On dynamic pricing*

Ilia Krasikov Rohit Lamba

October 2021

Abstract

This paper builds a theory of dynamic pricing for the sale of timed goods. The main friction is private and evolving valuation of the buyer prior to the date of consumption, which follows a Poisson process. 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. The tools developed here are shown to be useful in thinking about global incentives in dynamic mechanisms, and mapping dynamic pricing to the classic taxonomy of consumer-producer surplus and deadweight loss.

1 Introduction

Two-part tariff is a commonly used pricing strategy. Typically the first part is a flat payment that aims to extract some part of the total surplus and the second part segments the market of buyers. A uniform fee to enter an amusement park followed by individual prices for high-end rides is an oft invoked example. This paper studies trade of timed goods where the second component of the two-part tariff endogenously sorts buyers across timing of purchase. Examples to keep in mind are concerts, hotels and air travel.

Suppose a buyer wants to consume a good at a specified future date. The value of the good at the date of consumption is drawn from some prior and then evolves over time according to a Poisson process. The seller is not privy to this value. She posts a menu with an upfront payment and time dependent prices. All buyer types pay the first part. Then, as their values evolve, they sort themselves into different buckets of when (if at all) to trade along what turns out to be a continuously increasing price path.

The optimal pricing scheme within this class of instruments is completely characterized. The buyer faces an optimal stopping problem for any given pricing mechanism. The seller in turn solves a complex dynamic optimization problem of choosing the profit maximizing mechanism. This is

*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 Mikhail Golosov, 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.

reduced to the choice of a single threshold, viz. the final spot price at the date of consumption, from which the optimal price path and upfront payment are backward imputed. Next, a description of the pricing mechanism in the language of consumer-producer surplus and deadweight loss is presented. The flat fee transfers some part of consumer surplus, and the changing price schedule transfers some deadweight loss to producer surplus. Finally, using tools of dynamic mechanism design, it is shown that the allocation implemented by the pricing mechanism achieves the global (deterministic) optimum.

This paper builds on the two-period sequential screening problems pioneered by Courty and Li [2000], and developed further by Esö and Szentes [2007], Krähmer and Strausz [2015] and Akan, Ata, and Dana Jr. [2015], amongst others. They study an option pricing mechanism wherein the buyer picks a menu in the first period from which a price in the second period is then chosen. In contrast, we study a continuous-time finite horizon problem with one menu where an upfront fee and time dependent price path together determine the structure of price discrimination.

In a closely related paper, Deb [2014] studies the sale of a single good over an infinite horizon where the value of the buyer can change at most once through a Poisson shock. It provides a partial characterization of the optimum through two prices—an introductory and a final price. The buyer's choice as an optimal stopping problem with changing valuations finds resonance in Kruse and Strack [2015]. Also, related is the study of dynamic price discrimination in a repeated sales setting with Brownian shocks by Bergemann and Strack [2015]. Bergemann and Välimäki [2019] offers a review of dynamic mechanism design—the principal-agent problem with privately evolving agent's valuations falls into that realm.

A salient departure from much of this literature is that we study a model in which the standard first-order approach fails generically. For any prior distribution and interior Poisson arrival rate, the envelope theorem is not sufficient to guarantee incentive compatibility. Inspired from the study of static multidimensional screening (see Armstrong [2016]), we turn to indirect mechanisms to understand which class of constraints may be bind in the Myersonian characterizations of direct mechanisms.

We pick the Poisson (or renewal Markov) model for the evolution of valuations for two reasons. It is tractable enough to be solved in closed form and still interesting enough to deliver insights about the structure of binding constraints, and correspondingly the simple economics of dynamic price discrimination. In fact, it has been suggested as a canonical setup to understand the impact of global incentive constraints in dynamic mechanisms (see Pavan, Segal, and Toikka [2014] and Bergemann and Välimäki [2019]).

Moreover, a leading application of dynamic mechanism design is the study of optimal taxation in which types are assumed to be earnings (see Stantcheva [2020]). Recent work in labor economics has established that the earnings follows a *job-ladder* structure (see Arellano et al. [2017] and Guvenen et al. [2021]). This violates the standardly used $AR(1)$ family of processes for earnings (eg. Farhi and Werning [2013]); because earnings are sticky and feature some loss of history upon the arrival of a new event such as unemployment or promotion, properties approximated by the renewal Markov process. Analogously, for the sequential pricing model, the process captures what we call the *value-ladder*: Buyer's valuation do not change too frequently. And, conditional on change, lower valuation buyers are more likely to increase their valuation and high valuation buyers to see a drop in valuation.

Before delving into the general model, we explore a two period version. The solution to this problem helps to fix ideas about the general model. It also provides pedagogical value in understanding the economics of sequential pricing in a simple way, especially when global incentive constraints bind at the optimum.

2 An example

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 an "estimate" of the eventual consumption value with precision $1-\lambda$. The cost of production for the seller is zero. For simplicity of exposition, 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.

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. 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.[1] 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

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

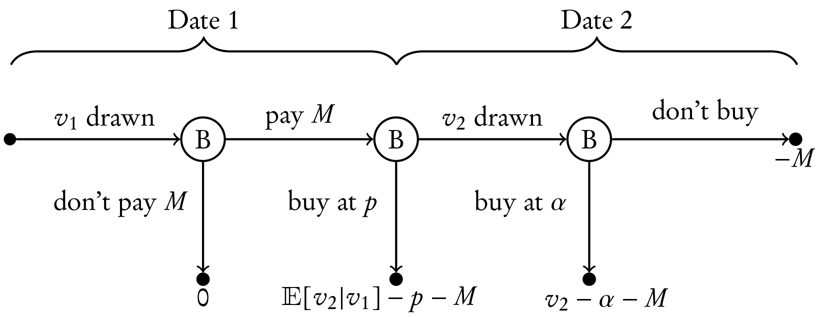

Figure 1: Timing of the contract $\langle M, p, \alpha \rangle$.

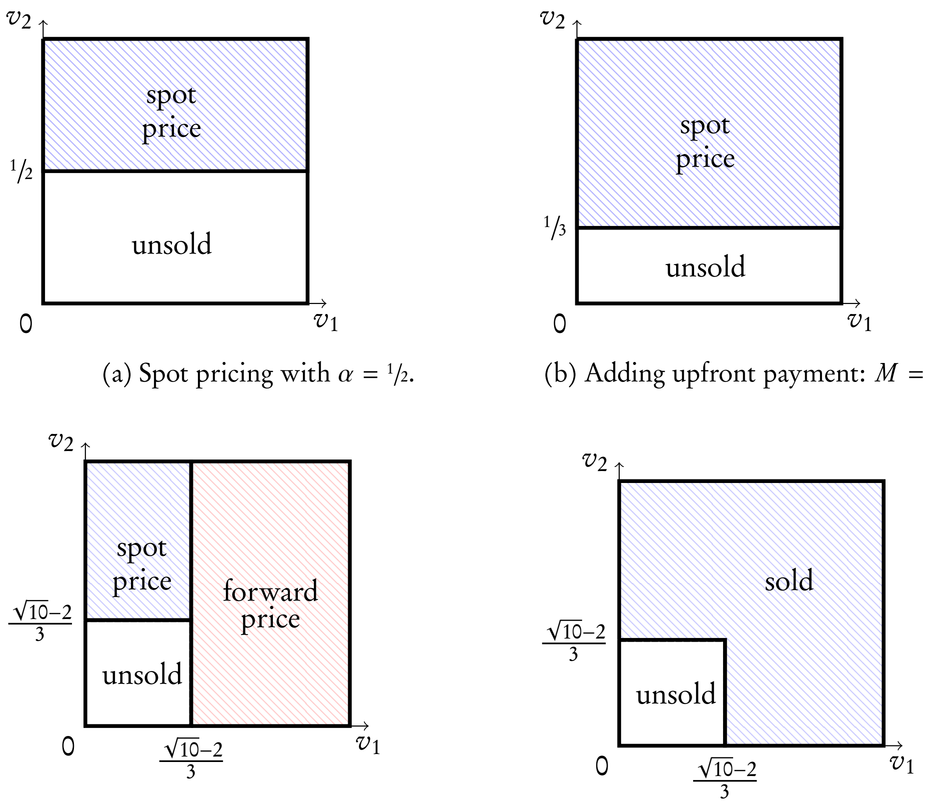

(a) Spot pricing with $\alpha = 1/2$.

(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 $\lambda = 1/2$.

expected profit of 1/3. So, all types agree to pay M, and as Figure 2b shows, trade takes place when $v_2 \geq 1/3$.

Going further, we ask: 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:

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

So, M is chosen to extract expected surplus of the lowest first period type, and not tinker with selection. 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. The no trade region forms a smaller square of side α (see Figure 2c).

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.

Two aspects of our construction need probing: Can the seller do better by making some other type $v_1 = \alpha'$ indifferent between waiting and trading in the first period, where $\alpha' > \alpha$ or $\alpha' < \alpha$. And, further still, can the seller do better by employing some other pricing mechanism altogether? Next, we show that the pricing mechanism characterized here is the optimal dynamic contract—the seller cannot do better, which answers both questions.

Optimal dynamic mechanism. To establish the global optimality of our pricing mechanism, we state the general two-period dynamic mechanism design problem, invoke the revelation principle, and then solve the problem to show that the optimal allocation is given by $q(v_1, v_2) = \mathbb{1}\left(\max\{v_1, v_2\} \geqslant \frac{\sqrt{10}-2}{3}\right)$.

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. 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 $q(v_1, v_2)$, the latter is termed as the envelope formula

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 ${}^1\!/_2 \times q(v_1, v_1)$. The latter is termed as the dynamic envelope formula; it gives the following expression for buyer's payoff:

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

In the first period convexity is necessary but not sufficient for incentive compatibility. To produce a condition on allocation, subtract $U_1(\hat{v}_1)$ from each side of the incentive constraint and substitute $U_1(.)$ and $U_2(\hat{v}_1, .)$ to get:

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

This condition, known as *integral monotonicity* (Pavan, Segal, and Toikka [2014]), characterizes incentive compatibility when information arrives gradually.

The next step is to write the seller's profit as a function of q. Because of linearity, profit is the difference between surplus and buyer's expected payoff. Surplus is simply $\mathbb{E}[v_2 q(v_1, v_2)]$, and buyer's expected payoff is $\mathbb{E}[U_1(v_1)]$. We of course 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 using integration by parts the seller's problem can finally be stated as:

$$\max_{q \in \{0,1\}} \quad \mathbb{E}[v_2 q(v_1, v_2)] - {}^1\!/_2 \times \int_0^1 (1 - v) q(v, v) dv.$$

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

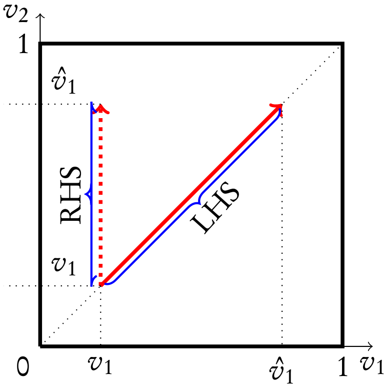

(a) Integral monotonicity

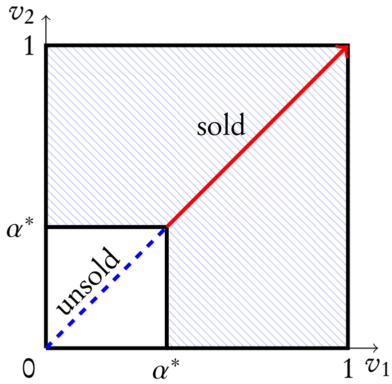

(b) Construction of the optimum

Figure 3: Optimal dynamic mechanism

should be monotonic.[2]

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 $v \geqslant \alpha$. This splits the diagonal in Figure 3b into trade (red-solid) and no-trade (blue-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. Moreover, 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

$$\mathbb{E}\left[v_2 \cdot \mathbb{1}\left(\max\{v_1, v_2\} \geqslant \alpha\right)\right] - {}^1\!/\!_2 \times \int_\alpha^1 (1-v)dv = {}^1\!/\!_2 \times \alpha[1-\alpha] + {}^1\!/\!_2 \times \int_\alpha^1 \frac{3v^2}{2} dv.$$

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, 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 randomization.[3] Now, we take these insights from the two-period model to a continuous time model with Poisson arrivals.

[2] The first-order approach would maximize the objective ignoring (i) and (ii). The solution 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).

[3] The proof of this claim is available from the authors upon request.

3 Setup

3.1 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, which is positive on $[0, 1]$

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.

At the outset, we consider the following pricing instrument: At the start the seller asks the buyer to make an upfront payment M and offers a menu of time-dependent prices $\mathbf{p} = (p_t)_{t=0}^T$.[4] The buyer can either opt out or pay the upfront payment; 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 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.[5] If the buyer does not buy the good till T, there is no trade and the upfront payment is lost as a sunk cost.

As is standard, 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

[4]A continuous vector of prices till time t is denoted 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 is denoted by $p^{[t,T]} = (p_s)_{s=t}^T$.

[5]Since there is no discounting, 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.

as a Markov decision problem:

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

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

Internalizing the optimal response of the buyer, the seller's problem consists of the choice of upfront payment M and price path $(p_t)_{t=0}^T$ to maximize her expected profit. Note that, 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.

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 $\alpha_t(p^{[t,T]}) \in [0,1]$ such that if $V_t \geqslant \alpha_t$, buy, else wait. The seller can then optimize over 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:*

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

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 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:

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

Equation ($\star$) parallels a similar expression we derived for the two-period model. It shows that the seller's optimization problem has been reduced to choice of a single variable: a threshold given by $\alpha^* := \arg\max_{\alpha\in[0,1]} \Pi(\alpha)$.[6] We now provide a classical price-theoretic explanation of how each of three individual components in the equation stack up to constitute the seller's profit.

5 The simple economics of dynamic pricing[7]

Suppose the seller ignores dynamics and offers the good for a spot price α. The buyer will buy the good 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 trade will happen with the following probability: $\mathbb{P}(V_T \geqslant \alpha) = 1 - F(\alpha)$.

We will regard this as the buyer's inverse demand function: $D(p) := 1 - F(p)$, where $D(p)$ is quantity demanded at price p.[8] 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 $\alpha[1 - F(\alpha)]$, the first term of Equation ($\star$).[9]

Note that even the buyer with the lowest 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 an upfront fee:

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

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

The inverse demand function conditional on $V_0 = 0$ is $D_0(p) := \mathbb{P}\left(V_T \geqslant p | V_0 = 0\right)$, 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:

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

[7] The term "simple economics" is a homage to Bulow and Roberts [1989] which provides a price theoretic interpretation for Myerson [1981]'s optimal auction problem.

[8] 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.

[9] 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. Also, linearity of D is assumed for simplicity of exposition, this would of course be exact if F is uniform.

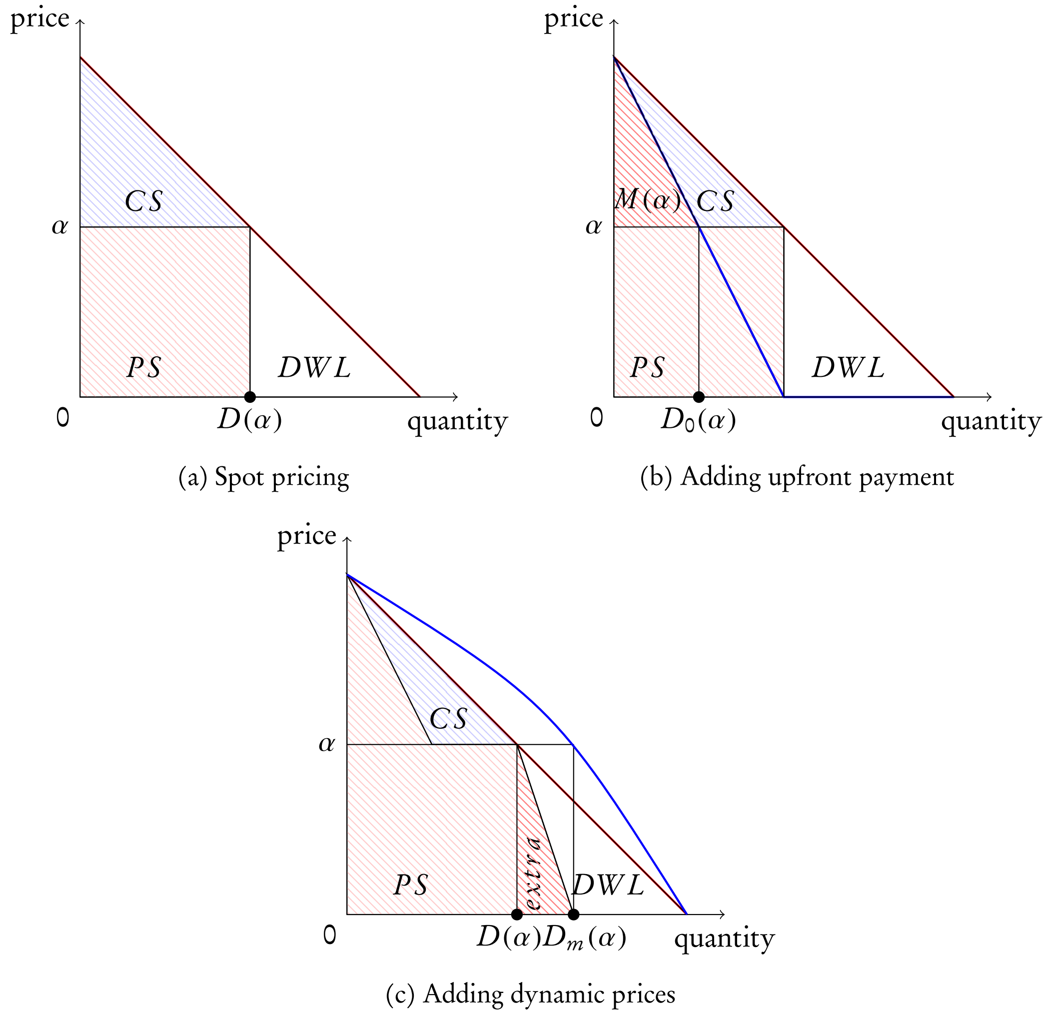

Figure 4: Decomposition of total surplus from trade into CS, PS and DWL

$M(\alpha)$ and the rest.[10] By construction, the seller extracts $M(\alpha)$ 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.

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 $\max\limits_{t\geqslant 0} V_t \geqslant \alpha$.

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 $D_m(p) = \mathbb{P}\left(\max\limits_{t\geq 0} V_t \geq p\right)$.[11,12] The change in

[10] 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.

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

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

DWL is as follows:

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

Here $\zeta := \mathbb{P}\big(\max_{t\geqslant 0} V_t \geqslant \alpha | V_T < \alpha\big)$ 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 [$\zeta \cdot$ static DWL], depicted in Figure 4c as "extra", is transferred from DWL to PS due to dynamic pricing.[13] The magnitude of this added profit for the seller is the third term in Equation ($\star$).

6 Comparative statics

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 multiplicatively. Thus, what matters for comparative statics is the normalized time λT.

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 non-increasing in* λT *with* $\lim_{\lambda T \to \infty} \alpha^* = 0$.

Part (b) states that as $T \to 0$, i.e. the time horizon shrinks, or as $\lambda \to 0$, i.e. values become perfectly persistent, the optimal contract coverages to its static benchmark. 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 seller can extract all surplus as profit using the upfront fee.

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

What about prices? The model predicts 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, prices for trade increase steadily from some initial value p_0^* to the final spot price p_T^*.

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* (see Dana Jr. [1998]). Typically these feature two types of buyers with idiosyncratic uncertainty in demand. Low valued buyers with higher uncertainty buy early and high valued buyers with lower uncertainty buy later. Here, in contrast, we have a continuous price path with a steadily declining discount. 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.[14]

7 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 $V_t \geqslant \alpha^* > \max_{s\in[0,t)} V_s$, 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.[15] 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.[16]

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

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

[14] 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. A die hard fan may receive some news later that he cannot be in town on the day of the concert, but it is a trade-off worth making given the advance purchase discount and current value estimate.

[15] 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.

[16] 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.

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,

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

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 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:*

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

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

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

Lemma 1 is a dynamic analog of the Myersonian characterization of incentive compatibility in static mechanism design (see Börgers [2015]). It is a counterpart to similar results established in discrete time by Pavan, Segal, and Toikka [2014] and Battaglini and Lamba [2019], and in continuous time with Brownian shocks by Bergemann and Strack [2015]. It pins down the space of allocations that can be implemented by some pricing strategy. And it expresses the buyer's (expected) payoffs as a function of the allocation rule, up to a constant. Equation (Env) represents the dynamic analog of the 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 types (Rochet [1987]).

The next natural restriction 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

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

So the buyer cannot be forced to continue the contract 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 following result states the global optimality of our pricing 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. The optimal choice of α then pins down the first-order optimum. This contract is obviously not incentive compatible.[17]

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 edges of length α. In the remaining histories $V^{[0,T]}$, where there exists t such that $V_t > \alpha$, the seller always wants trade, since the shadow price of incentives here is zero, and the surplus from trade can be extracted by binding the individual rationality constraint at time 0. Thus, for the optimal value of the threshold α^*, the allocation rule splits all histories into two classes, no-trade: $\max_{0\leqslant s\leqslant T} V_t < \alpha^*$, and trade: $\max_{0\leqslant s\leqslant T} V_t \geqslant \alpha^*$.

As a final thought, there are three levels of private information with the buyer: What is the initial type. whether there is a Poisson arrival, and in case of an arrival, what is the new type. If the second piece, 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; captured by Equation (Env) at $t = 0$. Since the draw of the new type upon an arrival is independent, and the news of 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 the intuition for why dynamics can be posed as being "irrelevant" under the assumption of sufficiency of the local approach (see Esö and Szentes [2017]), since only initial private information matters for dynamic information rents. We go beyond these class of models to demonstrate the relevance of dynamic interaction in mechanism design.

8 Final remarks

The model can be extended to incorporate a linear cost of production for the seller, and a common discount factor for the two players. The model studied by Deb [2014] in which the number of Poisson arrivals is restricted to be one can also be completely characterized using the pricing approach we explored here. Going beyond to more general stochastic processes may require a combination of menu choice and time dependent pricing. Exploring improvement in profits offered by randomization in dynamic pricing is also an open question.[18]

The largely normative exercise here is silent of the feasibility of an upfront payment in the

[17]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 $s \neq t$. The seller then optimizes over α.

[18]Details on the model with a single Poisson arrival, and improvements offered by randomization are available from the authors upon request.

pricing schedule. If it is not practical, the model demands the imposition of a financial or regulatory constraint. On the other hand, it maybe approximated by other instruments such as refundable prices or small but recurring membership fees.

9 Appendix

9.1 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 accounting for two possible cases: first when the buyer stops at the the best possible time along the constant path, and second when he does not stop at all along the constant path:

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

Note that $(\cdot)^+$ is omitted from G_t in the above equation. The reason is that the continuation value, W_t, 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 pays the upfront fee. 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 $e^{-\lambda(s-t)}\left[G_s(v) - G_s(0)\right] = e^{-\lambda(T-t)}v$ 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 $W_s(0) > \varepsilon + \lim_{\epsilon\downarrow 0}\sup_{r\in[s,s+\epsilon]} G_r(0)$ 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 ($\star$) 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 $\mathbb{E}[W_0(V_0) - W_0(0)] = e^{-\lambda T}\mathbb{E}[(V_0-\alpha)^+]$. What about total surplus? Since, $\alpha = \inf_{t\geqslant 0}\alpha_t$, the total surplus is bounded above by $\mathbb{E}[V_T \cdot \mathbb{1}(\max_{t\geqslant 0} V_t \geqslant \alpha)]$. Combining these two we obtain the following upper bound on seller's expected profit:

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

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

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

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 (Peskir and 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 (Peskir and 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 $\displaystyle\int_\alpha^1 v dF(v)$, and conditional on some $\tilde{T} \neq 0$ is:

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

By the law of iterated expectations:

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

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

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

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

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

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

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

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 $v \mapsto v f(v)$ 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 $\int_0^\alpha v dF(v)/\alpha$ with respect to α:

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

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, say $\hat{\alpha}$, 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 $\alpha_1^* > \alpha_2^*$ are the optimal thresholds for (λ_1, T_1) and (λ_2, T_2) with $\lambda_1 T_1 < \lambda_2 T_2$. Observe that the seller's profit can be rewritten in an integral form as it follows:

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

Then,

$$\int_{\alpha_i^*}^{\alpha_j^*}\left[e^{\lambda_i T_i\cdot F(v)}\left(v+\lambda_i T_i\int_0^v wdF(w)\right)f(v) - [1-F(v)]\right]dv \geqslant 0 \quad i,j = 1,2.$$

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

$$\int_{\alpha_1^*}^{\alpha_2^*} e^{\lambda_1 T_1\cdot F(v)}\left(v+\lambda_1 T_1\int_0^v wdF(w)\right)f(v)dv \geqslant \int_{\alpha_1^*}^{\alpha_2^*} e^{\lambda_2 T_2\cdot F(v)}\left(v+\lambda_1 T_1\int_0^v wdF(w)\right)f(v)dv,$$

which is a clear contradiction.

Since α^* is non-increasing 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.

$\square$

9.2 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. $\square$

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. □

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. Arellano, R. Blundell, and S. Bonhomme. Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica*, 85(3):693–734, 2017.

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.

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.

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.

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.

E. Farhi and I. Werning. Insurance and taxation over the life cycle. *Review of Economic Studies*, 80 (2):596–635, 2013.

F. Guvenen, F. Karahan, S. Ozkan, and J. Song. What do data on millions of u.s. workers reveal about lifecycle earnings dynamics? *Econometrica*, forthcoming, 2021.

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

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

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.

G. Peskir and A. N. Shiryaev. *Optimal Stopping and Free-Boundary Problems*. Springer Applied Maths, 2006.

J.-C. Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics*, 16(2):191–200, 1987.

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