting, fixing the physical payment at time T is without loss of generality. We could instead assume that that M is paid upfront, p_t is paid at time t, and the good is consumed at time T.

The buyer always pays M^* upfront and purchases the good at time t for price p_t^* iff $V_t \geqslant \alpha^* > \max_{s<t} V_s$.

Theorem 1 establishes that the optimal pricing strategy and the buyer's best response to it can be parametrized by a single variable, viz. the final price of (potential) trade: $p_T = \alpha$; we will term it the *spot price*. For any spot price α, the seller will optimally (backward) construct the path of prices and upfront payment as listed above. The buyer in turn will always make the upfront payment, and will stop and trade at price p_t, the first instant t at which his value is above the threshold α. In case $\max_{t \geqslant 0} V_t < \alpha$, there is no trade.

The rough intuition for why the threshold reduces to a constant is as follows: Suppose the threshold is (locally) decreasing. Then the buyer will always prefer to wait and learn his future valuations before trading. The seller can then increase her profit by flattening the threshold for those time periods, and as we saw in the two period example, increase her profit through greater price discrimination. Suppose instead that the threshold is (locally) increasing. Then, replacing this with an appropriately chosen constant threshold increases the area of trade, surplus produced from which can be extracted using upfront payment.

The formal proof presented in the appendix has four steps. The heart of the argument is the first step which relabels prices as a function of time dependent thresholds. Exploiting the intuition stated above, the thresholds are then replaced by a constant α. The second step argues the upfront payment must exactly equal the expected utility of the lowest initial type, $V_0 = 0$. These properties are then used to construct an upper bound on the seller's profit. The third step shows that $\langle M^*, \mathbf{p}^* \rangle$ defined above achieve this upper bound. Finally, the fourth step simplifies the seller's expected profit generated from these prices to be:

$$\Pi(\alpha) := \underbrace{\alpha\,[1 - F(\alpha)]}_{\text{spot pricing}} + \underbrace{\left(1 - e^{-\lambda T}\right) \int_{\alpha}^{1} [1 - F(v)]\, dv}_{\text{upfront payment}} + \underbrace{\left(1 - e^{-\lambda[1-F(\alpha)]T}\right) \int_{0}^{\alpha} v dF(v)}_{\text{dynamic market segmentation}}. \qquad (\star)$$

Equation ($\star$) shows that the seller's optimization problem has been reduced to choice of a single variable: a threshold given by $\alpha^* := \arg\max_{\alpha \in [0,1]} \Pi(\alpha)$.[10] In what follows, we provide a classical price-theoretic explanation of how each of three individual components in the equation stack up to constitute the seller's profit.

4.2 The simple economics of dynamic pricing[11]

Suppose that the seller ignores dynamics completely and offers the good for a spot price α. The buyer will buy the good if and only if his final value of consumption is larger than α, that is $V_T \geqslant \alpha$. Note that from an ex ante perspective V_T is (unconditionally) distributed according to F, thus the trade will happen with the following probability: $\mathbb{P}(V_T \geqslant \alpha) = 1 - F(\alpha)$.

We will regard this probability as the buyer's inverse demand function: $D(p) := 1 - F(p)$,

[10] In Corollary 1 we show that α^* is unique under standard restrictions on value distribution.

[11] The term "simple economics" is used as a homage to Bulow and Roberts [1989], which provides a price theoretic interpretation for optimal auctions (and more generally mechanism design) problem initiated by Myerson [1981].

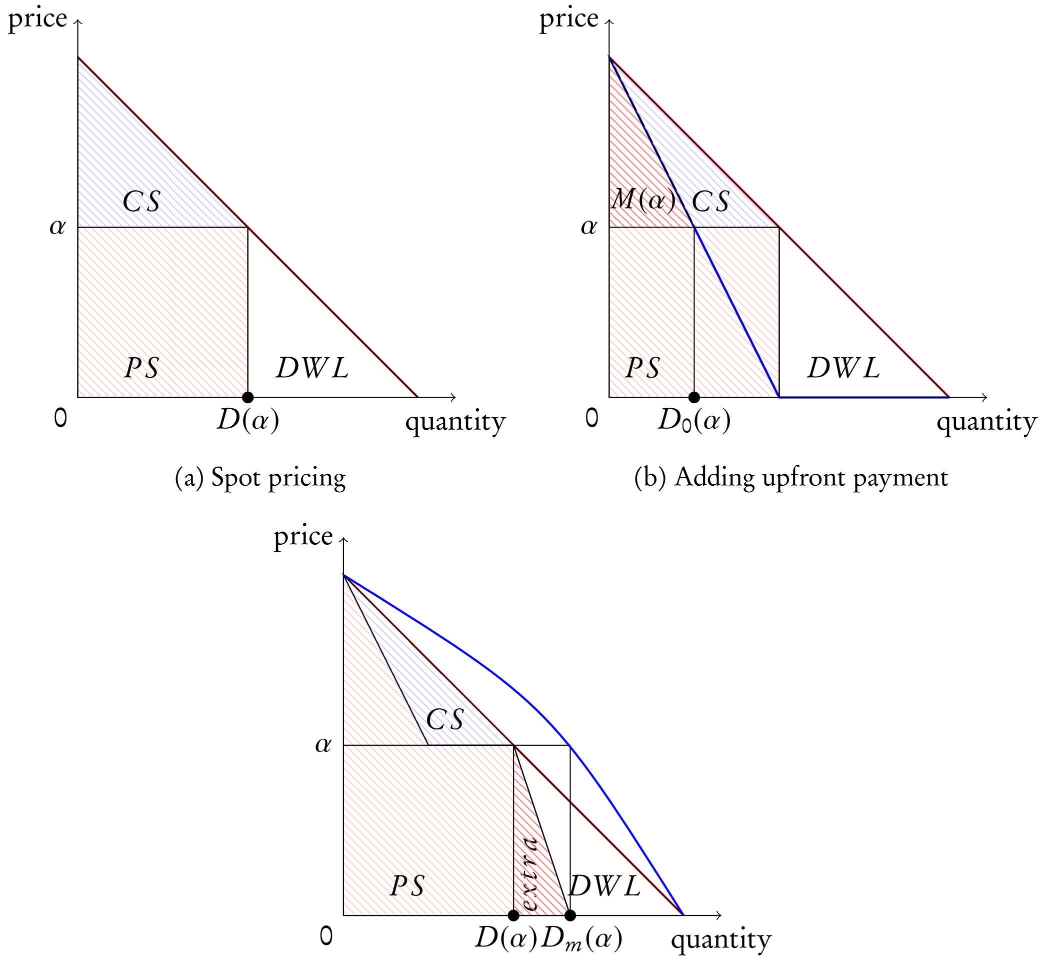

(a) Spot pricing

(b) Adding upfront payment

(c) Adding dynamic prices

Figure 4: Decomposition of total surplus from trade into CS, PS and DWL

where $D(p)$ is quantity demanded at price p.[12] Figure 4 depicts the buyer's demand function with "quantity" on the x-axis and price on the y-axis. In Figure 4a, for a fixed spot price α, the red area captures the seller's expected profit (also known as producer surplus, PS), the blue area captures the buyer's expected payoff (consumer surplus, CS) and the rest unshaded part is the deadweight loss (DWL) due to no trade, whenever $V_T < \alpha$. The seller's expected payoff from this static pricing strategy is $\alpha[1 - F(\alpha)]$, which is the first term of Equation ($\star$).[13]

In this static pricing mechanism, even the buyer with the lowest possible initial value, $V_0 = 0$, has a positive probability of ending with a final value, V_T, greater than α. This leaves a baseline positive expected surplus for all types of buyers. As a first step towards dynamic pricing, the seller can extract this consumer surplus through a positive upfront fee that still induces the buyer to always participate. Specifically, the seller can choose the upfront payment $M(\alpha)$ to be the expected

[12] As in standard demand theory, we can also think of the seller facing a population of buyers in which case $D(p)$ is the size of the market at price p.

[13] Two things to note about the figures: First, since the cost of production for the seller is assumed to be zero, trade is always efficient. As a consequence, the entire area of no trade forms the deadweight loss. Second, linearity of D is assumed for simplicity of exposition, this would of course be exact if F is uniform.

payoff of the lowest value buyer:

$$M(\alpha) := \mathbb{E}\left[(V_T - \alpha)^+ | V_0 = 0\right] = \left(1 - e^{-\lambda T}\right) \int_\alpha^1 [1 - F(v)]\, dv.$$

that is exactly the second term in Equation ($\star$).

The inverse demand function conditional on $V_0 = 0$ is $D_0(p) := \mathbb{P}(V_T \geqslant p | V_0 = 0)$, it is the blue curve in Figure 4b. The area under D_0 above the price line α, viz. the consumer surplus corresponding to D_0, is given by $M(\alpha)$. The original consumer surplus is now split into two parts: $M(\alpha)$ and the rest.[14] By construction, the seller can extract $M(\alpha)$ from the buyer without distorting his participation decision at time $t = 0$. Therefore, in Figure 4b, $M(\alpha)$ represents the transfer of an erstwhile component of consumer surplus to what is now a part of the producer surplus.

The pattern of trade is unaltered when the seller asks for $M(\alpha)$ upfront. A natural next question is this: *can the seller sell the good before the terminal date, thus decrease DWL, and increase her profit*? One strategy is to offer lower prices initially to induce the buyer with high enough values, say $V_t \geqslant \alpha_t$, to purchase the good early. The optimal price sequence is pinned down by a time independent threshold: make all buyer types $V_t \geqslant \alpha$ indifferent between making a purchase at t and waiting until the terminal date. Thus, trade take place early whenever $\max_{t\geqslant 0} V_t \geqslant \alpha$.

In this final step, increased region of trade moves some part of the erstwhile DWL to the producer surplus. The dynamic segmentation of the market can be visualized in Figure 4c; the blue curved line depicts a new demand function given by $D_m(p) = \mathbb{P}\left(\max_{t\geq 0} V_t \geq p\right)$.[15,16] The change in DWL can be expressed as follows:

$$\text{"extra"} = \underbrace{\mathbb{E}\left[V_T \cdot \mathbb{1}\left(V_T < \alpha \leqslant \max_{t\geqslant 0} V_t\right)\right]}_{\text{difference in static and dynamic DWL}} = \mathbb{P}\left(\max_{t\geqslant 0} V_t \geqslant \alpha | V_T < \alpha\right) \cdot \underbrace{\mathbb{E}\left[V_T \mathbb{1}\left(V_T < \alpha\right)\right]}_{\text{static DWL}}.$$

Here $\zeta := \mathbb{P}\left(\max_{t\geqslant 0} V_t \geqslant \alpha | V_T < \alpha\right)$ measures the fraction of trades that can be "recovered" by using dynamic pricing when the final spot price p_T is fixed at α. Using Bayes rule it can be shown that $\zeta = \frac{D_m(\alpha) - D(\alpha)}{1 - D(\alpha)}$, so the area [$\zeta$ · static DWL], depicted in Figure 4c as "extra", is transferred from DWL to PS due to dynamic pricing.[17] The magnitude of this extra profit for the seller exactly equals the third term in Equation ($\star$).

[14] It is easy to see that $\mathbb{P}(V_T \geqslant \alpha | V_0 = 0) < \mathbb{P}(V_T \geqslant \alpha) = 1 - F(\alpha)$, hence the new conditional demand function lies below the old unconditional one.

[15] $D_m(p)$ lies above the static demand curve $D(p)$, since $D_m(p) := 1 - F(p)e^{-\lambda[1-F(p)]T} \geqslant 1 - F(p) = D(p)$.

[16] Note that maximal ex post surplus is still V_T and *not* $\max_{t\geqslant 0} V_t$.

[17] $\zeta = \mathbb{P}\left(\max_{t\geqslant 0} V_t \geqslant \alpha \middle| V_T < \alpha\right) = \frac{\mathbb{P}\left(V_T < \alpha \leqslant \max_{t\geqslant 0} V_t\right)}{\mathbb{P}(V_T < \alpha)} = \frac{D_m(\alpha) - D(\alpha)}{1 - D(\alpha)}$.

4.3 Understanding the pricing strategy

In this section, we discuss the structure of the threshold α^*, and associatedly the comparative statics of the dynamic pricing mechanism with respect to the primitives of the model.

From Equation ($\star$), the first-order condition determining α^* is given by:

$$\frac{1 - F(\alpha)}{\alpha f(\alpha)} = e^{\lambda T \cdot F(\alpha)} \left(1 + \lambda T \cdot \int_0^{\alpha} \frac{v dF(v)}{\alpha}\right).$$

This equation only admits interior solutions, and the solution is unique under standard assumptions on the distribution of valuations, for example, monotonicity of the inverse hazard ratio. Moreover, it can be noted that time T and rate of transition λ enter symmetrically in this expression. Thus, what matters for comparative statics is the normalized time λT. to which we now turn.[18]

Corollary 1. *The optimal threshold satisfies the following properties:*

(a) $\alpha^* \in (0, 1)$ *and it is unique whenever* $v \mapsto \frac{f(v)}{1-F(v)}$ *or* $v \mapsto v f(v)$ *are nondecreasing.*

(b) α^* *converges to the optimal static spot price as* $\lambda T \to 0$.

(c) α^* *is strictly decreasing in* λT *with* $\lim_{\lambda T \to \infty} \alpha^* = 0$.

Part (b) states that as $T \to 0$, so that the time horizon shrinks, or as $\lambda \to 0$, so that values become perfectly persistent, the optimal contract coverages to its static benchmark. So, for example, with a uniform distribution over the unit interval, the best the seller can do is post a price of $1/2$. Part (c) says that α^* is positive and strictly decreasing in normalized time. As the date of consumption of the object goes further into the future, the (ex ante) probability of trade goes up. In particular, in the limit the good is always sold: when $T \to \infty$ the initial informational advantage of the agent goes to zero and when $\lambda \to \infty$ the stochastic process becomes iid. In both cases the efficient contract is optimal and the seller can extract all surplus as profit using the upfront fee.

We view the comparative static result in part (c) as an indication that our model should be considered as having maximum empirical relevance for smaller values of T: a flight that has to be taken in the next few months or a hotel booking in the next few weeks. Conceptually though it tells us that the ability to commit early to dynamic prices allows the seller to more effectively segment the market, thereby increasing total surplus, and extracting a larger fraction of it as profit.

4.4 Qualitative implications

The model here is of course stylistic, and hence the emphasis is mostly on learning the key economic forces at play. However, it does produce some predictions that can have potential empirical relevance. In this section we present two such qualitative results, and in the next section we discuss an estimation strategy to back out fundamentals from observables.

[18] The main reason for using a continuous time model is precisely that the expression for total profit given by Equation ($\star$) is simpler, solutions calculable in closed form, and comparative statics with respect to t and λ easier to report.

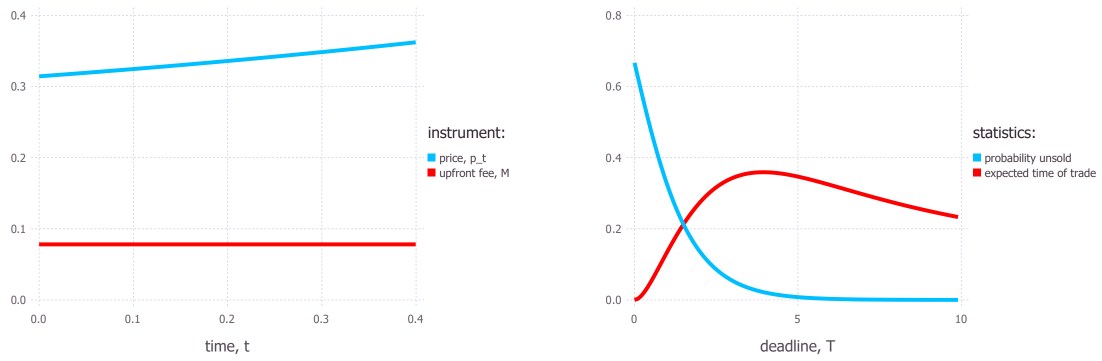

(a) Prices (b) Probability of sale

Figure 5: Evolution over time for $F(v) = v$, $\lambda = 1$.

The first prediction of the model is that the optimal price schedule is increasing over time. Theorem 1 states that optimal prices satisfy the following identity:

$$p_T^* - p_t^* = \left(1 - e^{-\lambda(T-t)}\right) \int_0^{\alpha^*} F(v)dv.$$

It can be checked that this difference between p_T^* and p_t^* is strictly decreasing in time. Thus, as shown in Figure 5a, prices for trade increase steadily from some initial value p_0^* to the final spot price p_T^*.[19]

Corollary 2. *$p_T^* - p_t^*$ is strictly decreasing in t.*

Such pricing schemes have been refereed to in the literature as *advance purchase discounts*. Typically a static or two-period model is explored where the seller offers two prices which sorts buyers into two categories. For example, Dana Jr. [1998] considers a static setup with competing sellers, and two types of buyers, each of whom have idiosyncratic uncertainty in demand. Low valued buyers with higher uncertainty buy early and high valued buyers with lower uncertainty buy later.[20]

The advance purchase discount in our setup works differently. There is a continuous price path and hence a steadily declining discount. The buyers are sorted along the first instant at which their value estimate crosses α^*. Those whose value estimate is above α^* early buy at a lower price, and those whose value estimate is below α^* initially but goes above the threshold later, buy then at a higher price. For example, die hard fans of an artist know early that they want to go to her concert, and new fans may decide to buy the ticket closer to the date when late information peaks their interest. Of course, a die hard fan may receive some news later on that he cannot be in town at the time of of the concert, but it is a trade-off worth making given the advance purchase discount and his current value estimate.

[19] Note also that the lower horizontal line in Figure 5a represents the upfront payment, M^*.

[20] See also Gale and Holmes [1993] for an analysis for advance purchase discounts in a static monopoly setup with capacity constraints.

The second prediction of the model is regarding the probability or the fraction of completed trades. Suppose the analyst has access to a large number of independent trades in identical environments. It is obvious that the fraction of trades completed by time t is increasing in t. Moreover, the random variable that captures this is given by $v_t := \max_{s \leqslant t} V_s$, and its distribution by: $\hat{F}_t(v) = F(v)e^{-\lambda[1-F(v)]t}$. The probability that the good (i) is sold by time t is $1 - \hat{F}_t(\alpha^*)$, and (ii) remains unsold is $\hat{F}_T(\alpha^*)$.

Relatedly, suppose the analyst only has access to the universe of completed sales– viz, trades are recorded but no trades are not. What can the analyst do then? The expected time of sale, given it happens, can be computed as

$$\int_0^T t d \left[\frac{1 - \hat{F}_t(\alpha^*)}{1 - \hat{F}_T(\alpha^*)} \right]$$

The analyst can then use this statistic to deduce a basic test of the model. These facts are documented in the next corollary.

Corollary 3.

(a) For a fixed T, the probability of no sale till time t, given by $\hat{F}_t(\alpha^*) = F(\alpha^*)e^{-\lambda[1-F(v)]t}$, is strictly decreasing in λt.

(b) The total probability of no sale, $\hat{F}_T(\alpha^*)$, is strictly decreasing in λT with $\lim_{\lambda T \to \infty} \hat{F}_T(\alpha^*) = 0$.

(c) Expected time of sale, conditional on sale, given by $E_{\lambda,T} = \int_0^T t d \left[\frac{1 - \hat{F}_t(\alpha^*)}{1 - \hat{F}_T(\alpha^*)} \right]$, is non-monotone:

$$\lim_{\lambda T \to 0} E_{\lambda,T} = \lim_{\lambda T \to \infty} E_{\lambda,T} = 0.$$

The first and second part are the exact counterpart of Corollary 1c: the probability of no trade is strictly decreasing in normalized time. The first part captures this at the intensive margin and the second part captures it at the extensive margin. The third part states that if the seller can vary the time at which she posts the good for sale, then the expected time at the which it is sold first goes up and then goes down.

The intuition for the third part is straightforward: Initially, when T is small, the seller wants to price discriminate by spreading the buyers into different buckets of time of purchase. However, as T become large, the private information of the buyer about a good he consumes far into the future is minimal. Hence, the seller wants to sell early, pay a small information rent, and extract most of the surplus through the upfront payment.

This non-monotonicity in expected sale time, depicted in Figure 5b, provides a simple testable implication of the model and optimality behavior by the agents. If the expected sale time is not hump-shaped, analyst can conclude that either the model is incomplete and/or the agents are not behaving optimally. In what follows, we provide a more extensive procedure to estimate fundamentals, which relies on the model specification but is valid for any menu of dynamic prices posted by the seller, optimal or not.

4.5 Estimating fundamentals

In this section, we wear the hat of an analyst who has access to the menu of prices posted by the seller, and a large set of trades by distinct buyers. Putting some structure on the data generating process, we ask: can the analyst (parametrically) estimate the distribution of valuations, (λ, F), from these observables? The attempt here is to provide an econometric grounding to the theoretical exercise, which may help connect to a burgeoning applied literature on dynamic pricing.[21]

Suppose the buyers have paid the upfront fees, M. To bring the model closer to practice, we assume that the seller updates prices at discrete time points $t = 0, \Delta, 2\Delta, \ldots, T$, where $\Delta > 0$ should be understood as a unit of observation, i.e., a day, week, etc. This sequence of prices is given by $\hat{\mathbf{p}} := \{\hat{p}_0, \hat{p}_1, \ldots, \hat{p}_{\bar{n}}\}$ where $\bar{n} := T/\Delta$.[22]

For the estimation exercise to be valid, some standard qualitative assumptions are made on economic environment: A large number of buyers enter the market (or are interested in purchasing the good) at time 0. Their initial values are drawn independently from the same distribution F. Each consecutive period, the values follow an identical and independent Markov process– they remain constant with probability $\beta := e^{-\lambda\Delta}$, and are redrawn from F with probability $1-\beta$. There are enough goods in the market so that buyers make their decisions independently.

The analyst only observes the buyers who trade, the ones who enter but do not trade are not observable. So, the dataset is composed of (i) I buyers who trade, indexed by $i = 1, \ldots, I$, (ii) the time of sale, indexed by $n_i \in \{0, 1, \ldots, \bar{n}\}$, and (iii) the price of sale, p_n.[23] Using this dataset, the analyst wants to estimate (β, F). To that end, we can first pin down the thresholds defined by the buyers' best response to the seller's pricing strategy.

Proposition 1. *Fix the deadline T, sequence of prices $\hat{\mathbf{p}}$, distribution F and persistence β. Define a sequence $\hat{\alpha} := \{\hat{\alpha}_n\}_{n=0}^{\bar{n}}$ by backward induction: $p_{\bar{n}} = \alpha_{\bar{n}}$ and*

$$\hat{p}_{n+1} - \hat{p}_n = \beta^{\bar{n}-n}(\hat{\alpha}_{n+1} - \hat{\alpha}_n) + \beta^{\bar{n}-1-n}(1-\beta)\int_0^{\min\limits_{k\geqslant n+1}\hat{\alpha}_k} F(v)dv, \quad \textit{for all } n = 0, \ldots, \bar{n}-1.$$

Let x_n be an unconditional probability of trade at time n for price $\hat{p}_n$. Then, the following holds:

a) If $\hat{\alpha}_n \neq \min\limits_{k\geqslant n} \hat{\alpha}_k$, then $x_n = 0$;

b) If $\hat{\alpha}_n$ is increasing in n, then $x_0 = 1 - F(\hat{\alpha}_0)$ and

$$x_n = \left(1 - \sum_{k=1}^{n-1} x_k\right)(1 - F(\hat{\alpha}_n))(1-\beta), \quad \textit{for all } n = 1, \ldots, \bar{n}.$$

Proposition 1 uses arguments analogous to Theorem 1 to characterize the buyer's optimal stopping strategy for any arbitrary sequence of prices $\hat{\mathbf{p}}$. It first constructs a sequence of thresholds

[21] This section introduces some new notation to define the econometric problem. A theoretically minded reader can skip the section at a first reading, without any loss of continuity.

[22] Two points to note here: First, the validity of the estimation exercise is independent of M. Second, for prices, we formally assume that the underlying continuous time price path $\mathbf{p}$ is piece-wise constant and left-continuous, that is $p_0 = \hat{p}_0$ and $p_t = \hat{p}_n$ for all $t \in ((n-1)\Delta, n\Delta]$ and $n = 0, \ldots, \bar{n}$.

[23] Since prices are piece-wise constant and left-continuous, buyers optimally purchase only at times: $t = 0, \Delta, \ldots, \bar{n}\Delta$.

$\hat{\alpha}$ such that the good is bought for $\hat{p}_n$ only if the value estimate at time n is above $\hat{\alpha}_n$ and $\hat{\alpha}_n = \min_{k\geqslant n} \hat{\alpha}_k$. In part a, it argues that there is no trade at time n if $\hat{\alpha}_n \neq \min_{k\geqslant n} \hat{\alpha}_k$; the price $\hat{p}_n$ is too large and the buyer has a strict incentive to wait. And, in part b, it establishes that if $\hat{\alpha}_n$ is increasing, then the trade probabilities $\mathbf{x} := \{x_n\}$ are uniquely identified by forward induction as a function of $\hat{\alpha}$.[24]

Now, we collect all the pieces to show how the analyst can estimate the stochastic process, and hence the size of the market, from the aforementioned dataset. For a sufficiently large sample of trades, the strong law of large numbers delivers that the fraction of trades at price $\hat{p}_n$ almost surely converges to the probability of trade, conditional on having a trade:

$$s_n := \frac{\sum_{i\in I} \mathbb{1}(n_i = n)}{I} \xrightarrow[I\to\infty]{a.s.} \frac{x_n}{\sum_{k=0}^{\bar{n}} x_k} =: \hat{x}_n.$$

Let $z := \sum_{k=0}^{\bar{n}} x_k$ be the total probability of trade, and note that $\mathbf{x} = z\hat{\mathbf{x}}$. Since the dataset does not directly report z, it has to be jointly identified with the stochastic process.

Define H to be the inverse of F, and $Y(z, \hat{\mathbf{x}}, \theta)$ as follows:

$$Y_0(z, \hat{\mathbf{x}}, \theta) := \hat{x}_0 \quad \text{and} \quad Y_n(z, \hat{\mathbf{x}}, \theta) := \frac{\hat{x}_n/(1-\beta)}{1 - z\sum_{k=1}^{n-1} \hat{x}_k}, \quad \text{for all } n = 1, \ldots, \bar{n}.$$

Invoking Proposition 1b, the thresholds can be re-written as follows:

$$\hat{\alpha}_n = H\left(1 - zY_n(z, \hat{\mathbf{x}}, \theta)\right), \quad \text{for all } n = 0, 1, \ldots, \bar{n}.$$

We now use this re-formulation of $\hat{\alpha}$ to construct an empirical analog for prices, which will then be matched with the dataset to produce a consistent parametric estimate of the distribution of values.

Suppose the stochastic process is specified parametrically as $F = F_\theta$ where $\theta \in \Theta \subseteq \mathbb{R}^{k+1}$ is a $k+1$-dimensional vector of parameters with the first entry $\theta_0 = \beta$ and k additional parameters. Note that both $z = z_\theta$ and $\hat{\mathbf{x}} = \hat{\mathbf{x}}_\theta$ depend on θ, as these are functions of the underlying stochastic process. Then, analogous to the equation in Proposition 1, the empirical analog of price vector, given by $\Gamma(z, \hat{\mathbf{x}}, \theta)$, is constructed as follows: $\Gamma_{\bar{n}}(z, \hat{\mathbf{x}}, \theta) := H_\theta\left(1 - zY_{\bar{n}}(z, \hat{\mathbf{x}}, \theta)\right)$; for $n = 1, \ldots, \bar{n}-1$,

$$\Gamma_n(z, \hat{\mathbf{x}}, \theta) := \beta^{\bar{n}-n}\left[H_\theta\left(1 - zY_{n+1}(z, \hat{\mathbf{x}}, \theta)\right) - H_\theta\left(1 - zY_n(z, \hat{\mathbf{x}}, \theta)\right)\right] + \beta^{\bar{n}-1-n}(1-\beta)\int_0^{H_\theta\left(1-zY_n(z,\hat{\mathbf{x}},\theta)\right)} F_\theta(v)dv;$$

$$\Gamma_0(z, \hat{\mathbf{x}}, \theta) := \beta^{\bar{n}}\left[H_\theta\left(1 - zY_1(z, \hat{\mathbf{x}}, \theta)\right) - H_\theta\left(1 - zY_0(z, \hat{\mathbf{x}}, \theta)\right)\right] + \beta^{\bar{n}-1}(1-\beta)\int_0^{H_\theta\left(1-zY_0(z,\hat{\mathbf{x}},\theta)\right)} F_\theta(v)dv,$$

By construction, for the true parameter $\hat{\theta}$, we have the following vector identity:

$$\mathbf{p} = \Gamma(z, \hat{\mathbf{x}}, \hat{\theta}).$$

[24] Trade probabilities fail to be unique when the buyer is indifferent between buying and waiting, which might happen when some thresholds coincide. We assume that the buyer breaks indifference in favor of the seller, i.e, buys the good. Then even for non-decreasing thresholds $\hat{\alpha}$, the trade probabilities $\mathbf{x}$ are unique and given by part b of Proposition 1.

where, of course, $z = z_{\hat{\theta}}$ and $\hat{\mathbf{x}} = \hat{\mathbf{x}}_{\hat{\theta}}$ are evaluated at the true value, $\hat{\theta}$. The following result then provides a consistent estimate of the fundamental of the data generating process, $\hat{\theta}$, under some specific assumption that we explicitly state and subsequently explain.

Proposition 2. *Fix the set of prices* $\mathbf{p}$. *Suppose that* Θ *is compact where* θ_0 *bounded away from 1. Let* $\tilde{z}, \tilde{\theta}$ *be the solution to*

$$\min_{z\in[0,1],\theta\in\Theta} \quad -\|\mathbf{p} - \Gamma(z, \mathbf{s}, \theta)\|_2^2 \quad \textit{subject to} \quad Y_0(z, \mathbf{s}, \theta) \geqslant Y_1(z, \mathbf{s}, \theta) \geqslant \ldots \geqslant Y_{\bar{n}}(z, \mathbf{s}, \theta).$$

Assume that the data generating process satisfies:

a) Trade in every period: $Y_0(z_{\hat{\theta}}, \hat{\mathbf{x}}_{\hat{\theta}}, \hat{\theta}) \geqslant Y_1(z_{\hat{\theta}}, \hat{\mathbf{x}}_{\hat{\theta}}, \hat{\theta}) \geqslant \ldots \geqslant Y_{\bar{n}}(z_{\hat{\theta}}, \hat{\mathbf{x}}_{\hat{\theta}}, \hat{\theta}) > 0$ *and* $z_{\hat{\theta}} > 0$*;*

b) Identification: $\Gamma(z_{\hat{\theta}}, \hat{\mathbf{x}}_{\hat{\theta}}, \hat{\theta}) \neq \Gamma(z, \hat{\mathbf{x}}_{\hat{\theta}}, \theta)$ *for all* $z \in [0, 1], \theta \in \Theta$ *such that* $Y_0(z, \hat{\mathbf{x}}_{\hat{\theta}}, \theta) \geqslant Y_1(z, \hat{\mathbf{x}}_{\hat{\theta}}, \theta) \geqslant \ldots \geqslant Y_{\bar{n}}(z, \hat{\mathbf{x}}_{\hat{\theta}}, \theta)$.

Then, both $\tilde{z}$ *and* $\tilde{\theta}$ *are strongly consistent estimators of* $z_{\hat{\theta}}$ *and* $\hat{\theta}$*, respectively.*

Proposition 2 describes the way to parametrically estimate the buyer's stochastic process. The first part of the result simply states the problem: it replaces $\hat{\mathbf{x}}$ with $\mathbf{s}$ in the empirical analog of prices, since the latter is observed in the dataset, and then minimizes the square of the distance between observed prices and their parametric estimates. This minimization exercise is carried out subject to the constraint that the empirical analog of thresholds, $\hat{\alpha}_n$, is decreasing over time.

Part a puts two related restrictions on the data generating process: First it ensures that thresholds for trade generated by the dataset are increasing, so the buyers' have an incentive to trade in each time interval Δ. Second, if some buyers trade in every period, the dataset is interior: $\mathbf{x} \gg 0$. This can be ensured by appropriately slicing the data set, that is, choosing Δ. Part b assumes that the solution to the minimization problem when s is replaced with $\hat{\mathbf{x}}$ produces a unique value of (z, θ) so the outcome of the exercise is a point estimate. This is routine for theoretical results on estimation. If, for instance, this assumption is violated, we will still have partial identification and interval estimates for $\hat{\theta}$.

Finally, three aspects of the estimation exercise are worth noting. First, it generates an empirical approximation of market size: If there were I trades and the probability of trade is estimated to be $\tilde{z}$, then $I/\tilde{z}$ measures the total number of potential buyers. In addition, since z is jointly estimated with θ, the analyst also has estimate of the the deadweight loss, given by $1 - \tilde{z}$. Second, our estimation strategy is parametric as it relies on matching $\bar{n} + 1$ moments to estimate $k + 2$ parameters. Non-parametric estimation would be possible if trades happen frequently (small Δ), because the number of moments would get large. In that case, one can take $\theta_1, ..., \theta_k$ to specify a probability mass function on a fine grid, and estimate the distribution non-parametrically. Third, we executed the estimation exercise for any arbitrary pricing strategy; the optimal one produces a constant threshold and satisfies the assumptions of Proposition 2.

5 Optimality: a mechanism design approach

In this section, we establish that the seller cannot achieve a profit higher than $\Pi(\alpha^*)$ through standard pricing mechanisms. More specifically, it is shown that the optimal deterministic dynamic mechanism implements the same allocation as before: trade at time t whenever $\max_{t\geqslant 0} V_t \geqslant \alpha^*$, where $\alpha^* \in [0,1]$ is threshold determined in Theorem 1.

Invoking the revelation principle, it is without loss of generality to focus on direct mechanisms. A dynamic (direct) mechanism is a history-dependent pair $\langle Q, \mathbf{P}\rangle$ such that $Q \in \{0,1\}$ is the allocation rule and $P_t \in \mathbb{R}$ is the cumulative payment at time t.[25] The agent reports his "type" at each "period"; that is, his strategy prescribes a report $\hat{V}_t \in [0,1]$ at each instance of time.[26]

A mechanism is *incentive compatible* if there is no reporting strategy which gives the buyer a strictly higher (expected) payoff than truthtelling:

$$\mathbb{E}\left[V_T Q\left(V^T\right) - \int_0^T dP_t\left(V^t\right)\right] \geqslant \mathbb{E}\left[V_T Q\left(\hat{V}^T\right) - \int_0^T dP_t\left(\hat{V}^t\right)\right] \quad \forall \hat{V}^T.$$

Fix a history $\hat{V}^{[0,t)}$ and $V_t = v$; define buyer's continuation payoff at this history $U_t\left(v|\hat{V}^{[0,t)}\right)$ as

$$U_t\left(v|\hat{V}^{[0,t)}\right) := \mathbb{E}\left[V_T Q\left(\hat{V}^{[0,t)}, V^{[t,T]}\right) - \int_t^T dP_s\left(\hat{V}^{[0,t)}, V^{[t,s]}\right)\middle| V_t = v\right].$$

Note that the buyer who misreported $\hat{V}^{[0,t)}$ in the past faces exactly the same incentive problem at time t as the buyer who happened to report $\hat{V}^{[0,t)}$ truthfully, that is the sequential rationality constraint is valid both "on and off path". Therefore, incentive compatibility can be restated as the requirement that for almost every $V^{[0,t)}$, $V_t = v$ and t,

$$U_t\left(v|V^{[0,t)}\right) \geqslant \mathbb{E}\left[V_T Q\left(V^{[0,t)}, \hat{V}^{[t,T]}\right) - \int_t^T dP_s\left(V^{[0,t)}, \hat{V}^{[t,s]}\right)\middle| V_t = v\right] \quad \forall \hat{V}^{[t,T]}.$$

Let $v^{[t,T]}$ denote the history where there are no Poisson arrivals and the type stays constant at v from t till T. The following lemma completely characterizes incentive compatibility.

Lemma 1. *A mechanism $\langle Q, P\rangle$ is incentive compatible if and only if (Env), (C) and (IM) hold for almost all $V^{[0,t)}$ and t:*

$$U_t\left(v|V^{[0,t)}\right) = U_t\left(0|V^{[0,t)}\right) + e^{-\lambda(T-t)}\int_0^v Q\left(V^{[0,t)}, w^{[t,T]}\right) dw \quad \text{a.e. } v; \tag{Env}$$

$$v \mapsto Q\left(V^{[0,t)}, v^{[t,T]}\right) \quad \text{is non-decreasing} \quad \text{a.e. } v; \tag{C}$$

$$\int_{\hat{v}}^v Q\left(V^{[0,t)}, w^{[t,T]}\right) dw \geqslant \int_{\hat{v}}^v Q\left(V^{[0,t)}, \hat{v}^{[t,t+\varepsilon)}, w^{[t+\varepsilon,T]}\right) dw \quad \text{a.e. } v, \hat{v} \quad \forall \varepsilon \leqslant T-t. \tag{IM}$$

[25] We require: (i) Q is measurable with respect to the sigma algebra generated by V^T, (ii) P_t is cadlag, adapted to the natural filtration and uniformly bounded.

[26] The process of reports is cadlag, adapted to the natural filtration, and almost surely constant with a finite number of jumps in any closed interval. Otherwise, the seller would detect a deviation.

Lemma 1 is the dynamic analog of the Myersonian characterization of incentive compatibility in static mechanism design (see Börgers [2015], Chapter 3), and is the continuous time counterpart to similar results established (in discrete time) by Pavan, Segal, and Toikka [2014] and Battaglini and Lamba [2019]. Here we build on Bergemann and Strack [2015], who state necessary conditions in a continuous time screening model, and provided sufficiency conditions to check for optimality.

Analogous to its static cousin, Lemma 1 pins down the space of allocations that can be implemented by some pricing strategy. Moreover, it shows that buyer's (expected) payoffs can be expressed as a function of allocation rule, up to a constant. Equation (Env) represents the dynamic analog of the widely used envelope formula; Equation (C) represents a monotonicity condition for constant histories, when there is no Poisson arrival after date t; and Equation (IM) is the integral monotonicity constraint that captures "global incentives" arising out of the multidimensionality of the mechanism design problem.

The next natural restriction that we impose is individual rationality. Formally, a mechanism $\langle Q, \mathbf{P}\rangle$ is *individually rational* if for almost every $V^{[0,t)}$, $V_t = v$ and t

$$U_t\left(v|V^{[0,t)}\right) \geqslant 0.$$

This restriction says that the buyer cannot be forced to continue the relationship with the principal when it is not in his own interest. In our quasi-linear setting with no (or equal) discounting, individual rationality binds only at the initial date and can be ignored thereafter. A mechanism that is incentive compatible and individually rational will be termed *implementable*.

Recall that $\Pi(\alpha^*)$ is the profit achieved by the optimal pricing mechanism (Theorem 1 and Equation ($\star$)). The main result of this section is that the seller cannot do better than $\Pi(\alpha^*)$ by using any other mechanism.

Theorem 2. *The seller's profit is at most $\Pi(\alpha^*)$ for any implementable mechanism.*

Start with a constant history $v^{[0,T]}$, where $V_0 = v$ and there are no Poisson arrivals. By Equation (C), there exists some threshold α such that trade takes place for $v \geqslant \alpha$ and not for $v < \alpha$. If we ignore Equation (IM), then using the envelope formula, the so-called first-order contract optimal produces trade whenever there is at least one arrival, and when the constant history is above α. The only no-trade region is along the constant history below α. Finally, the optimal choice of α pins down the first-order optimum. This contract is obviously not incentive compatible.[27]

In order to preserve incentive compatibility, the first-order optimal contract needs to be *ironed* along the following type of history: Consider $V^{[0,T]}$ such that $V_t < \alpha$ for all t. Since there is no trade for the history $\alpha^{[0,T]}$, in order to satisfy (IM), there must be no trade for these (pointwise) lower histories as well. Analogous to the two period example (see Figures 2c and 2d), in the "continuous hypercube" of all possible evolutions of valuations, an incentive compatible contract produces no trade in the bottom hypercube with a continuum of edges of length α. In the remaining histories $V^{[0,T]}$, where there exists t such that $V_t > \alpha$, the seller always wants trades, since

[27] Formally, the first-order optimal contract takes the form: $Q(v^{[0,T]}) = 1$ if $v \geqslant \alpha$ and $Q(v^{[0,T]}) = 0$ for $v < \alpha$, and $Q(V^{[0,T]}) = 1$ if $V_s \neq V_t$ for any $s \neq t$, where the seller then optimizes over α.

the shadow price of incentives here is zero, and the seller can extract the surplus from trade by binding the individual rationality constraint at time 0. Finally, seller chooses the optimal value of the threshold, α^*.

The prices that implement the optimal allocation can be constructed from Equation (Env). Since the optimal threshold for trade turns out to be a constant, the optimal allocation rule splits all histories into two classes, no-trade and trade, captured by $\max_{0 \leqslant s \leqslant T} V_t < \alpha^*$ and $\max_{0 \leqslant s \leqslant T} V_t \geqslant \alpha^*$, respectively. The latter class of histories is segmented into a continuum of $[0, T]$ buckets corresponding to the first instance at which the valuation jumps above α^*. The mapping of this construction into the pricing world, $\langle M, \mathbf{p} \rangle$, is then intuitive. The seller's optimal pricing strategy segments the valuation space precisely in way that when $\max_{s<t} V_t < \alpha^*$ and $V_t \geqslant \alpha^*$, trade happens at time t for price p_t^*.

As a final thought, there are three levels of private information possessed by the buyer: First, what is the initial type, second whether there is a Poisson arrival, i.e. whether the type has changed, and third, in case of an arrival, what is the new type. If the second level, the Poisson arrival, is publicly observed, then the local approach delivers the optimal contract. Information rent associated with the privateness of the initial draw is pinned down by the allocation along the constant history; this is captured by Equation (Env) at $t = 0$. Since the draw of the new type upon an arrival is independent of past types, and the fact of the arrival is now public, the buyer has no informational advantage over the seller beyond the initial type, hence no incentive constraints bind at the optimum except those captured by the envelope formula. This is also the rough intuition for why dynamics can be posed as being "irrelevant" under the assumption of sufficiency of the local or first-order approach (see Esö and Szentes [2017]), since only the initial private information matters for determining dynamic information rents. We go beyond these class of models to demonstrate the relevance of dynamics, especially in its economic content.

6 Alternate implementation: refunds

In the pricing mechanism employed thus far, $\langle M, \mathbf{p} \rangle$, M extracts consumer surplus modulo self-selection rents, and $\mathbf{p}$ draws from deadweight loss through dynamic market segmentation. While this is a largely normative exercise, a positive question can be thus posed: we might not observe upfront payments in many real world contexts such as travel tickets and hotels. We conceptualize such line of thinking by offering an alternate form of implementation through refund contracts, in the process emphasizing "two-partness" as being essential for the seller to achieve the optimum.

A *refund contract* constitutes two price paths: non-refundable and refundable. A non-refundable sale is final, in contrast a refundable sale can be returned for the prevailing refundable price. Formally, the seller commits to two price schedules $\mathbf{p^n} = (p_t^n)_{t=0}^T$ and $\mathbf{p^r} = (p_t^r)_{t=0}^T$. The buyer who bought the good at time t for p_t^r can return it at time $s \geqslant t$ and receive p_s^r. If the good is bought at price p_t^n, the sale is final and the good is consumed at date T. The following result states the optimal pricing mechanism and the buyer's best response to it.

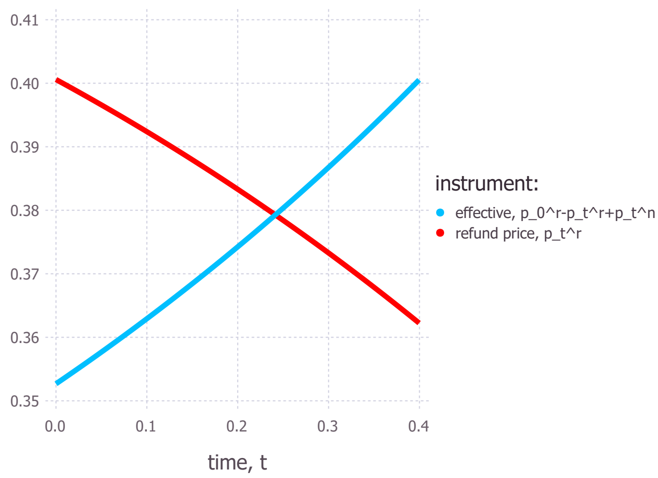

Figure 6: Optimal refund contract, $\langle \mathbf{p^r}, \mathbf{p^e} \rangle$, for $F(v) = v$ and $\lambda T = 0.4$.

Proposition 3. *Let α^* be the threshold in Theorem 1. The mechanism $\langle \mathbf{p^n}, \mathbf{p^r} \rangle$ defined by*

$$p_t^n = \mathbb{E}[V_T | V_t = \alpha^*], \quad p_t^r = \mathbb{E}[\max\{V_T, \alpha^*\} | V_t = \alpha^*]$$

achieves the same expected profit as $\langle M^*, \mathbf{p}^* \rangle$. At the start, the buyer pays the non-refundable price p_0^n. Then, he exercises the refund p_t^r at time t and buys at the non-refundable price p_t^n if $V_t \geqslant \alpha^* > \max_{s<t} V_s$. If $\max_{t \geqslant 0} V_t < \alpha^*$, the buyer claims the refund $p_T^r = \alpha^*$ at date T, and there is no trade.

It is easy to see that $p_t^r > p_t^n$ for all $t < T$, and $p_T^r = p_T^n = \alpha^*$, that is, the refundable price is always higher than the non-refundable price and both converge to the same spot price α^* at the time of consumption. Moreover, the optimal pricing strategy is structured in a way that the buyer wants to take the refundable price at time 0, and then ask for the refund and buy at the non-refundable price the first instance at which the value process jumps above α^*. In the case of no trade, that is $\max_{t \geqslant 0} V_t < \alpha^*$, the buyer ends up with a payoff of $-(p_0^r - \alpha^*)$ which is equivalent to $-M$ in terms of our benchmark pricing strategy.

Figure 6 plots the refundable price $\mathbf{p^r}$ and effective price $\mathbf{p^e}$, where the latter is defined as $p_t^e = p_0^r - p_t^r + p_t^n$, that is, the total price the buyer pays, conditional on buying the good at time t. It can be theoretically shown that $\mathbf{p^r}$ decreases with time and $\mathbf{p^e}$ increases with time. Again, the monotonicity of price paths in refund contracts generates introspectively plausible and arguably testable predictions. To be sure, there are facets of reality which may instigate the seller to have monotonic price paths that are outside the scope of the model. Even still, the model does offer some intuition that a decreasing refundable price and an increasing effective price can act as the two-parts of dynamic pricing– surplus extraction and dynamic market segmentation– that help the producer maximize her profit.[28]

[28] Another implementation that is quite intuitive is to break down the membership fees M over time in a subscription type mechanism. Therein, whenever $V_t \geqslant \alpha^* > \max_{s<t} V_s$, the buyer switches to a premium service and is delivered the good for sure at time T. Details are available upon request.

Previously the literature on sequential screening has explored option contracts as a way to achieve the optimum. Specifically, Courty and Li [2000] and Esö and Szentes [2007] consider two-period models in which the buyer selects into a menu in the first period by making a type specific payment, and then depending on the realization of type in the second period, selects from that menu. Courty and Li [2000] show that when the first period type assumes two values this has an alternate "refund" implementation. In a similar spirit, instead of seeking an an implementation that features [menus over menus.. over menus]$_T$, we look at two sequences of time dependent prices which together can be interpreted as a refund contract.

7 Gains from randomization

For the two-period model discussed in Section 2, the global optimum is deterministic. However, this is not generally true. Using the pricing approach, here we construct a stochastic improvement on the optimal contract described in Theorem 1. To provide a simple(r) intuition for what could otherwise be a complicated random mechanism, we interpret an allocation probability to be the fraction of good that is sold.

The seller uses these fractional sales and their refunds to dynamically acquire more information about the buyer's value, thereby improving the scope for price discrimination. For this section, assume that the good is divisible and can be split over time for eventual consumption at date T. An improvement over $\langle M^*, \mathbf{p}^* \rangle$ is constructed in two steps.

Step 1. In the first step, we split the good into two unequal parts of sizes $1 - z$ and z where z is a small number. The seller offers a mechanism $\langle \tilde{M}, \mathbf{p}^*, \mathbf{p^z} \rangle$, where for another small number ε,

$$\tilde{M} := z \cdot \mathbb{E}\left[(V_T - \alpha^* - \varepsilon)^+ | V_0 = 0 \right] + (1 - z) \cdot M^*.$$

is the required up front payment, M^* and $\mathbf{p}^*$ are as defined in Theorem 1, and $\mathbf{p^z}$ is given by

$$p_t^z := \alpha^* - \varepsilon - \left(1 - e^{-\lambda(T-t)}\right) \int_0^{\alpha^* - \varepsilon} F(v)dv,$$

which simply replaces α^* with $\alpha^* - \varepsilon$ in the definition of p_t^*. Upon payment of $\tilde{M}$, the buyer gets access to two price paths: he can buy the whole unit at price $z \cdot p_t^z + (1 - z) \cdot p_t^*$ or the fraction z at a marginally lower price $z \cdot p_t^z$. At any given instant, if the buyer has already bought the fraction z, he can buy the remaining part at price $(1 - z) \cdot p_t^*$.

It is easy to see that the pricing mechanism $\langle \tilde{M}, \mathbf{p}^*, \mathbf{p^z} \rangle$ is equivalent to *selling both fractions independently*. Thus, the basic ideas in Theorem 1 carry through. The buyer optimally chooses to (i) always make the upfront payment, (ii) buy the fraction "z" at the first instance when $V_t \in [\alpha^* - \varepsilon, \alpha^*)$, (iii) buy the remaining part "1-z" at the first instance when $V_t \geqslant \alpha^*$, and (iv) buy the whole unit at the first instance when $V_t \geqslant \alpha^*$.

This pricing mechanism implements the following allocation:

$$\tilde{Q}(V^T) := z \cdot \mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha^* - \varepsilon\right) + (1-z)\cdot \mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha^*\right).$$

Moreover, the net change of seller's profit $\tilde{D}(\varepsilon, z)$ by implementing $\tilde{Q}(V^T)$ instead of $Q(V^T) = \mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha^*\right)$ is given by

$$\tilde{D}(\varepsilon, z) := z \cdot \int_{\alpha^*}^{\alpha^*-\varepsilon} d\Pi(v),$$

where it can be checked: $\tilde{D}(0,0) = 0$, $\nabla\tilde{D}(0,0) = 0$, and $\nabla^2\tilde{D}(0,0) = 0$. Thus, in total for Step 1, although splitting the good and marginally increasing the trade probability is payoff-neutral, at least up to the second order, it permits the seller to relax the buyer's incentive constraints by offering more flexible terms of trade.

Step 2. In the second step we will use the fraction z as a "signal" about the buyer's value which links the sale of "z" and "$1-z$" across time. To do that, the seller will introduce a *buyback* option to the pricing mechanism in Step 1, and this new mechanism will be shown to dominate the original one in terms of the seller's profits.

Since the construction is somewhat involved, we outline the prices and their roles, formal definitions are provided in the appendix. As before, an upfront payment of $\tilde{M}$ will be charged, and the whole unit will be priced at $z \cdot p_t^z + (1-z)\cdot p_t^*$. The fraction z is now sold at price $z \cdot \hat{p}_t^z$, where $p_t^z < \hat{p}_t^z < p_t^*$. If z is bought, the buyer now has two options: he can either purchase the remaining part at a slightly lower price $(1-z)\cdot \hat{p}_t^{1-z}$ where $\hat{p}_t^{1-z} < p_t^*$ or return fraction z for refund r_t. In the latter case, after the buyback for price r_t, the good is made available again for price p_t^b, where $p_t^b < p_t^*$. This "timing" of the mechanism is summarized in Figure 7. To sum up,

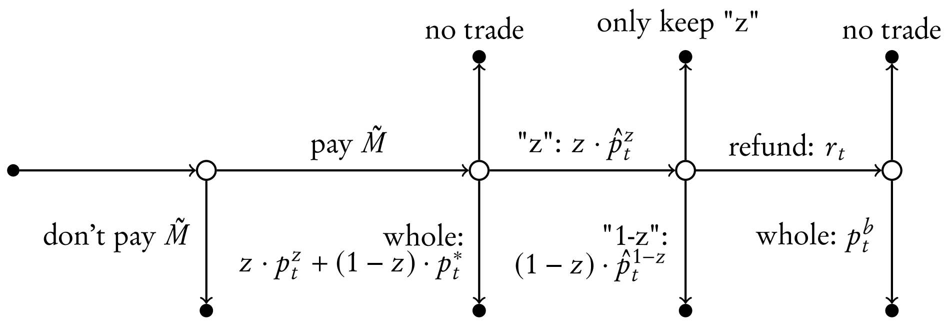

Figure 7: Timing of the fractional buyback pricing mechanism.

the main conceptual innovation here is that the seller can buyback the fraction z of the good and then condition future prices on it, thereby linking sales across time.

In the appendix, we construct a specific set of prices $\hat{p}_t^z, \hat{p}_t^{1-z}, p_t^b, r_t$, derive the allocation which it implements, and establish that the seller's profit is strictly higher for this new allocation. The next result formally summarizes the construction.

Proposition 4. *Fix three small numbers* z, ε *and* $\delta > 0$. *There exists a pricing mechanism* $\langle \tilde{M}, \mathbf{p}^*, \mathbf{p}^\mathbf{z}, \hat{\mathbf{p}}^\mathbf{z}, \hat{\mathbf{p}}^{1-\mathbf{z}}, \mathbf{p}^\mathbf{b}, \mathbf{r}\rangle$ *such that the buyer's optimal decision is to as follows:*

1. *always make the upfront payment* $\tilde{M}$*;*
2. *if fraction* z *has not been bought yet:*
   - *buy fraction* z *at price* $z \cdot \hat{p}_t^z$ *at the first instance with* $V_t \in [\alpha^* - \varepsilon, \alpha^*)$,
   - *buy the whole unit at price* $z \cdot p_t^z + (1 - z) \cdot p_t^*$ *at the first instance when* $V_t \geqslant \alpha^*$*;*
3. *if fraction* z *has been bought:*
   - *buy the remaining part* $1 - z$ *at price* $(1 - z) \cdot \hat{p}_t^{1-z}$ *at the first instance when* $V_t \geqslant \alpha^*$,
   - *claim the refund* r_t *at the first instance when* $V_t < \delta$*;*
4. *if the fraction* z *has been bought and the buyback exercised:*
   - *buy the whole unit at price* p_t^b *at the first instance when* $V_t \geqslant z\delta + (1 - z)\alpha^*$.

And, suppose that either $v \mapsto v f(v)$ *or* $v \mapsto \frac{f(v)}{1-F(v)}$ *is non-decreasing, then there exists* $\bar{\delta}$ *such that for all* $\delta < \bar{\delta}$, *the net change in seller's profit from using the this pricing mechanism, say* $\hat{D}(\varepsilon, z)$, *satisfies the following:*

$$\hat{D}(0, 0) = 0, \quad \nabla \hat{D}(0, 0) = 0, \quad \nabla^2 \hat{D}(0, 0) > 0.$$

Proposition 4 outlines a new channel of dynamic gains which, to the best of our knowledge, has not yet been explored in the literature: The seller can increase her profit by splitting the good into two parts and conditioning the prices on the history of past purchases.

To conceptualize the construction, start with the baseline pricing mechanism and suppose that the buyer has not yet bought the good. This only tells the seller that the buyer's value is below the common threshold α^*. If the seller had more precise information, she would be able to offer more flexible terms of trade, perhaps reducing the future prices to increase the probability of trade. Acquiring information is not straightforward because buyers have to self-select to provide it. The market would need further segmentation– our baseline mechanism is not rich enough.

One incentive compatible way to obtain additional information that also turns out to be profitable is to split the good into two parts, one much smaller than the other. The seller offers the smaller fraction at a lower price that can be returned for a small refund. The buyer claiming the refund signals that his value has dropped below a certain (low) threshold. The seller can then increase probability of trade by offering a lower price. It turns out that the benefit of adjusting the price is higher than the cost of offering the buyback.

8 Final remarks

This paper introduces a model of dynamic pricing where the valuation of the buyer for eventual consumption at a fixed date T evolves according to a Poisson process. A dynamic pricing strategy akin to a two-part tariff is completely characterized, where the first part extracts total surplus modulo self-selection rents, and the second part offers a continuously increasing price path that segments the market of buyers along their timing of purchase. The allocation rule implemented by this pricing mechanism is shown to be the (deterministic) optimum of the general dynamic mechanism design problem. An alternate implementation in the form refund contracts is offered. The gains from randomization are explained through a narrative of information acquisition.

From a technical perspective, the approach to identify the family of binding global incentive constraints might prove useful in other related economic environments. From an applied perspective, the restatement of the optimal pricing mechanism in terms of standard price-theoretic taxonomy helps explain the key forces in an intuitive way. Moreover, the qualitative predictions on the price and time of sale, and the estimation of fundamentals from observables can help connect what is largely a theoretical exercise here to some real world considerations of dynamic pricing.

In the appendix, two related models are explored: first allows for a linear cost of production for the seller, and second caps the number of Poisson arrivals at one. In both cases, analogous versions of Theorems 1 and 2 are presented– the optimal pricing mechanism and its global optimality. In the first case, interesting comparative statics emerge for the optimal threshold of trade with respect to the seller's cost. In the second case, the analyst can study dynamic pricing scenarios where information arrival is sparse. This latter case also shows that our approach is useful in understanding dynamic mechanisms even when the one-shot deviation principle for incentive compatibility fails.

There are many other directions in which this paper, and the research agenda more broadly can be judiciously expanded. First, we considered a model of single good with a fixed date of consumption. It would be interesting to look at repeated sales and flexible timing of consumption.[29] Second, we considered a specific stochastic process for the change in valuations which lends tractability. Exploring optimal or approximately optimal contracts for more general processes would be an obvious next step. Third, a fruitful extension is towards multiple buyers– characterizing dynamic auctions with evolving valuations. Fourth, adding stochastically arriving buyers and capacity considerations will help connect the basic analysis to the revenue management literature and potentially better connect the model to empirical applications.[30] Fifth, we assumed the information structure here to be exogenous, it would be interesting to equip the seller with an information design problem, and understand better its implications in dynamic pricing.[31] Sixth, we assumed here that the buyer does not face any financial constraints– adding a limited liability type restriction, which is realistic in many scenarios, will limit the use of upfront payment and modify the optimal allocation rule.[32]

[29] The repeated sales problem can be mapped into the dynamic Mirrleesian taxation problem, see Stantcheva [2020] for a recent review of the literature.

[30] Most empirical work in dynamic pricing assumes no commitment on the seller's side, and randomly changing market size on the buyer's side. See for example, McAfee and te Velde [2006] and Williams [2018] for airline pricing; Sweeting [2012] for ticket sales for Major League Baseball; and Chen and Sheldon [2016] study pricing for ride-sharing.

[31] Li and Shi [2017] consider a two-period model of sequential screening where the seller can disclose some additional information to the buyer; they find that full information disclosure is not optimal. Hinnosaar and Kawai [2019] show in a two -period model that a particular refund contract robustly implements the optimum across information structures.

[32] Krähmer and Strausz [2015] and Krasikov and Lamba [2019] consider such models with financial constraints in different contexts.

9 Appendix

First, missing proofs from the main text are provided. Then two related models, one with linear cost of production for the seller, another where the number of Poisson arrivals is capped, are explored.

9.1 Details and proofs for the two-period model

Here we show that the values chose for M, p and α are optimal within that class. Assuming M has been paid, start with constructing the second price p. Type v_1 buyer's expected payoff from waiting is simply $1/2 \times (v_1 - \alpha)^+ + 1/2 \times \mathbb{E}\left[(v_2 - \alpha)^+\right]$, and his payoff from trading at date 1 is given by $1/2 \times v_1 + 1/2 \times \mathbb{E}[v_2] - p$. Without loss of generality, rewrite p in terms of another threshold α':

$$p = 1/2 \times \alpha' + 1/2 \times \left(\mathbb{E}[v_2] - \mathbb{E}\left[(v_2 - \alpha)^+\right]\right).$$

Then, the expected payoff for the buyer from trading at date 1 is $1/2 \times (v_1 - \alpha') + 1/2 \times \mathbb{E}\left[(v_2 - \alpha)^+\right]$. With this re-writing of p in terms of α', the second term is the same in the payoff from accepting trade at date 1 and waiting for date 2. So, the buyer's choice is between $(v_1 - \alpha')$ and $(v_1 - \alpha)^+$, and therefore, the aggregated expected payoff of buyer of type v_1 (modulo M) is:

$$1/2 \max\{v_1 - \alpha', v_1 - \alpha\}^+ + 1/2 \times \mathbb{E}\left[(v_2 - \alpha)^+\right] = 1/2 \times (v_1 - \min\{\alpha', \alpha\})^+ + 1/2 \times \mathbb{E}\left[(v_2 - \alpha)^+\right].$$

Does he wait for α or self-select into α' today? This is the crux of our construction.

We claim that the seller would optimally induce the same threshold in both periods, i.e. $\alpha' = \alpha$. The intuition is as follows: If $\alpha' > \alpha$ then the buyer would never trade at date 1, so it is moot to have the instrument p in the first place. Thus, we can without loss of generality assume $\alpha' \leqslant \alpha$. Now, if $\alpha' < \alpha$, then, the buyer waits when $v_1 < \alpha'$ and takes the forward price when $v_1 > \alpha'$; the marginal type $v_1 = \alpha'$ is indifferent between these two options. Reducing α slightly increases the buyer's payoff in the second period uniformly for all v_1, and all of this is extracted as surplus upfront by increasing M. This strictly improves the seller's profit.

Since $\alpha' = \alpha$. the buyer waits when $v_1 < \alpha$ and for any $v_1 \geqslant \alpha$ is indifferent between both options. It is easy to see in the latter case the seller will break the indifference in favor of the forward price. Allowing the seller to break indifferences is in the standard spirit of equilibrium characterization for the seller can always give a slightly lower forward price to strictly break buyer's indifference and approximate the optimal profit.[33]

In the final step, we argue that the seller will choose $M = \mathbb{E}\left[(v_2 - \alpha)^+ | v_1 = 0\right]$ so that every buyer type pays it. At a high level if M is lower than this value the seller leaves too much surplus on the table and if it is higher than this value, then M is trying to screen buyers in the first period, which is the job of the forward price. In a construction similar to the forward price, let $M = 1/2 \times \beta + \mathbb{E}\left[(v_2 - \alpha)^+ | v_1 = 0\right] = 1/2 \times \beta + 1/2 \times \mathbb{E}[(v_2 - \alpha)^+]$, for some β, where it is easy to see

[33] If $v_1 \geqslant \alpha$, the expected profit from making the buyer wait is given by $\alpha \mathbb{P}(v_2 \geqslant \alpha | v_1 \geqslant \alpha) = \alpha[1/2 + 1/2 \times (1 - \alpha)] = 1/2 \times \alpha + 1/2 \times \alpha(1 - \alpha)$ which is strictly less than $p = 1/2 \times \alpha + 1/4 - 1/4 \times (1 - \alpha)^2$.

that $\mathbb{E}\left[(v_2-\alpha)^+ \,|v_1=0\right] = 1/2 \times \mathbb{E}\left[(v_2-\alpha)^+\right]$. Then, subtracting M from the above expression, the expected payoff of the buyer of type v_1 upon paying M is given by:

$$1/2 \times (v_1-\alpha)^+ + 1/2 \times \mathbb{E}\left[(v_2-\alpha)^+\right] - M = 1/2 \times (v_1-\alpha)^+ - 1/2 \times \beta.$$

At the optimum, we claim $\beta = 0$. It is obvious that $\beta < 0$ simply leaves surplus on the table for the buyer, so it cannot be profit maximizing. Suppose $\beta > 0$. Then, the buyer with type $v_1 \leqslant \alpha+\beta$ refuses to pay the upfront fees and opts out of the mechanism. And, all types $v_1 > \alpha+\beta$ opt-in and buy at the forward price. Thus, by allowing $\beta > 0$, the spot price (or waiting to trade in the second period) becomes redundant.

The following change in the contract then improves the seller's profit without changing the buyer's incentives: let $M = 1/2 \times \mathbb{E}\left[(v_2-(\alpha+\beta))^+\right] = \mathbb{E}\left[(v_2-(\alpha+\beta))^+ \,|v_1=0\right]$, the new spot price to be $\alpha+\beta$ and the new forward price that makes all all types $v_1 \geqslant \alpha' = \alpha+\beta$ indifferent between waiting or trading at date 1. The area of trade now strictly increases: in the earlier case the no-trade region in the (v_1, v_2) space was $[0, \alpha+\beta] \times [0,1]$ and for this modified contract the no-trade region is the square $[0, \alpha+\beta] \times [0, \alpha+\beta]$. Since, the expected payoff of the buyer remains the same and total surplus expands, the seller is strictly better off. Thus, we must have $\beta = 0$.

9.2 Proofs for the pricing mechanism $\langle M, \mathbf{p}\rangle$

Proof of Theorem 1. Recollect that the gains process is given by $G_t(v) = \mathbb{E}[V_T|V_t = v] - p_t$. The buyer can always refuse to trade upon stopping, thus the effective gain is $(G_t(v))^+$, where $a^+ = \max\{0, a\}$. And, the value function of the buyer W_t at any time t is then given by

$$W_t(v) = \sup_{\tau \in [t,T]} \mathbb{E}\left[(G_\tau(V_\tau))^+ \mid V_t = v\right].$$

where "sup" is taken over all stopping times larger than t.

Since the gain process is Markov in time and current valuation, there is no loss from using Markov strategies. It is easy to see that a strategy corresponding to $V_t = v$ can be written recursively using only constant paths:

stop and trade/not trade at $s_t(v) \in [t, T]$ whenever there is no arrival in $[t, s_t(v)]$, and

if there is an arrival at $r \in [t, s_t(v)]$, then continue with a strategy $s_r(v')$ corresponding to $V_r = v'$.

Optimizing over such recursive strategies is a much simpler task. In particular, the buyer's value function admits a natural representation where the buyer is choosing the best (deterministic) time to stop along the persistent path, i.e. before an arrival:

$$W_t(v) = \sup_{s \in [t,T]} e^{-\lambda(s-t)}[G_s(v)]^+ + \lambda \int_t^s e^{-\lambda(r-t)}\mathbb{E}\left[W_r(V_r)\right] dr.$$

We can re-write this equation to explicitly account for the two possible cases: first when the buyer stops at the the best possible time along the constant path, and second when does not stop at all along the constant path:

$$W_t(v) = \max\left\{ \underbrace{\sup_{s\in[t,T]} e^{-\lambda(s-t)}G_s(v) + \lambda\int_t^s e^{-\lambda(r-t)}\mathbb{E}\left[W_r(V_r)\right]dr}_{\text{trade at } s\,\in\,[t,T]}, \underbrace{\lambda\int_t^T e^{-\lambda(r-t)}\mathbb{E}\left[W_r(V_r)\right]dr}_{\text{no trade in } [t,T]} \right\}. \tag{†}$$

Note that $(\cdot)^+$ is omitted from G in the above equation. The reason is that the continuation value, W, is always non-negative, therefore it is suboptimal to stop and refuse to trade before the terminal date. The first term in Equation (†) is the value from having trade along the constant path in $[t, T]$, whereas the latter corresponds to no trade.

Equation (†) defines the buyer's decision problem after he makes the upfront payment. At the initial date, the buyer's type V_0 decides either to pay the upfront fee and receive $W_0(V_0) - M$ or opt out.

We prove this theorem in four steps. First, we find an upper bound on the seller's expected profit as a function of the buyer's expected payoff: in Step I- we solve for the buyer's value function, and in Step II- we consider the buyer's incentives to make the upfront payment. Then, we show that this upper bound is achieved by a pricing scheme in a specific class. Finally, we derive the optimal pricing strategy and exhibit the closed form of α^*.

Step I. Fix the pricing scheme $\langle M, \mathbf{p}\rangle$ and let $W_t(v)$ be the buyer's value function as it is defined in (†). The key to our construction is to relabel the prices. Specifically, define a number α_t by the following equation:

$$p_t = e^{-\lambda(T-t)}\alpha_t + \left(1 - e^{-\lambda(T-t)}\right)\int_0^1 w\,dF(w) - W_t(0).$$

so that the gains process can then be rewritten as

$$\begin{aligned} G_t(v) &= \mathbb{E}[V_T | V_t = v] - p_t = \\ &= e^{-\lambda(T-t)}v + \left(1 - e^{-\lambda(T-t)}\right)\int_0^1 w\,dF(w) - p_t = \\ &= e^{-\lambda(T-t)}\left(v - \alpha_t\right) + W_t(0). \end{aligned}$$

Now, by definition, the value function must dominate the gain process:

$$W_t(0) \geqslant G_t(0) = -e^{-\lambda(T-t)}\alpha_t + W_t(0).$$

Therefore, the threshold must be nonnegative, that is $\alpha_t \geqslant 0$.

In what follows we show how to solve for the spread $W_t(v) - W_t(0)$ as a function of the thresholds. There are two cases to consider. First, suppose that the buyer with the lowest valuation prefers to trade, that is the max in Equation (†) is achieved at the first term:

$$W_t(0) = \sup_{s\in[t,T]} e^{-\lambda(s-t)}G_s(0) + \lambda\int_t^s e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr \geqslant \lambda\int_t^T e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr.$$

Substitute for $G_s(v)$ and $G_s(0)$ to get $e^{-\lambda(s-t)}[G_s(v) - G_s(0)] = e^{-\lambda(T-t)}v$ for all $s \in [t,T]$, which is independent of s. Thus, both buyers of types v and 0 at time t follow the same strategy of when to stop and trade. This gives us the following expression for the spread:

$$W_t(v) - W_t(0) = e^{-\lambda(T-t)}v.$$

Next, consider the case when trade is strictly dominated for the buyer with the lowest valuation. Formally, there exists $\varepsilon > 0$ such that

$$W_t(0) = \lambda\int_t^T e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr > \varepsilon + e^{-\lambda(s-t)}G_s(0) + \lambda\int_t^s e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr \quad \forall s\in[t,T].$$

Substituting for the gain process, we have

$$\lambda\int_t^T e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr > \varepsilon - e^{-\lambda(T-t)}\alpha_s + e^{-\lambda(s-t)}W_s(0) + \lambda\int_t^s e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr \quad \forall s\in[t,T].$$

By definition, the value function $W_s(0)$ at any moment $s \in [t,T]$ is weakly higher than the value of no trade $\lambda\int_s^T e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr$, thus

$$\lambda\int_t^T e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr > \varepsilon - e^{-\lambda(T-t)}\alpha_s + \lambda\int_t^T e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr \quad \forall s\in[t,T].$$

Conclude that $\inf_{s\geqslant t}\alpha_s > \varepsilon > 0$, which implies that $W_s(0) > \varepsilon + \sup_{\epsilon\downarrow 0}\lim_{r\in[s,s+\epsilon]} G_r(0)$ for all $s \in [t,T]$. It follows that the buyer with the lowest valuation prefers to continue until the final date. Then, $W_t(0)$ can be rewritten as

$$W_t(0) = e^{-\lambda(s-t)}W_s(0) + \int_t^s \lambda e^{-\lambda(r-t)}\mathbb{E}[W_r(V_r)]\,dr, \quad \forall s\in[t,T].$$

Substitute this to (⋆) for $v > 0$ and take the sup to obtain the following estimate:

$$W_t(v) - W_t(0) = e^{-\lambda(T-t)}\left(v - \inf_{s\geqslant t}\alpha_s\right)^+.$$

Finally, note that if $\inf_{s\geqslant t}\alpha_s = 0$, then trade cannot be strictly dominated, thus $W_t(v) - W_t(0) = e^{-\lambda(T-t)}v$. The above expression captures both possible cases simultaneously. Going forward, define a unique threshold α to the inf of all thresholds: $\alpha := \inf_{t\geqslant 0}\alpha_t$.

Step II. Next, we consider the buyer's decision to make the upfront payment. It is convenient to reparametrize M as $M = W_0(0) + e^{-\lambda T}\beta$ for $\beta\in\mathbb{R}$, where note that $W_0(0)$ is the expected utility of the lowest value buyer at time 0. It is immediately clear that $\beta < 0$ leaves surplus on the table for the buyer without affecting any self-selection constraints. So, it suffices to look at $\beta \geqslant 0$. We separately study $\beta = 0$ and $\beta > 0$.

Consider $\beta = 0$. In this case, the buyer will agree to make the upfront payment irrespective of V_0. Recall that $W_t(v) - W_t(0) = e^{-\lambda(T-t)}(v - \inf_{s \geqslant t} \alpha_s)^+$. Thus, the buyer with value $V_t = v$ will agree to make a purchase at t or a moment later only if $W_t(v) = \lim_{\epsilon \downarrow 0} \sup_{s \in [t, t+\epsilon]} G_s(v)$. Equivalently: $v \geqslant \inf_{s \geqslant t} \alpha_s = \lim_{\epsilon \downarrow 0} \inf_{s \in [t, t+\epsilon]} \alpha_s$. In particular, note that there is no trade whenever $\max_{t \geqslant 0} V_t < \alpha$.

Now, the buyer's expected net payoff is $\mathbb{E}\left[W_0(V_0) - W_0(0)\right] = e^{-\lambda T}\mathbb{E}\left[(V_0 - \alpha)^+\right]$. What about total surplus? Since, $\alpha = \inf_{t \geqslant 0} \alpha_t$, the total surplus is bounded above by $\mathbb{E}\left[V_T \cdot \mathbb{1}\left(\max_{t \geqslant 0} V_t \geqslant \alpha\right)\right]$. Combining these two we obtain the following upper bound on seller's expected profit:

$$\Pi(\alpha) = \underbrace{\mathbb{E}\left[V_T \cdot \mathbb{1}\left(\max_{t \geqslant 0} V_t \geqslant \alpha\right)\right]}_{\text{maximal surplus}} - \underbrace{e^{-\lambda T}\mathbb{E}\left[(V_0 - \alpha)^+\right]}_{\text{buyer's expected payoff}}.$$

Next, let $\beta > 0$. It is easy to see that the buyer will agree to pay the upfront payment only if $V_0 \geqslant \beta + \alpha$, therefore the seller's profit is as at most

$$\underbrace{\mathbb{E}\left[V_T \cdot \mathbb{1}(V_0 \geqslant \beta + \alpha)\right]}_{\text{maximal surplus}} - \underbrace{e^{-\lambda T}\mathbb{E}\left[(V_0 - \beta - \alpha)^+\right]}_{\text{buyer's expected payoff}}.$$

The reader can verify that this is at most $\Pi(\alpha + \beta)$.

Step III. In this step, we show that for any $\alpha \geqslant 0$ there is a pricing strategy that achieves the upper bound $\Pi(\alpha)$. It is without loss to assume that $\alpha \leqslant 1$, otherwise $\Pi(\alpha) = 0$ and the problem becomes trivial. Define $\langle M, \mathbf{p} \rangle$ by

$$M = \mathbb{E}\left[(V_T - \alpha)^+ | V_0 = 0\right], \quad p_t = \alpha - \left(1 - e^{-\lambda(T-t)}\right)\int_0^\alpha F(v)dv.$$

Substitute p_t into the gain process to get:

$$p_t = e^{-\lambda(T-t)}\alpha_t + \left(1 - e^{-\lambda(T-t)}\right)\int_0^1 w dF(w) - W_t(0).$$

so that the gains process can then be rewritten as

$$\begin{aligned} G_t(v) = \mathbb{E}[V_T | V_t = v] - p_t = \\ = e^{-\lambda(T-t)}(v - \alpha) + \left(1 - e^{-\lambda(T-t)}\right)\int_\alpha^1 [1 - F(w)]\, dw \leqslant \\ \leqslant e^{-\lambda(T-t)}(v - \alpha)^+ + \left(1 - e^{-\lambda(T-t)}\right)\int_\alpha^1 [1 - F(w)]\, dw = \\ = \mathbb{E}\left[(V_T - \alpha)^+ | V_t = v\right]. \end{aligned}$$

Note that $\mathbb{E}\left[(V_T - \alpha)^+ | V_t\right]$ is the martingale that dominates the gains process. On the other hand, $\mathbb{E}\left[(V_T-\alpha)^+|V_t\right]$ is also the value from stopping only at the final date whenever $V_T \geqslant \alpha = p_T$. By definition, the value function, $W_t(V_t)$, is the smallest supermartingale dominating the gain process (cite Optimal Stopping and Freeboundary Problems by Peskir, Shiryaev (2006)), thus

$$W_t(v) = \mathbb{E}\left[(V_T-\alpha)^+|V_t = v\right].$$

Clearly, $W_0(v) \geqslant W_0(0) = M$, so the buyer is incentivized to always make the upfront payment. Moreover, the value $W_t(v)$ can be achieved by the smallest stopping time (cite Optimal Stopping and Freeboundary Problems by Peskir, Shiryaev (2006)): stop at the first instance at which $W_t(v) = G_t(v)$, that is $V_t \geqslant \alpha > \max\limits_{s<t} V_s$. Conclude that trade happens whenever $\max\limits_{t\geqslant 0} V_t \geqslant \alpha$, so the seller can obtain the profit of $\Pi(\alpha)$ by using the aforementioned pricing scheme.

Step IV. To conclude the proof, we derive the threshold α^* which maximizes $\Pi(\alpha)$. First of all, recollect that the buyer's expected payoff is simply $e^{-\lambda T}\int_\alpha^1 [1-F(v)]\,dv$

To compute the expected gains from trade, we need to introduce several auxiliary objects. Let $\tilde{T}$ be the time of the latest arrival in $[0,T]$. This is a random variable distributed on $[0,T]$ with a mass point at $t=0$:

$$\mathbb{P}\left(\tilde{T} \leqslant t\right) = e^{-\lambda T} + \int_0^t \lambda e^{-\lambda(T-s)}ds.$$

Next, denote the cdf of $\max\limits_{s\leq t} V_s$ by $\hat{F}_t(v) := F(v)e^{-\lambda[1-F(v)]t}$. It is easy to see that the expected gains from trade conditional on $\tilde{T}=0$ is $\displaystyle\int_\alpha^1 vdF(v)$, and conditional on some $\tilde{T}\neq 0$ is:

$$\mathbb{E}\left[V_T \mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha\right)|\tilde{T}\right] = \hat{F}_{\tilde{T}}(\alpha)\int_\alpha^1 vdF(v) + \left[1-\hat{F}_{\tilde{T}}(\alpha)\right]\int_0^1 vdF(v).$$

By the law of iterated expectations:

$$\mathbb{E}\left[\mathbb{E}\left[V_T\mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha\right)|\tilde{T}\right]\right] = \int_0^T \left[\hat{F}_{\tilde{T}}(\alpha)\int_\alpha^1 vdF(v) + \left[1-\hat{F}_{\tilde{T}}(\alpha)\right]\int_0^1 vdF(v)\right]\lambda e^{-\lambda(T-\tilde{T})}d\tilde{T} +$$
$$+e^{-\lambda T}\int_\alpha^1 vdF(v) = e^{-\lambda[1-F(\alpha)T]}\int_\alpha^1 vdF(v) + \left(1-e^{-\lambda[1-F(\alpha)]T}\right)\int_0^1 vdF(v).$$

After some rearrangements, the seller's profit can be expressed as in Equation (⋆):

$$\Pi(\alpha) = e^{-\lambda T}\alpha\left[1-F(\alpha)\right] + \left(1-e^{-\lambda T}\right)\int_\alpha^1 vdF(v) + \left(1-e^{-\lambda[1-F(\alpha)]T}\right)\int_0^\alpha vdF(v).$$

The optimal threshold α^* solves $\max\limits_{\alpha\in[0,1]} \Pi(\alpha)$. The first-order condition for α^* is given by

$$\frac{1-F(\alpha^*)}{\alpha^* f(\alpha^*)} = e^{\lambda T\cdot F(\alpha^*)}\left(1+\lambda T\cdot\int_0^{\alpha^*}\frac{vdF(v)}{\alpha^*}\right).$$

Corollary 1 proves that the optimal threshold α^* is well-defined, moreover it is unique under standard monotonicity assumptions. □

Proof of Corollary 1.

Parts (a), (b) Recall that the first-order condition can be written as

$$\frac{1-F(\alpha)}{\alpha f(\alpha)} = e^{\lambda T\cdot F(\alpha)}\left(1+\lambda T\cdot\int_0^\alpha \frac{vdF(v)}{\alpha}\right).$$

Clearly, the left hand side diverges to infinity, whereas the right hand side converges to zero as $\alpha\to 0$. On the other hand, the left hand side converges to zero, whereas the left hands goes to a strictly positive number as $\alpha\to 1$. Conclude that the optimal threshold is characterized by the first-order condition, thus it must lie within $(0,1)$.

Next, we show that the threshold is unique when $v\mapsto vf(v)$ is non-decreasing. Note that the left hand side is strictly decreasing in α, we claim that the right hand side is strictly increasing. To see it, differentiate $\int_0^\alpha vdF(v)/\alpha$ with respect to α:

$$\frac{d}{d\alpha}\int_0^\alpha \frac{vdF(v)}{\alpha} = f(\alpha)-\frac{\int_0^\alpha vf(v)dv}{\alpha^2} = \int_0^\alpha \frac{v}{\alpha^2}d(vf(v)) \geqslant 0$$

where we used integration by parts to obtain the last expression. Since $e^{\lambda T\cdot F(\alpha)}$ is strictly increasing, the whole right hand side is strictly increasing. By the mean value theorem, there exists unique α^* satisfying the first order condition.

Before showing that the threshold is unique when $v\mapsto\frac{1-F(v)}{f(v)}$ is non-increasing, we need to establish Part (b). The equation that defines the static threshold, say $\hat{\alpha}$, is

$$\frac{1-F(\alpha)}{\alpha f(\alpha)} = 1 \leqslant e^{\lambda T\cdot F(\alpha)}\left(1+\lambda T\cdot\int_0^\alpha \frac{vdF(v)}{\alpha}\right)$$

with equality if and only if $\lambda T=0$. Specifically, this arguments proves that α^* converges to $\hat{\alpha}$ as $\lambda T\to 0$.

Now, suppose that the inverse hazard ratio is non-increasing, then the static fixed price is uniquely pinned down as an intersection of two monotone functions, namely $\frac{1-F(\alpha)}{f(\alpha)}$ and α. It follows that α^* must be less than the static optimal fixed price, because

$$\alpha \leqslant \alpha e^{\lambda T\cdot F(\alpha)}\left(1+\lambda T\cdot\int_0^\alpha \frac{vdF(v)}{\alpha}\right).$$

We shall show that $v\mapsto vf(v)$ is non-decreasing on $[0,\hat{\alpha}]$ given the monotone inverse hazard rate, then uniqueness of α^* will follow from the argument above. Take $\beta\leqslant\alpha\leqslant\hat{\alpha}$. By monotonicity of the inverse hazard ratio, $\frac{d}{dv}(v[1-F(v)])>0$ for $v\leqslant\hat{\alpha}$, thus $\beta[1-F(\beta)]\leqslant\alpha[1-F(\alpha)]$ and

$$\frac{\alpha}{\beta}\geqslant\frac{1-F(\beta)}{1-F(\alpha)}\geqslant\frac{f(\beta)}{f(\alpha)}$$

Conclude that $\alpha f(\alpha)\geqslant\beta f(\beta)$.

Part (c). It remains to show that α^* is strictly decreasing in λT. By the way of contradiction, assume that $\alpha_1^*\geqslant\alpha_2^*$ are the optimal thresholds for (λ_1,T_1) and (λ_2,T_2) with $\lambda_1T_1<\lambda_2T_2$. Observe that the seller's profit can be rewritten in an integral form as it follows:

$$\Pi(\alpha) = e^{-\lambda T} \int_{\alpha}^{1} \left[e^{\lambda T \cdot F(v)} \left(v + \lambda T \int_{0}^{v} w dF(w) \right) f(v) - [1 - F(v)] \right] dv.$$

Then,

$$\int_{\alpha_i^*}^{\alpha_j^*} \left[e^{\lambda_i T_i \cdot F(v)} \left(v + \lambda_i T_i \int_{0}^{v} w dF(w) \right) f(v) - [1 - F(v)] \right] dv \geqslant 0 \quad i, j = 1, 2.$$

Add up two inequalities for $i = 1$, $j = 2$ and $i = 2$, $j = 1$ to obtain that

$$\int_{\alpha_1^*}^{\alpha_2^*} e^{\lambda_1 T_1 \cdot F(v)} \left(v + \lambda_1 T_1 \int_{0}^{v} w dF(w) \right) f(v) dv \geqslant \int_{\alpha_1^*}^{\alpha_2^*} e^{\lambda_2 T_2 \cdot F(v)} \left(v + \lambda_1 T_1 \int_{0}^{v} w dF(w) \right) f(v) dv,$$

which is a clear contradiction.

Since α^* is strictly decreasing in λT, it must converge as $\lambda T \to \infty$. Clearly, it cannot converge to a strictly positive number, because it will violate the first-order condition for sufficiently large λT.

□

Proof of Corollary 3. Part (a). Note that $\hat{F}_T(v)$ is strictly decreasing in λT for every fixed v. By Part (c) of Corollary 1, α^* is also strictly decreasing. It follows that $\hat{F}_T(\alpha^*)$ is strictly decreasing as a composition of two monotone functions.

Part (b). Part (c) of Corollary 1 implies that $1 - \hat{F}_t(\alpha^*)$ converges to 1, thus there will be a mass point at $t = 0$ and

$$\lim_{\lambda T \to \infty} \int_{0}^{T} t d \left(\frac{1 - \hat{F}_t(\alpha^*)}{1 - \hat{F}_T(\alpha^*)} \right) = 0.$$

On the other hand, Part (b) of Corollary 1 implies that $1 - \hat{F}_t(\alpha^*)$ converges to $1 - \hat{F}_t(\hat{\alpha})$ where $\hat{\alpha} \in (0, 1)$ is the static optimal fixed price. Since $1 - \hat{F}_0(\hat{\alpha}) = F(\hat{\alpha})$ is bounded away from zeros, the result then trivially follows:

$$\lim_{\lambda T \to 0} \int_{0}^{T} t d \left(\frac{1 - \hat{F}_t(\alpha^*)}{1 - \hat{F}_T(\alpha^*)} \right) = 0.$$

□

9.3 Proofs for estimating fundamentals

Proof of Proposition 1. The buyer's problem is a stopping problem with the gains process $G_t(v) = \mathbb{E}[V_T | V_t = v] - p_t$ where $\mathbf{p}$ is piece-wise constant and left-continuous, i.e., $p_0 = \hat{p}_0$ and $p_t = \hat{p}_n$ for all $t \in ((n-1)\Delta, n\Delta]$ and $n = 0, ..., \bar{n}$. Recall that the buyer's value function is given by

$$W_t(v) = \sup_{\tau \in [t, T]} \mathbb{E}\left[(G_\tau(V_\tau))^+ | V_t = v \right].$$

Since the gains process is right-continuous in t, the buyer's problem admits a solution. It is easy to see that the buyer optimally can stop only at the discrete time points, i.e., $t = 0, \Delta, ..., \bar{n}\Delta$. For example, stopping in $(0, \Delta)$ is dominated by stopping exactly at Δ, because the price is the same in both cases, but the buyer has a more accurate estimate of V_T in the latter case. It follows that there is no loss of generality to recast the buyer's stopping problem in discrete time: $W_{\bar{n}\Delta}(v) = (v - \hat{p}_{\bar{n}})^+$ and

$$W_{n\Delta}(v) = \max\left\{G_{n\Delta}(v), \mathbb{E}\left[W_{(n+1)\Delta}(V_{(n+1)\Delta}) | V_{n\Delta} = v\right]\right\}, \quad \text{for all } n = 0, ..., \bar{n} - 1.$$

We now solve the aforementioned stopping problem in a closed form. Similarly to the proof of Theorem xxx, the argument is based on relabeling the prices using auxiliary threshold variables. Let $\hat{\alpha}$ be defined as in the statement of the proposition, that is $\alpha_{\bar{n}} = p_{\bar{n}}$ and

$$\hat{p}_{n+1} - \hat{p}_n = \beta^{\bar{n}-n}(\hat{\alpha}_{n+1} - \hat{\alpha}_n) + \beta^{\bar{n}-1-n}(1-\beta)\int_0^{\min_{k \geqslant n+1} \hat{\alpha}_k} F(v)dv, \quad \text{for all } n = 0, \ldots, \bar{n} - 1.$$

For $n < \bar{n}$, this can be re-written as

$$\hat{p}_n = \beta^{\bar{n}-n}\hat{\alpha}_n + (1-\beta)\sum_{k=1}^{\bar{n}-n} \beta^{\bar{n}-n-k}\left(\mathbb{E}\left[V_T\right] - \mathbb{E}\left[\left(V_T - \min_{m \geqslant n+k} \hat{\alpha}_m\right)^+\right]\right).$$

We claim that the buyer's value function at time $t = n\Delta$ satisfies the following:

$$W_{n\Delta}(v) = \beta^{\bar{n}-n}(v - \min_{k \geqslant n} \hat{\alpha}_k)^+ + (1-\beta)\sum_{k=1}^{\bar{n}-n} \beta^{\bar{n}-n-k}\mathbb{E}\left[\left(V_T - \min_{m \geqslant n+k} \hat{\alpha}_m\right)^+\right].$$

The proof is by induction. First of all, note that $W_{\bar{n}\Delta}(v) = (v - \hat{\alpha}_{\bar{n}})$ and $\hat{p}_{\bar{n}-1} = \beta\hat{\alpha}_{\bar{n}-1} + (1 - \beta)\left(\mathbb{E}\left[V_T\right] - \mathbb{E}\left[W_{\bar{n}\Delta}(V_T)\right]\right)$, thus

$$\begin{aligned} W_{(\bar{n}-1)\Delta}(v) &= \max\left\{G_{(\bar{n}-1)\Delta}(v), \mathbb{E}\left[W_{\bar{n}\Delta}(V_{\bar{n}\Delta}) | V_{(\bar{n}-1)\Delta} = v\right]\right\} = \\ &= \beta\max\{v - \hat{\alpha}_{\bar{n}-1}, (v - \hat{\alpha}_{\bar{n}})^+\} + (1-\beta)\mathbb{E}\left[W_{\bar{n}\Delta}(V_T)\right] = \\ &= \beta(v - \min_{k \geqslant \bar{n}-1} \hat{\alpha}_k)^+ + (1-\beta)\mathbb{E}\left[(V_T - \hat{\alpha}_{\bar{n}})^+\right]. \end{aligned}$$

We verified our assertion for $n = \bar{n} - 1$. Next, consider n satisfying $0 < n < \bar{n}$; and suppose that

$$W_{n\Delta}(v) = \beta^{\bar{n}-n}(v - \min_{k \geqslant n} \hat{\alpha}_k)^+ + (1-\beta)\sum_{k=1}^{\bar{n}-n} \beta^{\bar{n}-n-k}\mathbb{E}\left[\left(V_T - \min_{m \geqslant n+k} \hat{\alpha}_m\right)^+\right].$$

We shall find the value function $W_{(n-1)\Delta}$. By the same argument as above, we have that

$$W_{(n-1)\Delta}(v) = \max\left\{G_{(n-1)\Delta}(v), \mathbb{E}\left[W_{n\Delta}(V_{n\Delta})|V_{(n-1)\Delta} = v\right]\right\} =$$
$$= \beta^{\bar{n}-(n-1)} \max\{v - \hat{\alpha}_{n-1}, (v - \min_{k\geqslant n} \hat{\alpha}_k)^+\} + (1-\beta) \sum_{k=1}^{\bar{n}-(n-1)} \beta^{\bar{n}-(n-1)-k} \mathbb{E}\left[\left(V_T - \min_{m\geqslant n-1+k} \hat{\alpha}_m\right)^+\right] =$$
$$= \beta^{\bar{n}-(n-1)}(v - \min_{k\geqslant n-1} \hat{\alpha}_k)^+ + (1-\beta) \sum_{k=1}^{\bar{n}-(n-1)} \beta^{\bar{n}-(n-1)-k} \mathbb{E}\left[\left(V_T - \min_{m\geqslant n-1+k} \hat{\alpha}_m\right)^+\right],$$

which completes the induction step.

We now prove the claim of our proposition. The reader can verify that at any discrete time point n, we have that

$$W_{n\Delta}(v) - G_{n\Delta}(v) = \beta^{\bar{n}-n}\left[(v - \min_{k\geqslant n} \hat{\alpha}_k)^+ - (v - \hat{\alpha}_n)\right] \geqslant 0.$$

This difference equals to if and only if $v \geqslant \hat{\alpha}_n = \min_{k\geqslant n} \hat{\alpha}_k$. Specifically, if $\hat{\alpha}_n \neq \min_{k\geqslant n} \hat{\alpha}_k$, then the buyer never stops at $t = n\Delta$, thus there is no trade for price $\hat{p}_n$, i.e, $x_n = 0$. This proves Part a) of the proposition.

As for Part b) of the proposition, suppose that $\hat{\alpha}_n < \min_{k>n} \hat{\alpha}_k$ for every $n < \bar{n}$. Then, at any time $t = n\Delta$, $W_{n\Delta}(v) - G_{n\Delta}(v) = 0$ if and only if $v \geqslant \hat{\alpha}_n$. It follows that the optimal stopping time is unique: the buyer purchases the good for $\hat{p}_n$ if he still has not bought the good and $V_{n\Delta} > \hat{p}_n$. Note that the trade happens at the outset with probability $x_0 = 1 - F(\hat{\alpha}_0)$. We now compute probability of trade at time $t = n\Delta$ with $n > 0$, which requires $V_{k\Delta} < \hat{\alpha}_{n-1}$ for all $k < n$ and $V_{n\Delta} \geqslant \hat{\alpha}_n$. Since $\hat{\alpha}_n \geqslant \hat{\alpha}_{n-1}$, conditional on not having trade in the first $n-1$ periods, trade at $t = n\Delta$ happens only if the buyer gets a new draw larger than $\hat{\alpha}_n$. This proves that $x_n = \left(1 - \sum_{k=1}^{n-1} x_k\right)(1-\beta)(1-F(\hat{\alpha}_n))$.

If $\hat{\alpha}$ is non-decreasing, then the buyer's best response is not unique, i.e., at any time $t = n\Delta$ for $n < \bar{n}$ the buyer with value $v \geqslant \hat{\alpha}_n$ is indifferent between stopping and continuing whenever $\hat{\alpha}_n = \min_{k>n} \hat{\alpha}_k$. Yet, the smallest stopping time is unique, and the trade probabilities implied by it coincide with those identified above.

□

Proof of Proposition 2. Recall that by the strong law or large numbers, each s_n is a strongly consistent estimator of $\hat{x}_{n,\hat{\theta}}$.

Fix $\varepsilon \in (0, z_{\hat{\theta}})$ and note that $Y(z_{\hat{\theta}} - \varepsilon, \hat{\mathbf{x}}, \hat{\theta}) \gg Y(z_{\hat{\theta}}, \hat{\mathbf{x}}, \hat{\theta}) \geqslant 0$. Consider a sample path for which $\mathbf{s}$ converges to $\hat{\mathbf{x}}_{\hat{\theta}}$. Then, there exists $\bar{I}^{\varepsilon}$ such that for every $I \geqslant \bar{I}^{\varepsilon}$ the set $z \in [0,1]$, $\theta \in \Theta$ with $Y(z, \mathbf{s}, \theta) \geqslant 0$ is contains $(z_{\hat{\theta}} - \varepsilon, \hat{\theta})$. Since both Γ and Π are continuous functions in all three variables, for every $I \geqslant \bar{I}^{z_{\hat{\theta}}}$ the optimization admits a solution, say $\tilde{z}$ and $\tilde{\theta}$; and for every $I \geqslant \bar{I}^{\varepsilon}$,

$$-\|\mathbf{p} - \Gamma(\tilde{z}, \mathbf{s}, \tilde{\theta})\|_2^2 \leqslant -\|\Gamma(z_{\hat{\theta}}, \hat{\mathbf{x}}_{\hat{\theta}}, \hat{\theta}) - \Gamma(z_{\hat{\theta}} - \varepsilon, \mathbf{s}, \hat{\theta})\|_2^2$$

Letting $\varepsilon \downarrow 0$, we obtain that $\Gamma(\tilde{z}, s, \tilde{\theta})$ converges to $\Gamma(z_{\hat{\theta}}, \hat{\mathbf{x}}_{\hat{\theta}}, \hat{\theta})$. By the second assumption, z and θ are point-identified: the only way for these two to match is to have $\tilde{z}$ and $\tilde{\theta}$ to converge to $z_{\hat{\theta}}$ and

$\hat{\theta}$. Since $\mathbf{s}$ is strongly consistent for $\hat{\mathbf{x}}$, so are $\tilde{z}$ and $\tilde{\theta}$. □

9.4 Proofs for the general dynamic mechanism design problem

Proof of Lemma 1. By the standard dynamic programming arguments, incentive compatibility can be rewritten using the one-shot deviation principle where the buyer with $V_t = v$ chooses a constant misreport $\hat{v}$ which he will follow for $\varepsilon > 0$ or until the first arrival. In other words, there is no loss to look at the following family of deviations parametrized by a length of time $\varepsilon > 0$:

- $V_t = v$ misreports $\hat{v} \neq v$ and continues to misreport until $\min\{t + \varepsilon, T\}$,
- the buyer switches to truth-telling at $t + \varepsilon$ if it is lower than T or right after the first arrival.

We first consider $\varepsilon > T - t$ that delivers two necessary conditions, namely (Env) and (IM). Fix $V^{[0,t)}$, $V_t = v$, then the deviation to $\hat{v} \neq v$ is unprofitable for large ε if

$$U_t\left(v|V^{[0,t)}\right) = e^{-\lambda(T-t)}\left[vQ\left(V^{[0,t)}, v^{[t,T]}\right) - \int_t^T dP_s\left(V^{[0,t)}, v^{[t,s]}\right)\right] + \int_t^T \lambda e^{-\lambda(s-t)}\mathbb{E}\left[U_s\left(V_s|V^{[0,t)}, v^{[t,s)}\right)\right] ds \geqslant$$
$$\geqslant e^{-\lambda(T-t)}\left[vQ\left(V^{[0,t)}, \hat{v}^{[t,T]}\right) - \int_t^T dP_s\left(V^{[0,t)}, \hat{v}^{[t,s]}\right)\right] + \int_t^T \lambda e^{-\lambda(s-t)}\mathbb{E}\left[U_s\left(V_s|V^{[0,t)}, \hat{v}^{[t,s)}\right)\right] ds$$

where the first term captures the case of no arrival until T, whereas the latter refers to the case of an earlier arrival. Subtract $U_t\left(\hat{v}|V^{[0,t)}\right)$ from both sides to obtain the following:

$$U_t\left(v|V^{[0,t)}\right) - U_t\left(\hat{v}|V^{[0,t)}\right) \geqslant (v - \hat{v})e^{-\lambda(T-t)}Q\left(V^{[0,t)}, \hat{v}^{[t,T]}\right).$$

The standard envelope argument implies that $U_t\left(\cdot|V^{[0,t)}\right)$ is convex, thus almost everywhere differentiable with the derivative given by $e^{-\lambda(T-t)}Q\left(V^{[0,T)}, v^{[t,T]}\right)$. Note that (Env) can be obtained by integrating this derivative from 0 to v. Moreover, convexity of $U_t\left(\cdot|V^{[0,t)}\right)$ is equivalent to (C).

Next, we look at small $\varepsilon > 0$. Again, fixing $V^{[0,t)}$, $V_t = v$, the deviation to $\hat{v} \neq v$ is unprofitable for $\varepsilon \leqslant T - t$ when

$$U_t\left(v|V^{[0,t)}\right) = e^{-\lambda\varepsilon}U_{t+\varepsilon}\left(v|V^{[0,t)}, v^{[t,t+\varepsilon)}\right) + \int_t^{t+\varepsilon} \lambda e^{-\lambda(s-t)}\mathbb{E}\left[U_s\left(V_s|V^{[0,t)}, v^{[t,s)}\right)\right] ds \geqslant$$
$$\geqslant e^{-\lambda\varepsilon}U_{t+\varepsilon}\left(v|V^{[0,t)}, \hat{v}^{[t,t+\varepsilon)}\right) + \int_t^{t+\varepsilon} \lambda e^{-\lambda(s-t)}\mathbb{E}\left[U_s\left(V_s|V^{[0,t)}, \hat{v}^{[t,s)}\right)\right] ds.$$

where the first term captures the case of no arrival until $t + \varepsilon$, whereas the latter refers to the case of an earlier arrival. Subtract $U_t\left(\hat{v}|V^{[0,t)}\right)$ from both sides to obtain the following expression:

$$U_t\left(v|V^{[0,t)}\right) - U_t\left(\hat{v}|V^{[0,t)}\right) \geqslant e^{-\lambda\varepsilon}\left[U_{t+\varepsilon}\left(v|V^{[0,t)}, \hat{v}^{[t,t+\varepsilon)}\right) - U_{t+\varepsilon}\left(\hat{v}|V^{[0,t)}, \hat{v}^{[t,t+\varepsilon)}\right)\right].$$

Using (Env), rewrite this as

$$\int_{\hat{v}}^{v} Q\left(V^{[0,t)}, w^{[t,T]}\right) dw \geqslant e^{-\lambda\varepsilon}\int_{\hat{v}}^{v} Q\left(V^{[0,t)}, \hat{v}^{[t,t+\varepsilon)}, w^{[t+\varepsilon,T]}\right) dw.$$

Clearly, if there is no profitable deviation for $\varepsilon > 0$, then all deviations are deterred for $\varepsilon' \in (\varepsilon, T-t]$ as well. Combining this observation with the above expression yields (IM).

To sum up, we described the necessary and sufficient conditions to deter any deviation at time t after observing $V^{[0,t)}$ and $V_t = v$. An incentive compatible mechanism must satisfy these only almost everywhere, though it is without loss to ask (Env), (C) and (IM) to hold pointwise. □

Proof of Theorem 2. Take any incentive compatible mechanism $\langle Q, \mathbf{P}\rangle$. By Lemma 1, this mechanism must satisfy (C) at the initial date:

$$v \mapsto Q\left(v^{[0,T]}\right) \text{ is non-decreasing.}$$

Since $Q \in \{0,1\}$, there exists $\alpha \in [0,1]$ such that

$$Q\left(v^{[0,T]}\right) = \begin{cases} 1 & v > \alpha, \\ 0 & v < \alpha. \end{cases}$$

We next derive an upper bound on seller's profit as a function α and show that it is less than $\Pi(\alpha^*)$. First of all, write the seller's expected profit as a difference between the total surplus and the buyer's ex ante payoff, that is

$$\mathbb{E}\left[V_T Q(V^T)\right] - \mathbb{E}\left[U_0(V_0)\right].$$

Using Lemma 1, specifically (Env), solve for the buyer's expected payoff as a function of payoff to $V_0 = 0$ and α:

$$\mathbb{E}\left[U_0(V_0)\right] = U_0(0) + e^{-\lambda T} \int_0^1 [1 - F(v)]\, Q\left(v^{[0,T]}\right) dv \geqslant e^{-\lambda T} \int_\alpha^1 [1 - F(v)]\, dv.$$

The inequality follows from individual rationality of the buyer with $V_0 = 0$.

Now, we bound the surplus by showing that there is no trade whenever $\max_{t \geqslant 0} V_t < \alpha$, that is $Q\left(V^T\right) = 0$. To begin, recall that any history can be represented as a finite sequence $\{(\tau_n, X_n)\}_{n=0}^{N_T}$ where τ_n is the time of n-th arrival and X_n is the value sampled at that moment. Our argument is based on induction over the number of arrivals.

Consider V^T with only one arrival, that is $N_T = 1$, and $\max\{X_0, X_1\} < \alpha$. There are two cases to look at, namely $X_0 > X_1$ and $X_0 < X_1$, because $X_0 = X_1$ has been established before.

- For $X_0 > X_1$: $Q\left(X_0^{[0,T]}\right) = 0 \geqslant Q\left(X_0^{[0,\tau_1)}, X_1^{[\tau_1,T]}\right)$ by (C).
- For $X_0 < X_1$: $\int_{X_1}^{\alpha} Q\left(w^{[0,T]}\right) dw = 0 \geqslant \int_{X_1}^{\alpha} Q\left(X_0^{[0,\tau_1)}, w^{[\tau_1,T]}\right) dw$ by (IM).

Conclude that $Q(V^T) = 0$.

By induction, suppose that there is no trade for all V^T with the number of arrivals N_T less than $K \geqslant 2$ and $\max\{X_0, \ldots, X_N\} < \alpha$. Consider V^T with $N_T = K+1$ arrivals and $\max\{X_0, \ldots, X_N\} < \alpha$, and, again, we distinguish between two cases.

- For $X_K > X_{K+1}$: $Q\left(V^{[0,\tau_K)}, X_K^{[\tau_K,T]}\right) = 0 \geqslant Q\left(V^{[0,\tau_K)}, X_K^{[\tau_K,\tau_{K+1})}, X_{K+1}^{[\tau_{K+1},T]}\right)$ by (C).
- For $X_K < X_{K+1}$: $\int_{X_{K+1}}^{\alpha} Q\left(V^{[0,\tau_K)}, w^{[\tau_K,T]}\right) dw = 0 \geqslant \int_{X_{K+1}}^{\alpha} Q\left(V^{[0,\tau_K)}, X_K^{[\tau_K,\tau_{K+1})}, w^{[\tau_{K+1},T]}\right) dw$ by (IM)

Conclude that $Q(V^T) = 0$, thus there is no trade whenever $\max_{t\geqslant 0} V_t < \alpha$.

It follows that the surplus is at most $\mathbb{E}\left[V_T \cdot \mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha\right)\right]$, which implies that the seller's profit is not higher than

$$\mathbb{E}\left[V_T \cdot \mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha\right)\right] - e^{-\lambda T}\int_{\alpha}^{1} [1 - F(v)] \,.$$

In the proof of Theorem 1, we showed that this equals to $\Pi(\alpha)$ as defined in Equation ($\star$). Of course, $\Pi(\alpha) \leqslant \max_{\alpha\in[0,1]} \Pi(\alpha) = \Pi(\alpha^*)$ that concludes the proof. □

9.5 Proofs for alternate implementation

Proof of Proposition 3. The proof is similar to Step II in the proof of Theorem 1. We again turn the buyer's decision problem into a stopping problem. Note that it is without loss to assume that the buyer gets the refundable good at $t = 0$ for p_0^r, because he can always immediately return it and get p_0^r back.

Since the buyer can stop and refuse to switch to the non-refundable good, the gain process is defined as $(\mathbb{E}[V_T|V_t = v] - p_t^n)^+ + p_t^r$. Let $W_t(v)$ be the buyer's value function at t with $V_t = v$ when he owns the refundable good. Waiting until the final date is always feasible, thus

$$W_t(v) \geqslant \mathbb{E}\left[\max\{V_T, \alpha^*\}|V_t = v\right].$$

The process on the right is a martingale and it dominates the gain process as

$$\mathbb{E}\left[\max\{V_T, \alpha^*\}|V_t = v\right] \geqslant \left(\mathbb{E}[V_T|V_t = v] - p_t^n\right)^+ + p_t^r = e^{\lambda(T-t)}v + \left(1 - e^{-\lambda(T-t)}\right)\mathbb{E}\left[\max\{V_T, \alpha^*\}\right].$$

By definition, $W_t(v)$ is the smallest supermartingale dominating the gain process, therefore $W_t(v) = \mathbb{E}\left[\max\{V_T, \alpha^*\}|V_t = v\right]$. The smallest optimal stopping time is to stop and switch at the first instance with $V_t \geqslant \alpha^*$.

It remains to show that the buyer has incentives to purchase the refundable good at $t = 0$, indeed:

$$W_0(v) - p_0^r = e^{-\lambda T}(v - \alpha)^+ \geqslant 0.$$

Observe that the pattern of trade and buyer's ex ante payoff coincide exactly with those implemented by the benchmark mechanism $\langle M^*, \mathbf{p}^*\rangle$. By the dynamic mechanism argument (see proof of Theorem 2), the seller's obtains exactly the same profit, which is $\Pi(\alpha^*)$. □

9.6 Proofs for gains from stochastic mechanisms

Proof of Proposition 4. The proof has two steps: first we define the set of prices which implements the buyer's decision described in the statement, and then we compute the seller's profit.

Step I. We work backwards: let p_t^b be given by

$$p_t^b := z\delta + (1-z)\alpha^* - \left(1 - e^{-\lambda(T-t)}\right) \int_0^{z\delta+(1-z)\alpha^*} F(v)dv.$$

The buyer who claimed the refund can either wait or purchase the whole unit for p_t^b. This can be thought as a stopping problem with the following gain process:

$$\mathbb{E}[V_T|V_t = v] - p_t^b = e^{-\lambda(T-t)}[v - (z\delta + (1-z)\alpha^*)] + \left(1 - e^{-\lambda(T-t)}\right) \mathbb{E}\left[(V_T - (z\delta + (1-z)\alpha^*))^+\right].$$

Note that waiting until the final date yields $\mathbb{E}\left[(V_T - (z\delta + (1-z)\alpha^*))^+ | V_t = v\right]$, which is a martingale dominating the gain process. Thus, the buyer's optimal value upon claiming the refund is $\mathbb{E}\left[(V_T - (z\delta + (1-z)\alpha^*))^+ | V_t = v\right]$ and stopping for $V_t \geqslant z\delta + (1-z)\alpha^*$ is the smallest optimal stopping time.

Next, we study the buyer's incentives before claiming the refund but after purchasing the fraction. Define p_t^{1-z} by

$$\hat{p}_t^{1-z} := p_t^* - \frac{z}{1-z}\left(1 - e^{-\lambda(T-t)}\right) \int_0^{\delta} F(v)dv.$$

It follows that the gain from buying the remaining part is given by

$$\mathbb{E}[V_T|V_t = v] - (1-z)\cdot \hat{p}_t^{1-z} = z \cdot \left(\mathbb{E}[V_T|V_t = v] + \left(1 - e^{-\lambda(T-t)}\right)\int_0^{\delta} F(v)dv\right) + \\ + (1-z)\cdot\left(e^{-\lambda(T-t)}(v - \alpha^*) + \left(1 - e^{-\lambda(T-t)}\right)\mathbb{E}\left[(V_T - \alpha^*)^+\right]\right).$$

Of course, the buyer can exercise the refund r_t, that is

$$r_t := z\cdot\left(e^{-\lambda(T-t)}\delta + \left(1 - e^{-\lambda(T-t)}\right)\int_0^{\delta} F(v)dv + \left(1 - e^{-\lambda(T-t)}\right)\int_0^{z\delta+(1-z)\alpha^*}[1 - F(v)dv]\right) + \\ +(1-z)\cdot\left(1 - e^{-\lambda(T-t)}\right)\int_{\alpha}^{z\delta+(1-z)\alpha^*}[1 - F(v)]dv.$$

Note that the gain from exercising refund equals to the following:

$$\mathbb{E}\left[(V_T - (z\delta + (1-z)\alpha^*))^+ | V_t = v\right] + r_t = (v - (z\delta + (1-z)\alpha^*))^+ + (1-z)\cdot\left(1 - e^{-\lambda(T-t)}\right)\mathbb{E}\left[(V_T - \alpha^*)^+\right] + \\ + z\cdot\left(e^{-\lambda(T-t)}\delta + \left(1 - e^{-\lambda(T-t)}\right)\int_0^{\delta} F(v)dv + \left(1 - e^{-\lambda(T-t)}\right)\mathbb{E}[V_T]\right).$$

The buyer's problem is again a stopping problem, though the buyer can stop either by buying the remaining part or exercising the refund. A feasible policy is to wait until the deadline.

At $t = T$, buying the remaining fraction yields $v - (1-z)\alpha^*$, whereas asking for the refund- $(v - (z\delta + (1-z)\alpha^*))^+ + z\delta$. In addition, the buyer can abstain in which case he will consume only the original fraction and receive zv. It is easy to see that the buyer strictly prefers to buy the remaining part for $v > \alpha^*$ and to ask for the refund when $v < \delta$, thus his maximal payoff is given by

$$z \cdot v + (1-z) \cdot (v - \alpha^*)^+ + z \cdot (\delta - v)^+$$

Conclude that waiting until T gives the following expected payoff:

$$z \cdot \mathbb{E}\left[V_T | V_t = v\right] + (1-z) \cdot \mathbb{E}\left[(V_T - \alpha^*)^+ | V_t = v\right] + z \cdot \mathbb{E}\left[(\delta - V_T)^+ | V_t = v\right]$$

that clearly dominates both gains: from exercising the refund and from buying the remaining fraction, thus this value is the buyer's optimal value in the stopping problem. It is routine to verify that the smallest optimal stopping time is to ask for the refund when $V_t < \delta$ and buy the remaining part with $V_t \geqslant \alpha^*$.

Finally, define $\hat{p}_t^z$ to be

$$\hat{p}_t^z := p_t^z + \left(1 - e^{-\lambda(T-t)}\right) \int_0^\delta F(v) dv.$$

Therefore, the gain from buying the fraction is simply

$$z \cdot \mathbb{E}\left[V_T | V_t = v\right] + (1-z) \cdot \mathbb{E}\left[(V_T - \alpha^*)^+ | V_t = v\right] + z \cdot \mathbb{E}\left[(\delta - V_T)^+ | V_t = v\right] - \hat{p}_t^z =$$
$$= z \cdot \left(e^{-\lambda(T-t)}(v - \alpha^* - \varepsilon) + \left(1 - e^{-\lambda(T-t)}\right)\mathbb{E}\left[(V_T - \alpha^* - \varepsilon)^+\right] + (\delta - v)^+\right) + (1-z) \cdot \mathbb{E}\left[(V_T - \alpha^*)^+ | V_t = v\right]$$

. Using the same argument as above: by waiting the buyer can guarantee himself

$$z \cdot \mathbb{E}\left[(V_T - \alpha^* - \varepsilon)^+ | V_t = v\right] + (1-z)\mathbb{E}\left[(V_T - \alpha^*)^+ | V_t = v\right].$$

which equals exactly to his value as it dominates both gain processes (buying the whole unit and buying the fraction). The smallest stopping time prescribes to self-select into the whole unit if $V_t \geqslant \alpha^*$ and purchase the fraction when $V_t \in [\alpha^* - \varepsilon, \alpha^*)$.

Observe that the buyer's payoff before making the upfront payment is the same as in the first step, thus the buyer always agrees to participate in the mechanism. Moreover, a change in the seller's profit as compared to the first case is only due to a change in surplus. In what follows we compute the net change of seller's profit and show that is strictly positive for small ε, z and δ.

Part II. Let $\hat{Q}(V^T)$ be the allocation implemented by our pricing strategy. It is convenient to represent a history V^T as a finite sequence $\{(\tau_n, X_n)\}_{n=0}^{N_T}$ where τ_n is the time of n-th arrival and X_n is the value sampled at that moment. Using these notations, the net change in a trade probability can be written as

$$\hat{Q}(V^T) - \tilde{Q}(V^T) = \mathbb{1}\left(X_0 \in [\alpha^* - \varepsilon, \alpha), X_1 < \alpha^*, \max_{t\geqslant 0} V_t \in [z\delta + (1-z)\alpha^*, \alpha^*]\right) - \\ - z \cdot \mathbb{1}\left(X_0 \in [\alpha^* - \varepsilon, \alpha), X_1 < \alpha^*, \max_{t\geqslant 0} V_t \in [0, z\delta + (1-z)\alpha^*]\right).$$

Our goal is to compute a change of the seller's profit $\hat{D}(\varepsilon, z)$, that is

$$\hat{D}(\varepsilon, z) = \mathbb{E}\left[V_T \cdot \left(\hat{Q}(V^T) - \tilde{Q}(V^T)\right)\right].$$

Note that for $X_0, \ldots, X_N$ iid random variables, the following is true:

$$\mathbb{E}\left[X_N \mathbb{1}\left(\max\{X_0, \ldots, X_N\} \in [a, b]\right)\right] = \int_a^b v f(v) F^N(v) dv + \int_a^b \int_0^v f(w) dw dF^N(x) = \\ = \int_a^b d\left(F^N(v) \int_0^v w f(w) dw\right).$$

Using the above result, we can compute $\mathbb{E}\left[V_T\left(\hat{Q}(V^T) - \tilde{Q}(V^T)\right)\right]$ by conditioning on $N_T = N$:

$$\mathbb{E}\left[V_T\left(\hat{Q}(V^T) - \tilde{Q}(V^T)\right)|N_T = N\right] = \\ = \begin{cases} 0 & \text{for } N \leqslant 1, \\ [F(\alpha^*) - F(\alpha^* - \varepsilon)]F(\delta) \int_{z\delta+(1-z)\alpha^*}^{\alpha^*} d\left(F^{N-2}(v) \int_0^v w f(w) dw\right) - & \\ -[F(\alpha^*) - F(\alpha^* - \varepsilon)]F(\delta) \int_0^{z\delta+(1-z)\alpha^*} d\left(F^{N-2}(v) \int_0^v w f(w) dw\right) & \text{for } N \geqslant 2. \end{cases}$$

Recall that N_T has a Poisson distribution with a parameter λT, thus

$$\hat{D}(\varepsilon, z) = [F(\alpha^*) - F(\alpha^* - \varepsilon)]F(\delta)\left(\int_{z\delta+(1-z)\alpha^*}^{\alpha^*} dK(v) - z\int_0^{z\delta+(1-z)\alpha^*} dK(v)\right)$$

where $K(v) = \sum_{N=2}^{\infty} \frac{(\lambda T)^N e^{-\lambda T}}{N!} F^{N-2}(v) \int_0^v w f(w) dw$.

It is easy to see that $\hat{D}(0,0) = 0$, $\nabla \hat{D}(0,0) = 0$, $\nabla^2 \hat{D}(0,0) = 0$, except

$$\frac{\partial^2}{\partial\varepsilon\partial z}\hat{D}(0,0) = f(\alpha^*)F(\delta)\left[(\alpha^* - \delta)K'(\alpha^*) - K(\alpha^*)\right].$$

We shall show that our monotonicity assumptions imply that $\alpha^* K'(\alpha^*) > K(\alpha^*)$. Indeed, for any $N \geqslant 2$:

$$vK'(v) - K(v) = \sum_{N=2}^{\infty} \frac{(\lambda T)^N e^{-\lambda T}}{N!}\left[F^{N-2}(v)\int_0^v w d(w f(w)) + \left(F^{N-2}(v)\right)' \int_0^v w f(w) dw\right]$$

In Corollary 1, we established that monotone hazard rate implies that $v \mapsto v f(v)$ is non-decreasing on $[0, \alpha^*]$, thus

$$\alpha^* K'(\alpha^*) - K(\alpha^*) \geqslant \sum_{N=2}^{\infty} \frac{(\lambda T)^N e^{-\lambda T}}{N!} \left(F^{N-2}(v)\right)'\Big|_{v=\alpha^*} \int_0^{\alpha^*} w f(w) dw > 0.$$

By setting δ sufficiently low, we ensure that $(\alpha^* - \delta)K'(\alpha^*) - K(\alpha^*) > 0$, therefore the new mechanism strictly increases the seller's profit. □

9.7 Extension: linear cost of production

Suppose it costs $c \in [0, 1]$ for the seller to produce the good. In response to the mechanism $\langle M^*, \mathbf{p}^* \rangle$, trade still happens whenever $\max_{t \geqslant 0} V_t \geqslant \alpha^*$, but the associated profit is not optimal for the seller because the buyer does not internalize her cost: trade is inefficient whenever $V_T < c$. Here we show that a simple modification restores optimality. As before, the seller asks the buyer to make an upfront payment M and in return offers time-dependent prices p_t. However, in contrast to the benchmark mechanism, the buyer can return the good at date T and recover c.

Proposition 5. *Suppose it costs $c \in [0, 1]$ for the seller to produce the good. There exists a number $\alpha^c \in [c, 1]$ such that the optimal pricing strategy $\langle M^c, \mathbf{p^c} \rangle$ is as follows:*

$$M^c = \mathbb{E}\left[(V_T - \alpha^c)^+ | V_0 = 0\right], \quad p_t^c = \alpha^c - \left(1 - e^{-\lambda(T-t)}\right) \int_c^{\alpha^c} F(v) dv.$$

The buyer always makes the upfront payment and purchases the good at time t if $V_t \geqslant \alpha^c > \max_{s<t} V_s$. Moreover, the buyer claims the refund of c at time T if and only if $V_T < c$.

Proposition 5 is the analog to Theorem 1 with added linear cost of production for the seller. Next, we show that the allocation rule selected by $\langle M^c, \mathbf{p^c} \rangle$ is the optimal (deterministic) mechanism. The threshold α^c solves $\alpha^c := \arg \max_{\alpha \in [c,1]} \Pi^c(\alpha)$ that resembles Equation ($\star$):

$$\Pi^c(\alpha) := (\alpha - c)[1 - F(\alpha)] + \left(1 - e^{-\lambda T}\right) \int_\alpha^1 [1 - F(v)] dv + \left(1 - e^{-\lambda[1-F(\alpha)]T}\right) \int_c^\alpha (v - c) dF(v).$$

The analog to Theorem 2 thus follows.

Proposition 6. *Suppose it costs $c \in [0, 1]$ for the seller to produce the good. The seller's profit is at most $\Pi^c(\alpha^c)$ for any implementable mechanism.*

A natural comparative statics question to ask is this: how does α^c change with the c? By construction, $\alpha^c \geqslant c$, and it is intuitive to conclude that it is increasing in c. In Figure 8 we plot the α^c and $\beta^c := \frac{[F(\alpha^c) - F(c)]}{[1 - F(c)]}$, the latter is a measure of inefficiency of the mechanism expressed by the rate at which the threshold increases in the "quantile space" as a function of the seller's cost. The measure β^c is motivated by the fact that efficiency prescribes trade whenever the buyer's value is above c. As can be seen from Figure 8, even though α^c increases in c, the mechanism becomes progressively more efficient. This is because for higher values of c the seller would never serve

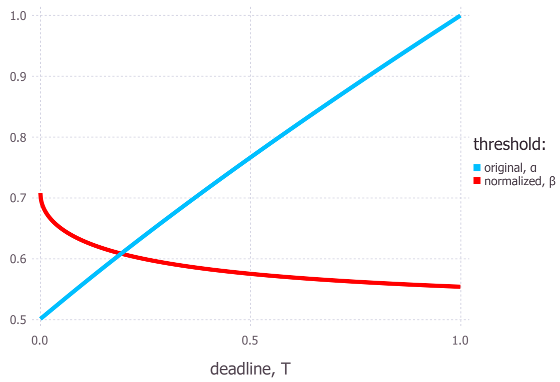

Figure 8: Thresholds α^c and $\beta^c = \frac{F(\alpha^c)-F(c)}{1-F(c)}$ as functions of c for $F(v) = \sqrt{v}$ and $\lambda T = 0.4$.

the buyer with low valuations, thus gains from price discrimination are smaller and the tradeoff between efficiency and rent extraction is less stringent.[34]

Proof of Proposition 5. The proof follows the same steps as the proof of Theorem 1. Instead of repeating our previous construction, we shall point out necessary adjustments. Clearly, the buyer's problem is again an instance of optimal stopping. Since the buyer can always recover c, a gain from stopping and trading at t with $V_t = v$ is $G_t(v) := \mathbb{E}[\max\{v, c\} | V_t = v] - p_t$.

Step I. As before, we first bound the seller's profit. Define the buyer's value function $W_t(v)$ as a solution to Equation (†), and then relabel the prices:

$$p_t = e^{-\lambda(T-t)}\alpha_t + \left(1 - e^{-\lambda(T-t)}\right)\mathbb{E}[\max\{v, c\}] - W_t(0).$$

It can be shown that $\alpha_t \geqslant c$ for all $t \geqslant 0$, and

$$W_t(v) - W_t(0) = e^{-\lambda(T-t)}\left(v - \inf_{s \geqslant t} \alpha_s\right)^+.$$

Let $\alpha := \inf_{t \geqslant 0} \alpha_t$ and write M as $M = e^{-\lambda(T-t)}\beta + W_t(0)$. The buyer's decision whether or not to pay the upfront payment is exactly as in Theorem 1. The buyer always enters the contract when $\beta = 0$, otherwise only the buyer with $V_0 \geqslant \alpha + \beta$ makes the payment. In the former case, the seller's profit is at most

$$\Pi^c(\alpha) = \mathbb{E}\left[(V_T - c)\mathbb{1}\left(\max_{t \geqslant 0} V_t \geqslant \alpha, V_T \geqslant c\right)\right] - e^{-\lambda T}\int_\alpha^1 [1 - F(v)].$$

In the latter it is at most $\Pi^c(\alpha + \beta)$.

Step II. Next, we claim that the bound $\Pi^c(\alpha)$ with $\alpha \geqslant c$ is achieved by the pricing strategy

[34] While this is only a numerical observation, we conjecture that this comparative static of monotonicity of α^c and β^c holds more generally.

identified in Proposition 5. Waiting until T is feasible, thus

$$W_t(v) \geqslant \mathbb{E}\left[(V_T-\alpha)^+ \,|V_t=v\right] \geqslant G_t(v) = e^{-\lambda(T-t)}(v-\alpha) + \left(1-e^{-\lambda(T-t)}\right)\mathbb{E}\left[(V_T-\alpha)^+\right].$$

Conclude that $W_t(v) = \mathbb{E}\left[(V_T-\alpha)^+ \,|V_t=v\right]$ and the smallest optimal stopping time is to stop at the first instance with $V_t \geqslant \alpha$. Of course, the buyer will return the good whenever $V_T < c$, this yields the profit of $\Pi^c(\alpha)$.

Step III. Write $\tilde{T}$ for the time of latest arrival, then for $\alpha \geqslant c$ the surplus is as it follows:

$$\mathbb{E}\left[\mathbb{E}\left[(V_T-c)^+\,\mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha\right)|\tilde{T}\right]\right] = \int_0^T \left[\hat{F}_{\tilde{T}}(\alpha)\int_\alpha^1 (v-c)dF(v) + \left[1-\hat{F}_{\tilde{T}}(\alpha)\right]\int_c^1 (v-c)dF(v)\right]\lambda e^{-\lambda(T-\tilde{T})}d\tilde{T} + \\ + e^{-\lambda T}\int_\alpha^1 (v-c)dF(v) = e^{-\lambda[1-F(\alpha)]T}\int_\alpha^1 (v-c)dF(v) + \left(1-e^{-\lambda[1-F(\alpha)]T}\right)\int_c^1 (v-c)dF(v).$$

Subtract the buyer's expected payoff $e^{-\lambda T}\int_\alpha^1[1-F(v)]dv$ and rearrange to obtain the following decomposition of $\Pi^c(\alpha)$:

$$\Pi^c(\alpha) = (\alpha-c)\left[1-F(\alpha)\right] + \left(1-e^{-\lambda T}\right)\int_\alpha^1 \left[1-F(v)\right]dv + \left(1-e^{-\lambda[1-F(\alpha)]T}\right)\int_c^\alpha (v-c)dF(v).$$

It only remains to choose the best threshold α^c, that is to solve $\max_{\alpha\in[c,1]} \Pi^c(\alpha)$. The first order condition for α^c is given by

$$\frac{1-F(\alpha^c)}{(\alpha^c-c)f(\alpha^c)} = e^{-\lambda T\cdot F(\alpha^c)}\left(1+\lambda T\int_c^{\alpha^c}\frac{(v-c)dF(v)}{\alpha^c-c}\right).$$

By the same argument as in Corollary 1, the threshold is well-defined by the first-order condition. Moreover, it is unique under the standard assumptions: a non-decreasing inverse hazard ratio.

□

Proof of Proposition 6. We only sketch the argument, because it is virtually identical to the proof of Theorem 2.

Positivity of cost does not alter the buyer's incentive constraints, thus Lemma 1 applies. Importantly, there still exists a number α such that there is no trade whenever $\max_{t\geqslant 0} V_t < \alpha$. In contrast, now the seller does not want to have trade for $V_T < c$.

If $\alpha \geqslant c$, then the following incentive compatible allocation maximizes surplus and keeps the buyer at the same level of rents, that is $\mathbb{E}\left[U_0(V_0)\right] = e^{-\lambda T}\mathbb{E}\left[(V_0-\alpha)^+\right]$:

$$Q(V^T) = \mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha, V_T \geqslant c\right),$$

which yields exactly $\Pi^c(\alpha) \leqslant \max_{\alpha\in[c,1]} \Pi^c(\alpha) = \Pi^c(\alpha^c)$.

On the other hand, any $\alpha < c$ yields necessarily leads to an inefficient trade, thus its profit is less than $\Pi^c(c) \leqslant \Pi^c(\alpha^c)$. □

9.8 Limited informational change: the case of a single arrival

It is possible that the analyst demands a model with limited informational change to put structure on a data set of dynamic prices and selling times. In the model we studied thus far, the value for consumption can change an arbitrary number of times with a given Poisson intensity. One way to enforce limited change of information is to have low values of λ and another way is to exogenously fix the number of times the value can change.

Here we consider the case where the buyer's value can change at most once in the time interval $[0, T]$, with Poisson intensity λ. The rest of the model is the same as before. This exercise is analogous to Deb [2014] with two differences: we have a finite time horizon and further he solves a relaxed problem where the Poisson arrival is observable (but not the new value). In Deb [2014], the seller commits to two price paths, an introductory one and another upon the arrival of the Poisson shock. Under some conditions on the primitives this delivers the optimum. We instead propose an upfront payment and a single continuously increasing price schedule and show that this *always* achieves the optimum. In the next result, we identify this pricing mechanism– it is "equivalent" to the baseline one, adjusting for the new model of informational change.

Proposition 7. *Suppose the buyer's value can change at most once. There exists $\alpha^s \in [0,1]$ such that the optimal pricing strategy, $\langle M^s, \mathbf{p}^s\rangle$, is the unique non-trivial solution (which always exists) to the following system:*

$$M^s = \lambda \int_0^T e^{-\lambda t}\left(\int_{p_t^s}^1 [1-F(v)]\, dv\right) dt, \quad \frac{dp_t^s}{dt} = \lambda \int_0^{p_t^s} F(v) dv \quad \text{and} \quad p_T^s = \alpha^s.$$

The buyer always makes the upfront payment and purchases only when he receives a new value. Specifically, the good is sold at time $t = 0$ whenever $V_0 \geqslant \alpha^s$, the good is sold at time $t > 0$ whenever $V_0 < \alpha^s$ and $V_t \geqslant p_t^s$. If $V_0 < \alpha^s$ and $V_t < p_t^s$ for all t, there is no trade.

As before, the seller's optimization problem is reduced to a single dimension. We fix the spot price at α and choose the prices p_t^s so that the buyer with value α whose type has not changed is made indifferent between buying or waiting:

$$\underbrace{\mathbb{E}[V_T | V_t = \alpha] - p_t^s}_{\text{buy at } t} = \underbrace{\int_t^T \lambda e^{-\lambda(r-t)} \mathbb{E}\left[(V_r - p_r^s)^+\right]}_{\text{do not buy before a new arrival}}.$$

The left-hand side is the buyer's gain from "stopping" immediately, and the right-hand side is the payoff from not stopping until a new arrival. Intuitively, such prices incentivize the buyer to make a purchase only in the event of an informational change, which is necessary for screening. In the proof we show that prices are pinned down by the differential equation given in Proposition 7.

The seller's profit $\Pi^s(\alpha)$ admits the following representation, which parallels ($\star$):

$$\Pi^s(\alpha) := \alpha[1-F(\alpha)] + \left(1 - e^{-\lambda T}\right)\int_\alpha^1 [1-F(v)]dv + \lambda \int_0^T e^{-\lambda t}\left[-\int_\alpha^1 v dF(v) + F(\alpha)\int_{p_t^s}^1 v dF(v)\right] dt.$$

They key conceptual difference from ($\star$) is that p_t^s enters the profit function. As can be seen from the statement of Proposition 7, this manifests in a time dependent threshold on valuations, $V_t \geqslant p_t^s$, unlike Theorem 1. However, since p_t^s change continuously with α, the profit function is continuous in α, and thus there exists an optimal terminal threshold: $\alpha^s = \arg\max_{\alpha\in[0,1]} \Pi^s(\alpha)$. Next, we show that this pricing mechanism implements the global optimum.

Proposition 8. *Suppose the buyer's value can change at most once. The seller's profit is at most $\Pi^s(\alpha^s)$ for any implementable mechanism.*

Unlike the baseline, the model is almost impenetrable without an appropriate conjecture about the optimal allocation rule. The reason being that the stochastic process of valuations is no longer Markov in time and the current value. Typically, when we invoke the revelation principle in dynamic mechanisms where the agent reports "new information" each period, we implicitly exploit a full support assumptions on the Markov process. In this model with limited number of Poisson jumps this assumption is not satisfied. If the agent has lied in the past, in some situations, he simply cannot go back to truthtelling, because if an arrival has been reported when it didn't happen, there is no possibility of a new arrival according to that history. Thus, the equivalence between "on-path" and "off-path" incentive constraints breaks down, and double deviations need to be accounted for.

To prove Proposition 8 we use the optimal pricing mechanism described in Proposition 7 to make an educated guess about the set of constraints that would bind at the optimum. Start by considering a relaxed problem that includes two sets of incentive constraints: (i) buyer with $V_0 \neq \alpha^s$ must be honest about the arrival time, i.e. cannot report the arrival which did not happen and cannot delay reporting one which happened, and (ii) the marginal type at the inception, $V_0 = \alpha^s$, underreports and then claims arrival at some time t equal to $\mathbb{E}[V_T|V_t = \alpha^s]$, which is the fair estimate given no actual arrivals. The first set of constraints generate a sequence of thresholds $(\alpha_t(V_0))_{t\in[0,T]}$ so that the buyer never lies about the arrival time and buys at $t > 0$ if and only if there is an arrival at this time and $V_t \geqslant \alpha_t(V_0)$. The second set of constraints dynamically connect these thresholds across time and make them independent of V_0 so as in the benchmark case these are only time dependent. Here $\alpha^s = \alpha_0$ and α_t for $t > 0$ cannot be too low, otherwise the marginal buyer with $V_0 = \alpha^s$ would want to buyer earlier.

The prices determined in Proposition 7 execute an implementation rule that corresponds one-for-one with the thresholds mentioned above that are pinned down by Proposition 8. Since the pricing mechanism is implementable, the remaining incentive constraint (not included in the relaxed problem and there are many of them) are satisfied as well, which completes the argument. This technique can plausibly be extended to a general model of exogenously fixed finite number of arrivals, and other stochastic processes that fail the one-shot deviation principle.

Proof of Proposition 7. Rewrite the buyer's decision problem as a stopping problem. The difference is that a gain process is no longer Markov in a current value and time, we also need to take into account whether or not there was an arrival. Formally: the gain from stopping at t with $V_t = v$ after the arrival is $G_t^A(v) := v - p_t$, whereas the gain from stopping at t with $V_t = v$ before the arrival is $G_t^B(v) := \mathbb{E}[V_T|V_t = v] - p_t$.

We solve the problem backwards. Let $W_t^A(v)$ be the value function after the arrival, that is

$$W_t^A(v) := \left(\sup_{s\geqslant t} G_t^A(v)\right)^+ = \left(v - \inf_{s\geqslant t} p_s\right)^+.$$

Similarly, let $W_t^B(v)$ be the value before the arrival. $W_t^B(v)$ can be characterized by the expression which resembles Equation (†):

$$W_t^B(v) := \max\left\{\sup_{s\in[t,T]} e^{-\lambda(s-t)}G_s(v) + \int_t^s \lambda e^{-\lambda(r-t)}\mathbb{E}\left[W_r^A(V_r)\right]dr, \int_t^T \lambda e^{-\lambda(r-t)}\mathbb{E}\left[W_r^A(V_r)\right]dr\right\}.$$

Next, we follow the same steps as in the proof of Theorem 1.

Step I. First, we derive the upper bound on the seller's profit. Relabel the prices as it follows:

$$p_t = e^{-\lambda(T-t)}\alpha_t + \left(1 - e^{-\lambda(T-t)}\right)\mathbb{E}[V_T] - \int_t^T \lambda e^{-\lambda(r-t)}\mathbb{E}\left[W_r^A(V_r)\right]dr.$$

Substitute these prices into the expressions for $G_t^B(v)$ and $W_t^B(v)$:

$$\begin{aligned} G_t^B(v) &= e^{-\lambda(T-t)}(v - \alpha_t) + \int_t^T \lambda e^{-\lambda(r-t)}\mathbb{E}\left[W_r^A(V_r)\right]dr, \\ W_t^B(v) &= e^{-\lambda(T-t)}\left(v - \inf_{s\geqslant t}\alpha_s\right)^+ + \int_t^T \lambda e^{-\lambda(r-t)}\mathbb{E}\left[W_r^A(V_r)\right]dr. \end{aligned}$$

It is convenient to introduce the following auxiliary notations: $\alpha := \inf_{t\geqslant 0}\alpha_s$ and β defined implicitly as $M = e^{-\lambda T}\beta + \int_0^T \lambda e^{-\lambda t}\mathbb{E}\left[W_t^A(V_t)\right]dt$.

If $\beta = 0$, then the buyer will always participate and receive the expected payoff of $e^{-\lambda T}\mathbb{E}\left[(V_0 - \alpha)^+\right]$. We shall bound the seller's profit for given α. Suppose first that $\alpha \geqslant 0$ and define p_t^s by the following integral equation:

$$p_t^s = e^{-\lambda(T-t)}\alpha + \int_t^T \lambda e^{-\lambda(s-t)}\left(\mathbb{E}[V_T] - \mathbb{E}\left[(V_s - p_s^s)^+\right]\right)ds, \quad p_T^s = \alpha.$$

Note that the equation has a unique solution for any $\alpha \geq 0$. To see it, rewrite the equation as

$$e^{-\lambda t}p_t^s = e^{-\lambda T}\alpha + \int_t^T \lambda e^{-\lambda s}\int_0^{p_s^s}[1 - F(v)]dvds.$$

Take a derivate on each side with respect to time and rearrange to obtain the following:

$$(p_t^s)' = \lambda\int_0^{p_t^s} F(v)dv \geqslant 0.$$

This differential equation with $p_T = \alpha \geqslant 0$ always have a unique solution, see Wallach [1948].

Importantly, $p_t^s \leqslant p_t$ for all t as $\alpha \leqslant \alpha_t$ for all t. Conclude that there is no trade whenever

$$V_0 < \alpha \text{ if } \tilde{T} = 0 \quad \text{and} \quad V_0 < \alpha, V_T < \hat{p}^s_{\tilde{T}} \text{ if } \tilde{T} \neq 0$$

where $\tilde{T}$ is the arrival time. It follows that the seller's profit is at most $\Pi^s(\alpha)$:

$$\Pi^s(\alpha) = \mathbb{E}\left[V_T \cdot \left(\mathbb{1}\left(V_0 \geqslant \alpha, \tilde{T} = 0\right) + \mathbb{1}\left(V_0 < \alpha, \tilde{T} \neq 0, V_T \geqslant \hat{p}_{\tilde{T}}\right)\right)\right] - e^{-\lambda T}\mathbb{E}\left[(V_0 - \alpha)^+\right].$$

Since p_t^s is continuously increasing in α for all t, $\Pi^s(1) \geqslant \Pi^s(\alpha)$ for $\alpha > 1$.

For $\alpha \leqslant 0$: the seller's profit cannot be higher than the maximal surplus net the buyer's expected payoff, that is

$$\mathbb{E}[V_T] - e^{-\lambda T}\left(\mathbb{E}[V_0] + \alpha\right) \leqslant \Pi^s(0).$$

The inequality follows from the fact that $p_t^s = 0$ is the solution for $\alpha = 0$.

Next, we consider $\beta > 0$. In this case, the buyer will pay M if and only if $V_0 \geqslant \alpha + \beta$ and his expected payoff is

$$\mathbb{E}\left[(V_0 - \alpha - \beta)^+\right].$$

Since the surplus is at most $\mathbb{E}[V_T \cdot \mathbb{1}(V_0 \geqslant \alpha + \beta)]$, the seller's profit is at most $\Pi^s(\alpha + \beta)$.

Step II. We have described the upper bound on the seller's profit, that is $\sup_{\alpha \in [0,1]} \Pi^s(\alpha)$. Next, we show that this upper bound can be attained. Indeed, for fixed $\alpha \in [0, 1]$, let p_t^s be defined as in Step I. Since $(p_t^s)' \geqslant 0$, after the arrival, it is optimal for the buyer to stop at t if $V_t \geqslant p_t^s$. Moreover, before the arrival, it is optimal to stop at t if $V_t \geqslant \alpha$. This yields the smallest stopping time: $V_0 \geqslant \alpha$ and $V_0 < \alpha$ and $V_T \geqslant p^s_{\tilde{T}}$ when $\tilde{T} \neq 0$. Finally, note that p_t^s changes continuously with α, thus $\Pi^s(\alpha)$ admits a maximum α^s.

Step III. It remains only to solve for Π^s in a closed form:

$$\Pi^s(\alpha) = e^{-\lambda T}\alpha[1 - F(\alpha)] + F(\alpha)\int_0^T \lambda e^{-\lambda t}\left[\int_{p_t}^1 v dF(v)\right] dt =$$
$$= \alpha[1 - F(\alpha)] + \left(1 - e^{-\lambda T}\right)\int_\alpha^1 [1 - F(v)]dv + \lambda\int_0^T e^{-\lambda t}\left[-\int_\alpha^1 v dF(v) + F(\alpha)\int_{p_t}^1 v dF(v)\right] dt.$$

□

Proof of Proposition 8. Let $\tilde{T}$ be the arrival time. Since only one arrival is possible, a space of histories can be written as a union of two cases: (i) $X_0 = x_0$ and no arrival, (ii) $X_0 = x_0, X_1 = x_1$ with $\tilde{T} = t$. To save on notations, we write $Q(x_0) = Q(x_0^{[0,T]})$ and $Q_t(x_0, x_1) = Q\left(x_0^{[0,\tilde{T})}, x_1^{[\tilde{T},T]}\right)$ We will prove the result by considering a relaxation of the seller's problem. The first step is to derive formally the relevant set of buyer's incentive constraints.

Define by $U_t(x_0)$ the buyer's value at t given that x_0 reported truthfully and there was no arrival before. In addition, let $U_t(\hat{x}_0, x_1)$ be the buyer's value at t given that $\tilde{T} = t$ and x_1 was reported truthfully, $\hat{x}_0$ was the initial report. Clearly, the buyer who misreported at the initial date faces exactly the same incentive problem as the buyer who lied, thus there is no loss to assume that $\hat{x}_0 = x_0$.

To begin, the buyer can misreport $\hat{x}_0$ and switch to truth-telling only at the time of the arrival. Similarly to Lemma 1, we obtain that

$$U_0(x_0) - U_0(\hat{x}_0) = \int_{\hat{x}_0}^{x_0} Q(v)dv, \quad Q(v) \text{ is non-decreasing.}$$

Thus, the buyer's expected payoff is completely determined by the allocation along the persistent history. This already gives the following expression for the seller's profit:

$$e^{-\lambda T} \int_0^1 (vf(v) - [1 - F(v)])\, Q(v)dv + \int_0^T \lambda e^{-\lambda t} \mathbb{E}\left[x_1 Q_t(x_0, x_1)\right] dt.$$

where the second term corresponds to the surplus conditional on the arrival. To pin down $Q_t(x_0, x_1)$ we need to look at other possible deviations.

Of course, the buyer can misreport $\hat{x}_0$ and then switch to truth-telling only at the time of arrival or at some fixed time t, in which case he will report his best estimate of V_t- $\mathbb{E}[V_T | V_t = x_0]$, thus

$$U_0(x_0) - U_0(\hat{x}_0) \geqslant e^{-\lambda t} \left[U_t\left(\hat{x}_0, \mathbb{E}[V_T | V_t = x_0]\right) - U_t(\hat{x}_0)\right].$$

Multiple both sides by $e^{\lambda t}$ and rewrite this incentive constraint as

$$e^{-\lambda(T-t)} \int_{\hat{x}_0}^{x_0} Q(v)dv \geqslant \Big[U_t\left(\hat{x}_0, \mathbb{E}[V_T | V_t = x_0]\right) - U_t\left(\hat{x}_0, 0\right)\Big] + \Big[U_t\left(\hat{x}_0, 0\right) - U_t\left(\hat{x}_0\right)\Big].$$

Note that the left side is expressed only as a function of Q. The next step is to use some of the remaining incentive constraints to rewrite each term on the right as a function of Q_t.

Now, suppose that there was an arrival of x_1 at t. The buyer can misreport in two ways: either claim $\hat{x}_1 \neq x_1$ or wait for a bit longer. The incentive constraint associated to the former is similar to the static setting:

$$U_t(x_0, x_1) - U_t(x_0, 0) = \int_0^{x_1} Q_t(x_0, v)dv, \quad x_1 \mapsto Q_t(x_0, x_1) \text{ is non-decreasing.}$$

In addition, the buyer with $x_1 = 0$ can always wait until the final date which yields the following constraint:

$$\begin{aligned} U_t(x_0, 0) - U_t(x_0) &\geqslant \int_t^T \lambda e^{-\lambda(s-t)} \mathbb{E}\left[U_s(x_0, 0) - U_s(x_0, x_1)\right] ds = \\ &= - \int_t^T \lambda e^{-\lambda(s-t)} \left(\int_0^1 [1 - F(v)] Q_s(x_0, v)dv\right) ds \end{aligned}$$

where the last expression is obtained by integration by parts.

Combine all these constraints together and evaluate them at $x_0 = \alpha$ and $\hat{x}_0 < \alpha$ to get the following necessary condition:

$$\int_t^T \lambda e^{-\lambda(s-t)} \left(\int_0^1 [1 - F(v)] Q_s(\hat{x}_0, v) dv \right) ds \geqslant \int_0^{\mathbb{E}[V_T | V_t = \alpha]} Q_t(\hat{x}_0, v) dv. \tag{‡}$$

So, our goal is to find the allocation which maximizes the surplus and satisfies Equation (‡), formally:

$$\max_{(Q, Q_t)} \quad e^{-\lambda T} \int_0^1 (v f(v) - [1 - F(v)]) \, Q(v) dv + \int_0^T \lambda e^{-\lambda t} \mathbb{E}\left[x_1 Q_t(x_0, x_1)\right] dt \text{ subject to (‡),}$$
$$Q \text{ is non-decreasing, } x_1 \mapsto Q_t(x_0, x_1) \text{ is non-decreasing for almost all } x_0 \text{ and } t.$$

As before, in a deterministic mechanism there exists a number α such that $Q(x_0) = 1$ if and only if $x_0 \geqslant \alpha$. By the same reasoning, there exists $\alpha_t(x_0)$ such that $Q_t(x_0, x_1) = 1$ if and only if $x_1 \geqslant \alpha_t(x_0)$. Since Equation (‡) is required only for $\hat{x}_0 < \alpha$, it is optimal to have $Q_t(x_0, x_1) = 1$ whenever $x_0 \geqslant \alpha$.

On the other hand, for $\hat{x}_0 < \alpha$, the threshold of $\alpha_t(\hat{x}_0) = 0$ is not feasible as it always violates Equation (‡). To see it formally, rewrite the equation in the following way:

$$\alpha_t(\hat{x}_0) \geqslant e^{-\lambda(T-t)} \alpha + \int_t^T e^{-\lambda(s-t)} \left(\int_0^{\alpha_s(\hat{x}_0)} [1 - F(v)] dv \right) ds.$$

Note that it is optimal to choose the smallest thresholds, that is exactly p_t^s as defined in Proposition 7 which yields $Q_t(x_0, x_1) = 1$ if and only if $x_1 \geqslant p_t^s$ whenever $x_0 < \alpha$.

To sum up, we have described the allocation which solves the relaxed problem. It is incentive compatible, in particular as shown in Proposition 7 this allocation can be implemented by dynamic pricing.

□

References

M. Akan, B. Ata, and J. D. Dana Jr. Revenue management by sequential screening. *Journal of Economic Theory*, 159(B):728–774, 2015.

M. Armstrong. Nonlinear pricing. *Annual Review of Economics*, 8:583–614, 2016.

M. Battaglini and R. Lamba. Optimal dynamic contracting: the first-order approach and beyond. *Theoretical Economics*, forthcoming, 2019.

D. Bergemann and P. Strack. Dynamic revenue maximization: A continuous time approach. *Journal of Economic Theory*, 159:819–853, 2015.

D. Bergemann and J. Välimäki. Dynamic mechanism design: an introduction. *Journal of Economic Literature*, forthcoming, 2019.

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

R. Boleslavsky and M. Said. Progressive screening: Long-term contracting with a privately known stochastic process. *Review of Economic Studies*, 80(1):1–34, 2013.

T. Börgers. *An Introduction to the Theory of Mechanism Design*. Oxford university press, 2015.

J. Bulow and J. Roberts. The simple economics of optimal auctions. *Journal of Political Economy,*, 97(5):1060–1090, 1989.

M. K. Chen and M. Sheldon. Dynamic pricing in a labor market: Surge pricing and flexible work on the Uber platform. UCLA and University of Chicago, 2016.

P. Courty and H. Li. Sequential screening. *Review of Economic Studies*, 67(4):697–717, 2000.

J. D. Dana Jr. Advance-purchase discounts and price discrimination in competitive markets. *Journal of Political Economy,*, 106(2):395–422, 1998.

R. Deb. Intertemporal price discrimination with stochastic values. University of Toronto, 2014.

J. Ely, D. F. G. Garrett, and T. Hinnosaar. Overbooking. *Journal of the European Economic Association*, 15(6):1258–1301, 2017.

P. Esö and B. Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74(3):705–731, 2007.

P. Esö and B. Szentes. Dynamic contracting: an irrelevance theorem. *Theoretical Economics*, 12(1): 109–139, 2017.

I. L. Gale and T. J. Holmes. Advance-purchase discounts and monopoly allocation of capacity. *American Economic Review*, 83(1):135–146, 1993.

D. Garrett. Intertemporal price discrimination: dynamic arrivals and changing values. *American Economic Review*, 106(11):3275—3299, 2016.

T. Hinnosaar and K. Kawai. Robust pricing with refunds. Collegio Carlo Alberto and UNSW, Sydney, 2019.

D. Krähmer and R. Strausz. Optimal sales contracts with withdrawal rights. *Review of Economic Studies*, 82(2):762–790, 2015.

I. Krasikov and R. Lamba. A theory of dynamic contracting with financial constraints. Penn State University, 2019.

T. Kruse and P. Strack. Optimal stopping with private information. *Journal of Economic Theory*, 159(B):702–727, 2015.

J.-J. Laffont, H. Ossard, and Q. Vuong. Econometrics of first-price auctions. *Econometrica*, 63(4): 953–980, 1995.

H. Li and X. Shi. Discriminatory information disclosure. *American Economic Review*, 107(11): 3363–3385, 2017.

R. P. McAfee and V. te Velde. Dynamic pricing in the airline industry. Caltech, 2006.

R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

A. Pavan, I. Segal, and J. Toikka. Dynamic mechanism design: A Myersonian approach. *Econometrica*, 82(2):601–653, 2014.

S. Stantcheva. Dynamic taxation. *Annual Review of Economics*, forthcoming, 2020.

G. Stigler. *Theory of Price*. Macmillan, New York, 1987.

N. Stokey. Intertemporal price discrimination. *The Quarterly Journal of Economics*, 93(3):355–371, 1979.

A. Sweeting. Dynamic pricing behavior in perishable goods markets: Evidence from secondary markets for Major League Baseball tickets. *Journal of Political Economy*, 120(6):1133–1172, 2012.

H. Varian. Price discrimination. In R. Schmalensee and R. Willig, editors, *Handbook of Industrial Organization, Vol. 1*, chapter 10, pages 597–654. Elsevier, 1989.

K. R. Williams. Dynamic airline pricing and seat availability. Yale University, 2018.