On dynamic pricing*

Ilia Krasikov Rohit Lamba

November 2019

(Click for the latest version)

Abstract

There are many economic problems where the observable data consists of prices and selling times– airline tickets and hotel bookings are leading examples. This paper provides a dynamic pricing model to help better understand such scenarios: A seller wants to sell to a buyer a good with a fixed date of consumption. The buyer's value for it can change over time according to a Poisson process prior to the date. The seller posts a two-part tariff where the first part extracts the buyer's surplus modulo self-selection rents, and the second part sequentially segments the market in the spirit of second degree price discrimination through a continuously increasing price path. The buyer always pays the first part of the tariff and then solves an optimal stopping problem– which price to accept for trade. The solution of this pricing problem, solved in closed form, is shown to implement the optimal deterministic contract. In the process, a novel dynamic mechanism design problem is operationalized where the standard relaxed approach generically fails. Alternate implementations through refund and subscription contracts are also presented. The gains from randomization are explained through the channel of information acquisition, and the optimal contract for limited informational change modeled through a fixed number of Poisson arrivals is also solved.

1 Introduction

Dynamic price discrimination. 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 state of the then art on price discrimination, Varian [1989] wrote:

> "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*

*Krasikov: krasikovis.main@gmail.com, Tel Aviv University and Higher School of Economics Moscow; Lamba: rlamba@psu.edu; Pennsylvania State University. We are indebted to Vijay Krishna for his advice and guidance. We are also also grateful to Karl Schurter, Rakesh Vohra, and seminar participants at University of Pennsylvania, Chinese University of Hong Kong, Canadian Economic Theory Conference at McGill University, and NSF/NBER/CEME Conference at University of Chicago for their comments.

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 an air travel and booking a room on AirBnb. In both cases the timing of "consumption" is fixed, but the valuation for eventual consumption may change over time. So the seller can design a menu of sequential prices with an aim to monetize the difference between buyers who have high initial valuations and buy early, and those that initially have low valuations but might buy the good later if they learn their valuation has increased. Since the seller is not privy to the buyer's valuation, the changing price path executes the sorting endogenously through self-selection.

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 will extract the total surplus modulo the rents that need to be paid to satisfy sell-selection, and second part unpacks 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 has potential applications in other related models such as optimal taxation.

As is the norm in price discrimination, our model can be seen as a single seller-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-many buyers interaction where the distribution determines the size of market demand. To fix ideas we start by considering a simple and hopefully intuitive example.

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 λ, that is $v_2 = v_1$, and it will be re-drawn again from F with probability $1 - \lambda$, that is $v_2 \sim F$. The cost of production for the seller is zero. For simplicity of exposition assume F is uniform on $[0, 1]$ and $\lambda = 1/2$. What pricing mechanism(s) should the seller employ to maximize her profit?[1]

[1] To the best of our knowledge the economics of this simple example has not been studied in the literature.

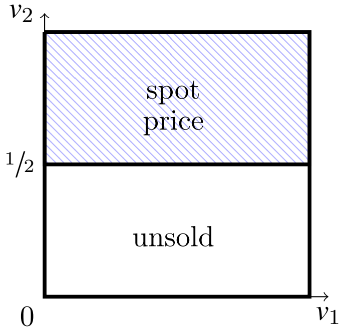

(a) Spot pricing with $\alpha = 1/2$.

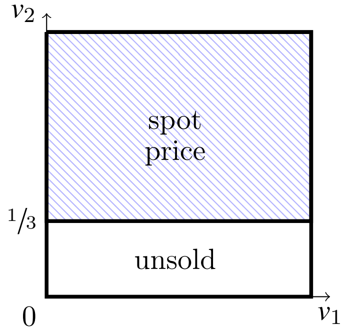

(b) Adding upfront payment: $M = 1/9, \alpha = 1/3$.

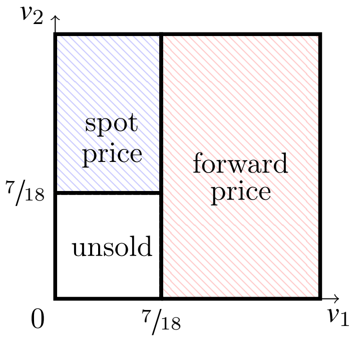

(c) Optimum:
$M = 121/1296, p = 455/1296, \alpha = 7/18$.

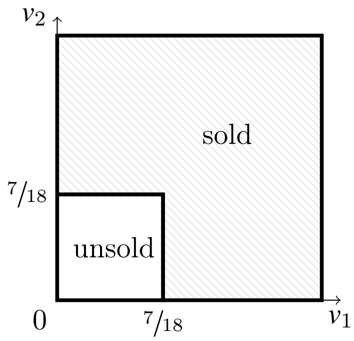

(d) Structure of trade at the optimum

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

Figure 1 represents the optimal trading regions for three 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 1 should be shaded.

In the first case, in Figure 1a, the seller ignores the dynamics of the problem and offers a spot price, say α. From an ex ante perspective v_2 is distributed uniformly on $[0,1]$, so the best the seller can do is to post a price of $\alpha = 1/2$. The spot price earns her a profit of $1/4$, since the good is sold with probability $1/2$ with a revenue of $1/2$.

The second contract, given by $\langle M, \alpha \rangle$, allows the seller to extract the buyer's surplus with an upfront payment: 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 optimal for the seller to put $\alpha = 1/3$, and charge the surplus of the lowest value customer at date 1 as the upfront payment: $M = \mathbb{E}\left[(v_2 - \alpha)^+ \mid v_1 = 0\right] = 1/9$.[2]. This gives an expected profit of $1/3$ to the seller. As shown in Figure 1b, trade takes place whenever $v_2 \geq 1/3$.

The contract $\langle M, \alpha \rangle$ is essentially a *two-part tariff*. It allows the seller to extract a uniform

Pavan, Segal, and Toikka [2014] and Bergemann and Välimäki [2019] mention it as an interesting case where the standard approach to dynamic mechanism design fails.

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

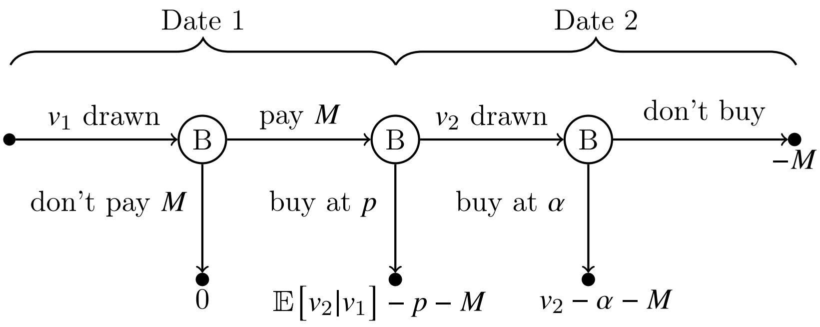

Figure 2: Timing of the contract $\langle M, p, \alpha \rangle$.

price M from all buyer types, and expands the frequency of trade by lowering the price from $1/2$ to $1/3$. But, since all types agree to pay M, who actually trades is decided at date 2 through the spot price α. Can it be profitable for the seller to screen buyers along both v_1 and v_2?

The contract space is now given by $\langle M, p, \alpha \rangle$, and its timing is depicted in Figure 2. At the start the buyer gets to choose between paying M upfront or ending the contract. If he pays M, he is granted access to two prices. At date 1, he can buy the good for p, or he can forgo p at date 1, and instead buy the good at date 2 for 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 or not.

M is again used purely for surplus extraction. The seller will optimally choose M to be the expected surplus of the buyer whose valuation at date 1 is zero. Intuitively, p must be lower than α, so that buyers with high enough realizations of v_1 are incentivized to buy at p and others wait till date 2 to trade if $v_2 \geq \alpha$. Moreover, in order to maximize profits, the seller wants to "bind the self-selection constraint": For a fixed α, she simultaneously chooses p and a threshold for v_1 so that all buyers whose valuations at date 1 one are above that threshold are exactly indifferent between paying p, and waiting for the realization of v_2 and deciding whether to trade at α. Finally, the seller would optimize by choosing the best value of α.

In what follows we develop this argument in more details. First of all, note that the buyer's expected payoff from waiting for the second period price conditional on v_1 is simply $1/2 \times (v_1 - \alpha)^+ + 1/2 \times \mathbb{E}\left[(v_2 - \alpha)^+\right]$. Importantly, the buyer's expected payoff from self-selecting into the first period price can always be expressed in a comparable form as a function of another threshold, say α': $1/2 \times (v_1 - \alpha') + 1/2 \times \mathbb{E}\left[(v_2 - \alpha)^+\right]$. To see it, simply relabel the first period price as $p = 1/2 \times \alpha' + 1/2 \times \left(\mathbb{E}[v_2] - \mathbb{E}\left[(v_2 - \alpha)^+\right]\right)$. This representation is a backbone of our construction, because the pair of threshold, namely α' and α, completely describes the buyer's incentives in the stopping problem.

It is easy to see that the buyer optimally choosing when to buy the good can guarantee himself the maximum of the aforementioned payoffs, that is

$$^1/_2 \times (v_1 - \min\{\alpha', \alpha\})^+ + {}^1/_2 \times \mathbb{E}\left[(v_2 - \alpha)^+\right]$$

We claim that the seller would optimally induce the same threshold in both periods, $\alpha' = \alpha$. On the one hand, the buyer never buy the good at the first date when $\alpha' > \alpha$, thus there is no loss to assume that $\alpha' \leqslant \alpha$. On the other hand, $\alpha' < \alpha$ is suboptimal as the seller can simultaneously decrease α, which improves efficiency, and extract all newly created surplus upfront by increasing M.

Finally, we argue that M is used purely for surplus extraction. Indeed, increasing M beyond $\mathbb{E}\left[(v_2 - \alpha)^+ | v_1 = 0\right]$ discourages the buyer with $v_1 < \alpha$ to participate effectively turning our pricing instrument into a forward sale: the buyer either buy at date 1 or does not buy at all, which is clearly suboptimal.

Conclude that the optimal pricing strategy is a member of simple family of contracts parametrized by the spot price: $M = M(\alpha) = \mathbb{E}\left[(v_2 - \alpha)^+ | v_1 = 0\right]$ and $p = p(\alpha) = \mathbb{E}\left[v_2 | v_1 = \alpha\right] - M(\alpha)$. We then maximize the seller's profit over $\alpha \in [0, 1]$ which yields the threshold $\alpha^* = {}^7/_{18}$ and an expected profit of approximately $0.387 > {}^1/_3$ to the seller. Trade takes place if and only if $\max\{v_1, v_2\} \geqslant {}^7/_{18}$, as shown in Figure 1c, and, perhaps, surprisingly the seller cannot obtain a higher profit by using any other selling procedure. Formally: $\max\{v_1, v_2\} \geqslant {}^7/_{18}$ is the best dynamic allocation for the seller. [3]

General model and global optimality. While some basic economic ideas are easily communicated through the two period model, in order to (i) generate testable implications for observable price paths, (ii) provide comparative statics for the changeability of buyer's valuation, (iii) quantify the gains from dynamic pricing, and (iv) understand general implications of time as an endogenous criterion of price discrimination, a richer model is required.

To that end, we consider a finite horizon continuous time model where the buyer's valuation is drawn from initially from a distribution F. At any given instant, the arrival of a Poisson shock with intensity λ changes the buyer's value. In case of an arrival, the new value is redrawn from F and in the absence of an arrival the value remains the same as before. This a natural extension of the example we presented above. The seller designs a dynamic pricing mechanism given by $\langle M, \mathbf{p}\rangle$, where M is an upfront payment which grants the buyer access to a sequence of prices $\mathbf{p} = (p_t)_{t=0}^T$.

M is of course chosen as the first part of the two-part tariff, with the objective of extracting buyer's surplus. After the payment of M, for a menu of prices $\mathbf{p}$, the buyer's strategy boils down to an optimal stopping problem. By choosing when to trade, the buyers self select themselves according to their evolving valuations in a third-degree price discrimination like exercise. The "fixed point" of this system implements the following allocation rule: there exists α^* such that the buyer trades at time t for price p_t^* iff $V_t \geqslant \alpha^* > \max_{s<t} V_s$. That is, the

[3]The proof can be found in Section ?? in the appendix.

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.[4]

The price path generated by this mechanism is shown to be smoothly increasing over time– it starts at some p_0^* and continuously increases to $p_T^* = \alpha^*$. Given "initial condition" $p_T^* = \alpha^*$, the price path is constructed by a backward system of indifference equations: at each point in time the buyer with value $V_t < \alpha^*$ strictly prefers to wait, and buyer with value $V_t \geqslant \alpha^*$ is made indifferent between waiting and buying at that point.

The closed form solution to the pricing problem is a novel addition to the literature on dynamic price discrimination. While its implications are intriguing in of themselves, these are provided further credence by establishing that the allocation rule it implements is actually the global (deterministic) optimum of the dynamic mechanism design problem. Invoking the revelation principle, we set up the general problem in continuous time, and state and prove a tight Myersonian characterization of incentive compatibility in terms of the dynamic envelope formula and monotonicity constraints. Then, in contrast to the literature, we establish the aforementioned allocation as the optimum. We would like to emphasize that the generic binding of global incentive constraints, is not a technical distraction, rather (analogous to multidimensional screening) it constitutes economic tradeoffs that are central to the understanding of dynamic pricing.

The simple economics of dynamic pricing. To understand the structure of dynamic pricing better, we offer a price theoretic interpretation of the dynamic mechanism design model.[5] For the renewal Markov (or Poisson) process we study, the ex ante distribution distribution of V_T, final value of consumption, is given by F. So, for a spot price α, trade happens with probability $1 - F(\alpha)$, which is regarded as the size of the market, represented by an inverse demand function: $D(p) = 1 - F(p)$. As is standard, any spot price $p_T = \alpha$ splits the area under the demand curve into three regions: consumer surplus, producer surplus and deadweight loss. The seller will chose the appropriate threshold α that maximizes producer surplus.

Next, by allowing the seller to charge an upfront payment M, at the intensive margin, some part of the consumer surplus is transferred to producer surplus, and at the extensive margin, by allowing the seller to reoptimize the threshold, total welfare is increased by expanding the region of trade. More specifically, $D_0(p) = \mathbb{P}(V_T \geqslant p | V_0 = 0)$ refers to the demand function of the lowest initial type and the consumer surplus for this demand function at "price α" is extracted as upfront payment.

Finally, by introducing a sequential price path, some part of the deadweight loss is transferred to producer surplus by adding buyers whose initial values could be higher than the threshold α but eventual value for consumption turns out to be lower. The relevant demand

[4]In the two period example presented above $\alpha^* = {}^7\!/\!{}_{18}$. More generally, α^* always exists and it is unique under standard regularity assumptions such as the monotone hazard rate.

[5]This exercise is inspired from **?**'s price-theoretic interpretation of optimal auctions alá Myerson [1981].

function here turns out to be $D_m(p) = \mathbb{P}\Big(\max_{t\geq 0} V_t \geq p\Big)$. The difference in the deadweight loss between D and D_m is precisely the added region of trade, and since the seller binds all the self-selection constraints, the buyer is indifferent and the seller gains all of the added surplus. By optimizing over the threshold with the added instruments of upfront payments and sequential prices, and total area of trade is again expanded through the extensive margin.

In conclusion, for a fixed threshold α, dynamic pricing allows the seller to extract some consumer surplus, and expand the region of trade, drawing in on deadweight loss. The optimal choice of the threshold, α^*, simultaneously maximizes the three components of the seller's profit– standard producer surplus, the part of consumer surplus extracted as upfront payment, and the part of deadweight loss recouped through market segmentation.

Bells and whistles. The basic analysis described above can be extended in various conceptual and theoretical directions. We report two such endeavors in the paper. First we show that two alternate mechanisms also implement the optimal allocation– refunds and subscription. In the former, the seller offers two sequences of prices $\langle \mathbf{p^n}, \mathbf{p^r}\rangle$, non-refundable and refundable respectively. The buyer is incentivzed to buy the good at the start at the refundable price p^r_0, and when first instant at which $V_t \geqslant \alpha^*$, the exercises the refund at price p^r_t and switches to the non-refundable price p^n_t. The subscription contract uses a combination of two continuous payments with pricing $\langle \mathbf{dm^s}, \mathbf{dm^p}, \mathbf{p^s}\rangle$. The buyer starts as a regular subscriber paying dm^s_t per unit of time. Then at the the first instance with $V_t \geqslant \alpha^*$, the buyer is incentivized to pay p^s_t and switch to the premium service which costs dm^p_t per unit of time.

Second, we extend our model to allow for linear costs for the seller. Suppose the seller incurs a cost c for producing the good. Then, if c is above threshold α^*, we have a problem: for $V_T \geq \alpha^*$ but $V_T < c$, trade would take place even if it is inefficient. It turns out that a simple fix to our mechanism restores optimality. The seller still put a menu $\langle M, \mathbf{p}\rangle$ where now the buyer who bought the good at time t can always return at the final date and receive c back. We solve for the optimal mechanism where the buyer gets to consume the good if $\max_{t\geqslant 0} V_t \geqslant \alpha^c$ and $V_T \geqslant c$.

We also illustrate the appeal of our dynamic pricing approach by (i) constructing a stochastic mechanism that improves upon the deterministic optimum, and (ii) solving a related model where the information change is exogenously restricted by the number of permissible Poisson arrivals. Both these problems are quite impenetrable through the standard approach of direct revelation mechanisms– in the first we have to construct a dynamic random mechanism and in the second the one-shot deviation property of incentive compatibility does not hold. In both cases, starting out with prices as opposed general dynamic mechanisms allows us to make progress on these hard problems.

Contribution to the literature. At the core, this paper is about price discrimination based on evolving valuations of the buyer which are private information. This literature has seen burgeoning interest over the last two decades, broadly under the rubric of dynamic contracts and mechanism design. Within that line of study, we look at the sale of a single timed good, which is referred to as sequential screening.[6]

The papers most closely related to our work are Courty and Li [2000], Esö and Szentes [2007], Deb [2014], and Kruse and Strack [2015]. The first two papers look at two period models where the good is sold at the end of the second period, and the buyer has some signal about his value in the first. They use specific implementations, that in the case of Courty and Li [2000], viz. refund contracts, is generalized here to longer time horizons, and that in the case of Esö and Szentes [2007], viz. European options, cannot be ported easily to general time horizons. Deb [2014] looks at the sale of a durable good in an infinite horizon framework with at most one Poisson shock and partially characterizes the optimum. Kruse and Strack [2015] looks at a finite horizon framework and characterizes allocations that can be implemented using threshold policies akin to our analysis.

Relative to this existing literature, the novelty of our paper is that it identifies a tractable environment in which (i) dynamic pricing plays the dual role of extracting buyer's surplus and sorting buyers endogenously along timing of purchase, (ii) the optimal (deterministic) contract can be precisely determined even though the first-order or local approach generically fails, (iii) a simple (and arguably intuitive) set of pricing instruments achieves the optimum for arbitrary time periods, and (iv) the optimal contract can be described in the familiar price-theoretic taxonomy.

2 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. We have in mind an airline ticket or a hotel booking for a specific date. Time is continuous and indexed by $t \in [0, T]$. For simplicity, we assume that both players do not discount payoffs. The buyer's valuation or more specifically signals of the actual valuation for the good follow a stationary Markov renewal process: $V_t = X_{N_t}$, where N_t is a Poisson process with an intensity λ, and X_n is a sequence of iid samples from a distribution F.

Physically, the process works as follows: a value V_0 is drawn at time zero from a distribution F. Then at each instant in time, V_t is the buyer's value 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– say a family emergency or an important meeting at work; 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. It is assumed that $V_t \in [0, 1]$ for all t, and that F admits a strictly positive density f over its support. We will often write V^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

[6] A richer discussion of the related literature is provided in Section 9. Here we give a quick overview of the ideas we build on, and the novelty of our approach.

structure of changing values prior to the realization of the actual payoff follows the sequential screening literature (see for example Courty and Li [2000]).

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 $\mathbf{p} = (p_t)_{t=0}^T$.[7] 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 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.[8] 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.

In familiar price theoretic vocabulary, the idea behind this pricing strategy is to provide the seller with a combination of *two-part tariffs* and *second degree price discrimination*. The upfront payment M and the set of prices $\mathbf{p}$ constitute the two parts of the tariff. Unpacking $\mathbf{p}$ gives us second degree price discrimination: the menu is designed to segment the market along the time dimension whereby buyers endogenously sort themselves according the evolution of their values.

3 A dynamic pricing strategy

For any 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 a problem such as this 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 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 sequences of price $(p_t)_{t=0}^T$ to maximize her expected profit.

A technical point to keep in mind is that in the solving the seller's problem, we will break

[7] 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 expresses as $\mathbf{p} = p^T = (p_t)_{t=0}^T$. The future set of prices will be denoted by $p^{[t,T]} = (p_s)_{s=t}^T$.

[8] 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.

buyer's indifference to stopping or not in favor of the seller. This is the "right" assumption in the sense that there is always a way to alter prices by a small amount and get a unique implementation with profits arbitrarily close to the optimal ones. So our results are not fragile to the assumption.

3.1 Optimal dynamic pricing

Stated in its full generality, the solution to the buyer's (and hence the seller's) problem seems complex. 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 $\beta_t\left(V_t, p^{[t,T]}\right) \in [0,1]$ such that if $p_t \leqslant \beta_t$, buy, else wait; or inversely, a threshold $\alpha_t\left(p^{[t,T]}\right) \in [0,1]$ such that buy if $V_t \geqslant \alpha_t$, 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 the 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 $V_t \geqslant \alpha^* > \max_{s<t} V_s$.

A complete proof of the theorem can be found in the appendix where it is shown that the optimal pricing strategy and the buyer's best response to it belong to a simple of class of contracts which 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 arbitrary spot price α, the seller will optimally (backward) construct the sequences of prices and upfront payment as listed in the theorem. The buyer in turn will stop (and trade) whenever $p_t \leq \beta_t$ (or $V_t \geq \alpha_t$). Since the $\mathbf{p}$ is itself a function of α, the buyer's threshold response to it, $(\alpha_t)_{t=0}^T$, is also simply a function of α. The bulk of work in the proof is to show (i) that the fixed point of the system settles onto a history independent threshold for the buyer, and (ii) that fixed threshold is the same as the final spot price $p_T = \alpha$.[9]

The allocation rule implemented by this construction has the following form: the first instant t at which the value of the buyer is above the threshold α, the buyer should trade at the price p_t, and in case $\max_{t\geqslant 0} V_t < \alpha$, there is no trade. The seller's expected profit from this pricing strategy can be decomposed into three terms:

$$\Pi(\alpha) = \underbrace{\alpha\left[1 - F(\alpha)\right]}_{\text{spot pricing}} + \underbrace{\left(1 - e^{-\lambda T}\right)\int_\alpha^1 \left[1 - F(v)\right]dv}_{\text{upfront payment}} + \underbrace{\left(1 - e^{-\lambda[1-F(\alpha)]T}\right)\int_0^\alpha v dF(v)}_{\text{dynamic market segmentation}} \qquad (\star)$$

[9]In the context of two period model depicted in Figures 1c and 1d, this means that the no trade region is (i) a lower rectangle, and (ii) that rectangle is actually a square.

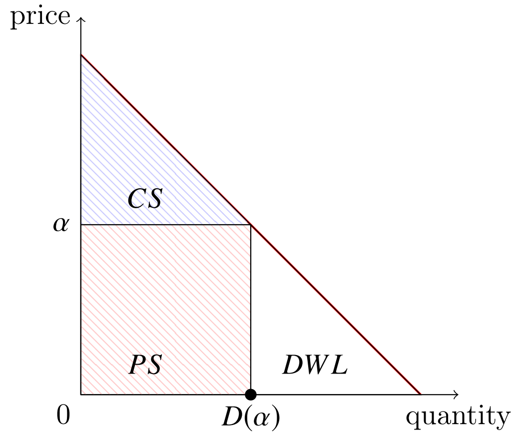

(a) Spot pricing

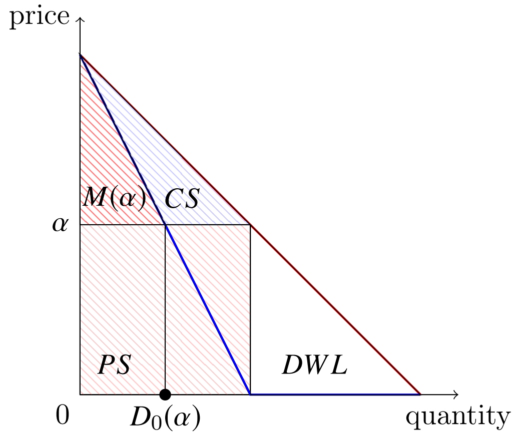

(b) Adding upfront payment

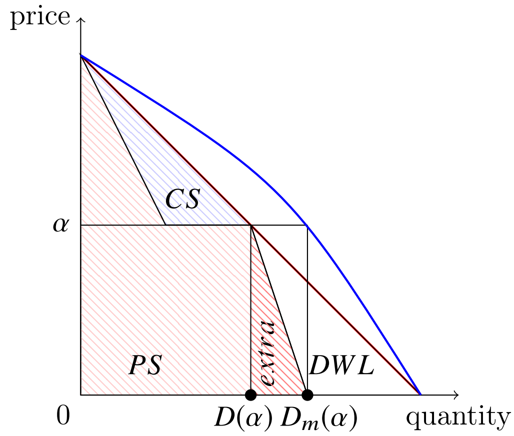

(c) Adding dynamic prices

Figure 3: Decomposition of total surplus from trade into CS, PS and DWL

The seller's optimization problem reduces to one dimension: she chooses a threshold given by by $\alpha^* = \arg\max_{\alpha\in[0,1]} \Pi(\alpha)$.[10] In what follows, we provide a brief explanation of how each of three individual components stack up to constitute the seller's profit expressed in Equation ($\star$).

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": $D(p) = 1-F(p)$, where $D(p)$ is quantity demanded at price p.[11] Figure 3 depicts the buyer's demand function with "quantity" on the x-axis and price on the y-axis. In Figure 3a, for a fixed spot price α,

[10] In Corollary 1 we show that α^* is unique under standard restrictions on value distribution, for example, $v \mapsto \frac{1-F(v)}{f(v)}$ is non-increasing.

[11] 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.

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 $V_T < \alpha$.[12] The seller's expected payoff from this static pricing strategy is $\alpha(1 - F(\alpha))$, which is the first term of Equation ($\star$).

In this static pricing mechanism, even the buyer with the lowest possible value 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 to be the expected payoff of the lowest value buyer:

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

which 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 depicted in Figure 3b. 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.[13] By construction, the seller can extract $M(\alpha)$ from the buyer without distorting his participation decision at time $t = 0$. Therefore, in Figure 3b, $M(\alpha)$ represents the transfer of an erstwhile component of consumer surplus to what is now a part of the producer surplus.

Now 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. As mentioned before, it turns to be optimal to substitute $\alpha_t = \alpha$ for all t. So the buyer with $V_t \geqslant \alpha$ is incentivized to purchase the good early, and trade takes place whenever $\max\limits_{t \geqslant 0} V_t \geqslant \alpha$.

To be precise the seller makes all buyer types $V_t \geqslant \alpha$ indifferent between making a purchase at t and waiting until the terminal date. Therefore, in this final step, increased region of trade moves some part of the erstwhile DWL to the producer surplus. The magnitude of this extra profit for the seller exactly equals the third term in Equation ($\star$)

The construction of dynamic market segmentation can be visualized in Figure 3c; the blue curved line depicts a new demand function given by $D_m(p) = \mathbb{P}\left(\max\limits_{t \geq 0} V_t \geq p\right)$.[14] To understand

[12] 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.

[13] 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.

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

a size of dynamic gains, rewrite a change in DWL as it follows:

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

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) is depicted in Figure 3c as "extra" is transferred from DWL to PS due to dynamic pricing.[15]

3.2 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:

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

This equation can only have 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 really matters is the "normalized time" λT.[16] Comparative statics for the optimal threshold are presented in the following result.

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* $v\mapsto vf(v)$ *are nondecreasing.*

(b) α^* *equals the static optimal fixed price whenever* $\lambda T=0$*.*

(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 model becomes static, or as $\lambda\to 0$, that is as values become perfectly persistent, the optimal spot price converges to the optimal (static) fixed price. So, for example, with a uniform distribution, the best the seller can do is post a price of $1/2$. Further, part (c) says that α^* is positive and strictly decreasing in normalized time. Thus, 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

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

[16] One of the reasons for using a continuous time model is precisely that the expression for total profit given by Equation ($\star$) is straightforward, and hence calculations with respect to time t and persistence λ become easier, and can be reported in closed form.

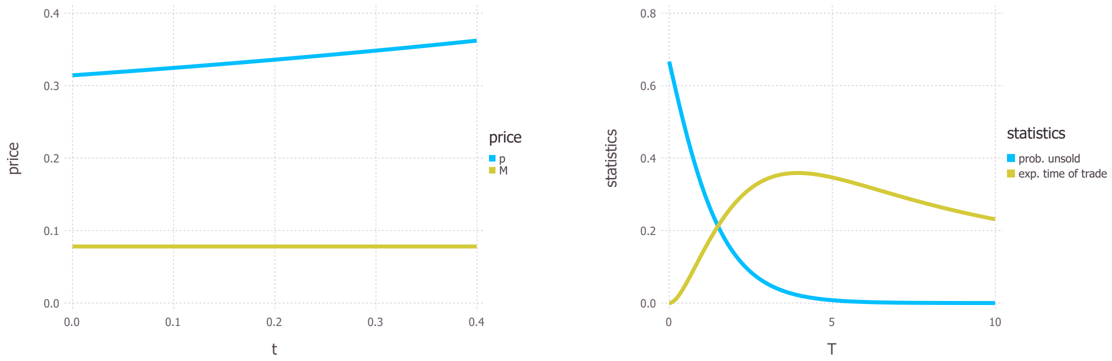

(a) Prices

(b) Probability of sale

Figure 4: Evolution over time. $F(v) = v$, $\lambda = 1$.

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 does throw light on the fact that the ability to commit to dynamic prices allows the seller to sequentially segment the market, thereby increasing total surplus, and extracting a larger fraction of it as profit.

Coming to prices, 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, hence, we term it the *forward discount*. As, shown in Figure 4a, prices for trade increase steadily from some initial value p_0^* to the final spot price p_T^*, and the forward discount decreases steadily from its highest initial value to zero.[17]

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

For the optimal pricing strategy, the seller may be interested in understanding the expected probability of sale of the good at any time t. Or for a given data set of price and selling times, an outside analyst may want to estimate some counterfactuals. To help with these questions, we want to compute the distribution of the random variable $v_t = \max_{s \leqslant t} V_s$, which is defined by

$$\hat{F}_t(v) = F(v)e^{-\lambda[1-F(v)]t}$$

It can then be noted that the probability that the good is sold by time t is $1 - \hat{F}_t(\alpha^*)$. A set of results complimentary then follow.

[17]Note also that the lower horizontal line in Figure 4a represents the upfront payment, M^*.

Corollary 3.

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

(b) Expected time of sale, conditional on sale, is non-monotone, specifically:

$$\lim_{\lambda T\to 0}\int_0^T td\left[\frac{1-\hat{F}_t(\alpha^*)}{1-\hat{F}_T(\alpha^*)}\right] = \lim_{\lambda T\to\infty}\int_0^T td\left[\frac{1-\hat{F}_t(\alpha^*)}{1-\hat{F}_T(\alpha^*)}\right] = 0$$

The first part is the exact counterpart of Corollary 1c: the probability of no trade is strictly decreasing in normalized time λT. The second part of the result provides a testable implication: If the analyst observes a large enough set of trades with recorded sale times, she can test whether the data fits the curve of expected sale times, as shown in Figure 4b.

While the pricing mechanism studied here seems intuitive enough to merit a theoretical analysis in its own right, Theorem 1 and the results that have followed after gain more relevance in the light of the fact that this mechanism is actually globally optimal, that is, no other set of "standard" pricing instruments can ensure a higher profit for the seller. We now turn to establishing this optimality.

4 Optimality: a mechanism design approach

In this section, we establish that the seller can not 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 it without 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 cumulative payment at time t.[18] The agent reports his marginal information each period; that is, his strategy prescribes a report $\hat{V}_t \in [0,1]$ at each instance of time.[19]

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

$$\mathbb{E}\left[V_T Q\left(V^{[0,T]}\right) - \int_0^T dP_t\left(V^{[0,t]}\right)\right] \geqslant \mathbb{E}\left[V_T Q\left(\hat{V}^{[0,T]}\right) - \int_0^T dP_t\left(\hat{V}^{[0,t]}\right)\right] \quad \forall \hat{V}$$

Fix a history $\hat{V}^{[0,t)}$ and $V_t = v$, then define buyer's continuation payoff as

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

[18] 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.

[19] The process of reports is cadlag, 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 and t

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

The following lemma completely characterizes incentive compatibility using this latter formulation.

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:*

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

$$v \mapsto Q\left(V^{[0,t)}, v^{[t,T]}\right) \qquad \textit{is non-decreasing} \tag{C}$$

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

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]. As can be noted in the proof, it is actually not immediately obvious how to extend the appropriate discrete time characterization directly in continuous time; this is a technical innovation novel to this paper.[20]

Analogous to its static cousin, Lemma 1 pins down the space of allocations which 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; and Equation (IM) is the integral monotonicity constraint which captures "global incentives" arising out of the multidimensionality of the mechanism design problem.

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

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

This restriction says that the buyer can not 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

[20] Here we build on Bergemann and Strack [2015], who have stated necessary conditions in a continuous time screening model, and provided sufficiency conditions to check for optimality.

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.*

Another way of stating this result is that given $\langle M, \mathbf{p} \rangle$, the principal cannot benefit from any additional pricing instruments. Since, the price vector $\mathbf{p}$ depends only on time, and not on the realization of buyer's values as in a standard mechanism design problem, the proof of the above result becomes non-trivial. In terms of the allocation rule, the optimal contract implements no trade if and only if $V_t \leq \alpha^* \; \forall \; t \leq T$, where again α^* in unique threshold for all possible sequence of values $V^{[0,T]}$. Analogous to the two period example (see Figures 1c and 1d), Theorems 1 and 2 establish that the optimal contract in the general continuous time model produces no trade in the bottom convex poltytope with a continuum of equal edges of length α^*. In the next section we show there exist other intuitive pricing mechanisms that implement the optimal allocation rule.

5 Alternate implementations

The key economic ideas from Sections 3 and 4 can be stated as follows: a combination of two-part tariff and second-degree price discrimination deliver the optimal profit in the underlying dynamic mechanism design problem. In our pricing mechanism, M and $\mathbf{p}$ formed the two parts of the tariffs and $\mathbf{p}$ was further unpacked to sort the buyers in the spirit of second-degree price discrimination. In this section, we show that there are two other dynamic pricing strategies- refund and subscription- which implement the same allocation rule, and hence achieve the optimal profit for the seller.

5.1 Refund

A *refund contract* constitute two sequence of prices: a non-refundable sequence and refundable sequence. A non-refundable sale is final, in contrast a refundable sale can be reimbursed for the prevailing refundable price. Formally, the seller commits to two price schedules $\mathbf{p^n} = \left(p_t^n\right)_{t=0}^T$ and $\mathbf{p^r} = \left(p_t^r\right)_{t=0}^T$. 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 at T. The following result states the optimal pricing mechanism and the buyer's best response to it.

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

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

(a) Optimal refund contract for $F(v) = \sqrt{v}$ (b) Optimal subscription contract for $F(v) = \sqrt{v}$

Figure 5: Alternate implementations, $\lambda T = 0.4$.

achieves the same expected profit as $\langle M^*, \mathbf{p}^* \rangle$*. At the start, the buyer pays the non-refundable price* p_0^r*. Then, he exercises the refund* p_t^r *at time* t *and switches to 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 $p_T^r = p_T^n = \alpha^*$. Figure 5a plots the refund contract (and the baseline price p_t^*) for a power distribution and normalized time of consumption $\lambda T = 1$. The economic idea behind the pricing structure relies on broadly the same principle: refundable price extracts the the buyer's surplus and non-refundable price endogenously sorts the market along timing of purchase.

It can be noted that the refundable price is always decreasing over time and the non-refundable price can be increasing or decreasing over time. In fact for the power distribution $F(v) = \sqrt{v}$, p_t^n is increasing over time, as shown in Figure 5a, and for the uniform distribution $F(v) = v$ it is decreasing. It makes intuitive sense that ticket prices would converge to the same level as we approach the date of consumption, but what is less obvious is why some class of prices would increase and others decrease over time.

The monotonicity of both sequence of prices in the refund contracts generates testable predictions. Should, for example, algorithmic prices for airline tickets or hotel stays increase or decrease over time? Should these paths be different for refundable and non-refundable prices? Our analysis is of course is missing capacity constraints– both airlines and hotels have fixed capacity that they can make available over time. However, it still allows us to separate the force of the seller segmenting the market along changing buyer valuations.

5.2 Subscription

The second alternate mechanism is to offer the buyer to become a regular subscriber which costs dm_t^s per instance of time or a premium subscriber which costs dm_t^p per instance of time. The regular service grants the buyer a right to purchase the good at the final date for a spot price α^*, in contrast the premium service awards the good for free. The buyer can become a regular subscriber only at the initial date. Then, he can either irreversibly discontinue his subscription or upgrade it to the premium level which costs p_t^s.

Proposition 2. *Let α^* be the threshold identified in Theorem 1. The mechanism $\langle dm^s, dm^p, p^s \rangle$ defined by*

$$dm_t^s = \lambda e^{-\lambda(T-t)} \int_{\alpha^*}^1 [1 - F(v)] dv dt, \quad dm_t^p = \lambda e^{-\lambda(T-t)} \int_0^1 [1 - F(v)] dv dt, \quad p_t^s = e^{-\lambda(T-t)} \alpha^*$$

achieves the same expected profit as $\langle M^*, p^* \rangle$. At the start, the buyer logs into the regular subscription. Then, he switches to the premium service at time t for price p_t^s if $V_t \geqslant \alpha^* > \max_{s<t} V_s$. If $\max_{t \geqslant 0} V_t < \alpha^*$, and pays dm^p thereafter.

Figure 5b plots the subscription mechanism for the power distribution and $\lambda T = 1$. At a technical level, this is the unique implementation for the optimal mechanism that binds the individual rationality constraint in each period. At a conceptual level, it connects to the widely observed subscription contracts that we observe in online platforms. While these are often motivated as experience goods, our model can be seen as a reduced form approach where the seller simply observes an exogenous stochastic distribution of taste changes in the past, as sets prices according to it.

6 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 implements the optimal contract. 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 3. *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 $\langle M^*, \mathbf{p}^\mathbf{c} \rangle$ 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 3 is the analog to Theorem 1 with added linear cost of production for the seller. Next, we show that the allocation rule selected by $\langle M^*, \mathbf{p}^\mathbf{c} \rangle$ is the optimal (deterministic)

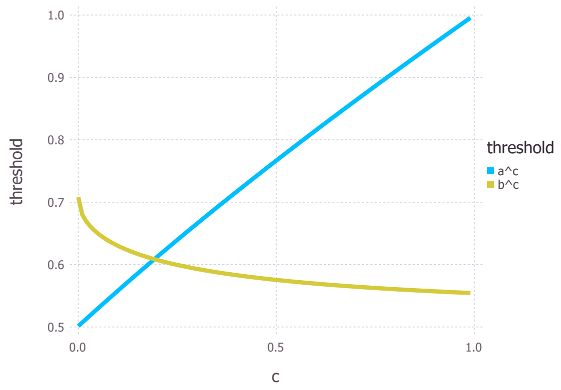

Figure 6: Thresholds α^c and $\beta^c = \frac{F(\alpha^c)-F(c)}{1-F(c)}$ as functions of c for $F(v) = \sqrt{v}$ and $\lambda T = 0.4$.

mechanism. The threshold α^c solves $\alpha^c = \arg\max_{\alpha \in [c,1]} \Pi^c(\alpha)$ which 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 4. *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 6 we plot the α^c and $\beta^c = {}^{[F(\alpha^c)-F(c)]}/_{[1-F(c)]}$, the latter is a measure of inefficiency of the mechanism expressed by the rate at which the threshold increases in the quantile of with an increase in c. This measure is motivated by the fact that efficiency prescribes trade whenever the buyer's value is above c. As can be seen from Figure 6, even though α^c increases in c, the mechanism becomes progressively more efficient. This is quite expected as for higher values of c, the seller would never serve the buyer with low valuations, thus gains from price discrimination are smaller and the tradeoff between efficiency and rent extraction is less stringent.

7 Gains from randomization

Characterizing randomization for static mechanism design can be quite complicated, and in dynamic mechanisms the task is only harder. Our approach here is to highlight the main economic message behind stochastic pricing. To that end, we use the language of price discrimination, where stochastic mechanisms can be thought of as the sale of fractions of the same good. Selling fractions is beneficial, because the seller can flexibly adjust terms of trade based on the history of past purchases. It provides her with additional information about changes in the buyer's valuation and improves 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. At 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 a 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 defined as

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

which simply replaces α^* with $\alpha^* - \varepsilon$ in the definition of p_t^*. Upon payment of $\tilde{M}$, the buyer gets access to two sequence of prices: 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. As a consequence, 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:

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

Moreover, the net change of seller's profit 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

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

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 a second order, it permits the seller to further 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 by the firm 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, the main conceptual innovation here is that the seller can buyback

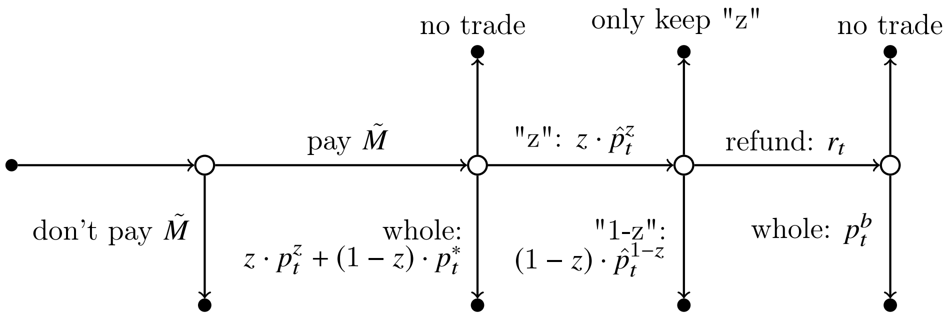

Figure 7: Timing of the fractional buyback pricing mechanism.

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 in this new allocation. The next result formally summarizes the construction.

Proposition 5. *Fix a small $\delta > 0$ and a small $\varepsilon > 0$. There exists a pricing mechanism $\langle \tilde{M}, \mathbf{p}^*, \mathbf{p^z}, \hat{\mathbf{p}}^{\mathbf{z}}, \hat{\mathbf{p}}^{\mathbf{1-z}}, \mathbf{p^b}, \mathbf{r}\rangle$ 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 $v \mapsto vf(v)$ 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 5 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 Limited informational change: the case of a single arrival

It is possible that the analyst (or econometrician) 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 time with with a given Poisson intensity. One way to enforce limited change of information is to have high 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 the same Poisson intensity λ. The rest of the model is identical.

Straightaway in the next result, we state the pricing mechanism that is "equivalent" to the baseline, adjusting for the new model of informational change. And, then we show that pricing mechanism implements the deterministic dynamic optimum.

Proposition 6. *Suppose the buyer's value can change at most once. There exists $\alpha^s \in [0,1]$ such that the optimal pricing strategy is the unique non-trivial solution to the following system:*

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

The buyer always makes the upfront payment and can make 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 $V_t \geqslant p_t^s$.

As before, the seller's problem is reduced to a one-dimensional optimization problem. We fix the spot price at α and choose the prices p_t^s so that the marginal buyer who's value $V_t = \alpha$

is indifferent between buying and continuing. In choosing the price path we focus exclusively on the the buyer who's value has not changed by time t, thus there is still some remaining uncertainty about the final value for trade. Intuitively, such prices incentivize the buyer to make a purchase only at the times of informational changes which is necessary to screen the buyer. In the proof we establish that these prices are pinned down for a fixed spot price by the differential equation given in Proposition 6 ensuring the marginal buyer is indifferent between accepting the trade and waiting until the last date:

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

The left-hand side is the buyer's gain from an immediate stopping, whereas the right-hand side is the payoff from not stopping until a new arrival.

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

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

The profit is continuous in α, because the prices p_t^s change continuously with α, and thus there exists an optimal threshold: $\alpha^s = \arg\max_{\alpha \in [0,1]} \Pi^s(\alpha)$.

Proposition 7. *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.*

Proposition 7 establishes that the simple pricing instrument achieves the highest possible profit in the set of all deterministic mechanisms the information changes are limited. Although, the result is parallel to one obtained in the model with an unlimited number of arrivals (see Theorem 2), the model itself is virtually unsolvable without an appropriate conjecture about the optimal mechanism. The reason is that the stochastic process of valuations with a single arrival is not Markov anymore, thus the incentive constraints are more complicated. In particular, it is not longer true that "on-path" incentive compatibility implies that the buyer wants to report his information "off-path". As a result the buyer's dynamic deviations are also needed to be taken into an account, for example, one where the buyer misreport his value at $t = 0$ and then misreports at some time t that he had an arrival.

To actually prove the result of Proposition 7 we make use of the optimal indirect pricing mechanism described in Proposition 6. First, we consider the relaxation of the buyer's problem using as the conjecture the allocation of Proposition 6 identified by the pricing approach. Specifically, the relaxation has only the following subset of incentive constraints: a) the buyer who had an arrival does not want to misreport it and b) the buyer with $V_0 = \alpha^s$ does not under-report his value at $t = 0$ and then stop at some t reporting $\mathbb{E}[V_T | V_t = \alpha^s]$. Second, we show that the surplus maximizing allocation subject to these two sets of constraints coincides with the one implemented by the optimal pricing strategy, therefore it is incentive compatible.

9 Final remarks

Related literature. Price discrimination has a rich history in economics, starting at least from Pigou [1920].[21] Its invocation as the instrument of market segmentation is now almost axiomatic. On the wide prevalence of price discrimination, Varian [1980] famously wrote: "Economists have belatedly come to recognize that the 'law of one price' is no law at all. Most retail markets are instead characterized by a rather large degree of price dispersion." We specifically study markets which feature prices that change over time.

One of the first papers to analyze dynamic pricing is Stokey [1979]. It engineers a simple environment where upon committing to a sequences of prices the monopolist finds it optimal to charge the optimal (static) monopoly price every period. Thus, the product is sold as soon as it is introduced in the market. It is a stark benchmark result on the unnecesisty of (dynamic) price discrimination. We depart from its main structure in many ways, most importantly by assuming that the buyers' valuations are private information and changeable over time.

The role of price discrimination when buyers have private valuations developed complimentary with the literature on mechanism design. While the "cross-section" gained much prominence through the study of bundling (see for example Adams and Yellen [1976], McAfee, McMillan, and Whinston [1989], Rochet and Choné [1998], and more recently **?**), the "time series" of it came to be studied later with a rise in interest in dynamic contracts and dynamic mechanism design. Our paper contributes on this latter body of work, reviewed in recent surveys by Krähmer and Strausz [2015] and Bergemann and Välimäki [2019]. Here we discuss a subset of papers which speak directly to our results.

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]. Almost all papers study two period models which use specific Markov processes to explain the use of evolving information as the primitive of price discrimination. We build on this agenda.

We explicitly solve for the optimal dynamic pricing mechanism $\langle \mathbf{M}, \mathbf{p} \rangle$, which is different from the implementation used in Courty and Li [2000], and then show that the indirect mechanism implements the optimal (deterministic) contract. This involves two novel and non-trivial optimization problems (dynamic pricing and dynamic mechanism design) that are then shown to be equivalent in allocations and payoffs. Courty and Li [2000] focussed on refund contracts to achieve the optimum, the appropriate generalization of which is presented in Section 5.1.

Previously the literature has also focussed on option contracts in sequential screening. These are akin to offering a contract $\langle \mathbf{M}, \alpha \rangle$ in our framework, where $\mathbf{M} = (M_v)_{v \in [0,1]}$ and $\alpha = (\alpha_v)_{v \in [0,1]}$ is a menu of upfront payments and a final strike prices. Esö and Szentes

[21] Interestingly, the classic political economy treatise from approximately the 2nd century BC called Arthshahtra (literally meaning the *study of money*), typically attributed to Kautilya, talks in some detail about offering a menu of loans varied by interest rates and maturity structure so as to discriminate potential consumers on the basis of their demand, risk, and capacity to repay (see Kautilya [1992]).

[2007] study a two period problem and use a menu of such options to implement the optimal allocation. In the first period, buyer's types are simply screened by their choice of option from the menu, and in the second period by whether they decide to exercise the option. It turns out in our setting a simpler menu is sufficient, which screen buyers' types only using timing of purchases. In fact, the implementation of Esö and Szentes [2007] can not be easily ported to a model with more than two periods, whereas our dynamic pricing strategy has a clear generalization to any time horizon.[22]

Deb [2014] studies the sale of a durable good in an infinite horizon where there is only one Poisson shock which changes the buyer's type. A partial characterization of the optimum is provided with an intuitive implementation of introductory pricing. His model is akin to the model we studied in Section 8.

In its framing of buyer's response as an optimal stopping problem, the paper builds on Kruse and Strack [2015], which considers the sale of a single good in a discrete time framework with changing buyer valuations. Using the envelope theorem, it characterizes the the class of allocations that can be implemented using threshold policies. We identify a tractable environment where the deterministic optimum is a threshold policy threshold policy, and it takes a very simple form.

At a conceptual level, we bring the literature on sequential screening, and dynamic mechanism more generally, closer to the taxonomy of price discrimination. The formulation of the seller's problem as a combination of two part tariff and second degree price discrimination pushes towards a classical industrial organization interpretation of problems. The explanation of the profit function as a combination of producer surplus, transferred consumer surplus, and recouped deadweight loss is a novel price theoretic interpretation of dynamic mechanism design.

10 Appendix

The missing proofs from the main text are provided in this section.

10.1 Solution to the two-period example

In this section we complete the two-period example presented in the introduction and show that our pricing strategy implements the optimal dynamic mechanism.

To begin, a direct dynamic mechanism is a history-dependent pair of an allocation $q \in \{0, 1\}$ and dynamic payments (p_1, p_2). In a mechanism the buyer reports his marginal information $\hat{v}_t$ in every period. As usual, by the revelation principle there is no loss from looking only at the set of direct incentive compatible mechanisms where the buyer cannot strictly

[22] Boleslavsky and Said [2013] study a model where the buyer's types are i.i.d conditional on the first period information. Thus, as in Esö and Szentes [2007] the first period and current type are sufficient statistics for allocative distortions and can be implemented by a menu of options. Our pricing strategy is shown to be useful even when this sufficient statistic is not available.

gain by misreporting his valuations. In what follows we develop a complete characterization of this set which is in the spirit of Myerson [1981] and Pavan et al. [2014].

Write $U_2(\hat{v}_1, v_2) = v_2 q(\hat{v}_1, v_2) - p_2(\hat{v}_1, v_2)$ for the buyer's second period payoff when v_2 is reported truthfully and v_1 was potentially misreported as $\hat{v}_1$. Note that 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 the incentive compatibility at $t = 2$ reads as it follows:

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

This constraint is quite standard and well-studied in the static mechanism design literature. 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 $U_1(v_1) = \mathbb{E}\left[U_2(v_1, v_2)\right]$. Since the second period incentive constraint ensures that the buyer always revert 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 $q(v_1, v_1)$. 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, 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. By linearity, this profit is a difference between surplus and the 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

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

The expression for the buyer's payoff is obtained by integration by parts using the fact $U'(v_1) = {}^1\!/_2 \times q(v_1, v_1)$ that for any incentive compatible mechanism. It is important to note that the buyer's payoff only depends on the allocation in the case of constant valuation, that is $v_1 = v_2$.

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) $v_2 \mapsto q(v_1, v_2)$ being non-decreasing.

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 $v \geqslant \alpha$. 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 $q(v_1, v_2) \leqslant q(v_1, v_1) = 0$.
- For $v_1 < v_2 < \alpha$: this is implied by integral monotonicity at $t = 1$:
$$\int_{v_1}^{v_2} q(v_1, v) dv \leqslant \int_{v_1}^{v_2} q(v,v) dv = 0$$

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 no higher than

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

It is routine to verify that the upper bound is maximized at $\alpha^* = {}^7\!/\!_{18}$ which yields exactly the same profit as the optimal pricing strategy which we introduced in the example.

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

Proof of Theorem 1. Before going to the proof, it is useful to restate the problem. Since the gain process is Markov in time and a 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 as

stop and trade/not trade at $s \in [t, T]$ whenever there is no arrival in $[t, s]$,

if there is an arrival at $r \in [t, s]$, then continue with a strategy corresponding to V_r

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:

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

Note that $(\cdot)^+$ is omitted in the above equation, the reason is that the continuation value 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 which he is facing after making the upfront payment. At the initial date, the buyer knowing V_0 decides either to pay the upfront fee and receive $W_0(V_0) - M$ or opt out.

We prove this theorem in three steps. First, we find an upper bound on the seller's expected profit as a function of the buyer's expected payoff. Second, 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:

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

In these new notations, the gain process can be rewritten as $G_t(v) = e^{-\lambda(T-t)}(v - \alpha_t) + W_t(0)$. By definition, the value function must dominate the gain process, thus

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

Clearly, 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. This will be then used to bound gains from trade and buyer's expected payoff.

There are two cases to look. First, suppose that the buyer with the lowest valuation has a strict incentive to trade. Formally, the max in Equation (†) is achieved at the first term:

$$W_t(0) > \lambda \int_t^T e^{-\lambda(s-t)}\mathbb{E}\left[W_s(V_s)\right] ds$$

Since $e^{-\lambda(s-t)}\left[G_s(v) - G_s(0)\right] = e^{-\lambda(T-t)}v$ is independent of s, the buyer with a higher valuation prefers to follow the same strategy. In other words, no trade is strictly dominant, which implies that $W_t(v) - W_t(0) = e^{-\lambda(T-t)}v$.

Next, suppose that $\inf\limits_{s\geqslant t}\alpha_s > 0$. We claim that the buyer with the lowest valuation prefers no trade, equivalently the max in Equation (†) is achieved at the second term:

$$W_t(0) = \lambda \int_t^T e^{-\lambda(s-t)}\mathbb{E}\left[W_s(V_s)\right] ds$$

To begin, note that $p_T = \alpha_T > 0$, thus the buyer should not trade at the final date: $W_T(0) = 0$. Now, by the way of contradiction, assume that the buyer stops at $s \in [t,T)$ or a moment later. This yields $\limsup\limits_{\varepsilon\downarrow 0} G_{s+\varepsilon}(0) = W_s(0) - e^{-\lambda(T-s)}\limsup\limits_{\varepsilon\downarrow 0}\alpha_{s+\varepsilon}$. By our assumption $\inf\limits_{s\geqslant t}\alpha_s > 0$, thus the gain from stopping is strictly less than $W_t(0)$ - a contradiction.

Since the buyer with the lowest valuation prefers to continue until the final date, $W_t(0)$ can be rewritten as

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

Subtract this from ($\star$) for $v > 0$ and take the sup to obtain the following estimate:

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

Note that this estimate covers the first case as well where we had $W_t(v) - W_t(0) = e^{-\lambda(T-t)} v$, because the necessary condition for having a strict incentive to trade is $\inf\limits_{s \geqslant t} \alpha_s = 0$.

Next, we consider the buyer's decision to make the upfront payment. To save on notations, let $\alpha = \inf\limits_{t \geqslant 0} \alpha_t$. It is also convenient to reparametrize M as $M = W_t(0) + e^{-\lambda(T-t)} \beta$ with $\beta \in \mathbb{R}$. Clearly, it suffices to look at $\beta \geqslant 0$, and 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 $W_t(v) - W_t(0) = e^{-\lambda(T-t)} \left(v - \inf\limits_{s \geqslant t} \alpha_s \right)$, thus the buyer with value $V_t = v$ will agree to make a purchase at t or a moment later only if $W_t(v) = \sup\limits_{\varepsilon \downarrow 0} G_{t+\varepsilon}(v)$. Equivalently: $v \geqslant -\inf\limits_{s \geqslant t} \alpha_s = \inf\limits_{\varepsilon \downarrow 0} \alpha_{t+\varepsilon}$. In particular, there is no trade whenever $\max\limits_{t \geqslant 0} V_t < \alpha$ and the buyer's expected net payoff is $\mathbb{E}\left[W_0(V_0) - W_0(0)\right]$. Combine these two to obtain the following bound on the seller's expected profit:

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

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

$$\underbrace{\mathbb{E}\left[V_T \mathbb{1}(V_0 \geqslant \beta + \alpha) \right]}_{\text{maximal surplus}} - \underbrace{\mathbb{E}\left[(V_0 - \beta - \alpha)^+ \right]}_{\text{buyer's expected payoff}}$$

Clearly, this is at most $\Pi(\alpha + \beta)$.

Step II. At this step, we show that for any $\alpha \geqslant 0$ there is a pricing strategy which 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 $\langle M, p \rangle$ 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:

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

Note that $\mathbb{E}\left[(V_T-\alpha)^+|V_t = v\right]$ is the martingale dominating the gain process. On the other hand, $\mathbb{E}\left[(V_T-\alpha)^+|V_t = v\right]$ is the value from stopping only at the final date whenever $V_T \geqslant \alpha = p_T$. By definition, the value function is the smallest supermartingale dominating the gain process, 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: stop at the first instance of $W_t(v) = G_t(v)$, that is $V_t \geqslant \alpha > \max_{s<t} V_s$. Conclude that trade happens whenever $\max_{t\geqslant 0} V_t \geqslant \alpha$, so the seller can obtain the profit of $\Pi(\alpha)$ by using the aforementioned pricing scheme.

Step III. To conclude the proof, we derive the threshold α^* which maximizes $\Pi(\alpha)$. First of all, observe that the buyer's expceted payoff is simply $\int_\alpha^1 [1-F(v)]\,dv$.

The expression of expected gains from trade is a bit more complicated. To compute the surplus efficiently, we need to introduce several auxiliary objects. Let $\tilde{T}$ be the time of 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 $\max_{s\leq t} V_s$ 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 $\tilde{T}=0$ are $\displaystyle\int_\alpha^1 v dF(v)$, and conditional on some $\tilde{T}\neq 0$:

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

By the law of iterated expectations:

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

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

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

The optimal threshold α^* solves $\max_{\alpha\in[0,1]} \Pi(\alpha)$. The first 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 infinite, 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 which 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 the unique 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 $[0,\hat{\alpha}]$ which will imply uniqueness of α^*. 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 increasing 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 $\lambda_1 T_1 < \lambda_2 T_2$. Observe that the seller's profit can be rewritten in an integral form as it follows:

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

It follows that

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

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 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.

□

Proof of Corollary 3. Part (a). Note that $\hat{F}_T(v)$ is strictly decreasing in λT for any 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:

$$\lim_{\lambda T \to 0} \int_0^T t d \left(\frac{1-\hat{F}_t(\alpha^*)}{1-\hat{F}_T(\alpha^*)} \right) = 0$$

□

10.3 Proofs for the general dynamic mechanism design problem

Proof of Lemma 1. By the standard dynamic programming argument, 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 $\varepsilon > 0$:

- $V_t = v$ 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 $\varepsilon > T - t$ which 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

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

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

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

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 $\varepsilon' > \varepsilon$ 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:

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

By our assumption, Q is either 0 or 1, thus 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 $\max\limits_{t \geqslant 0} V_t < \alpha$, 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 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 $Q(V) = 0$.

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

- For $X_{N-1} > X_N$: $Q\left(V^{[0,\tau_{N-1})}, X_{N-1}^{[\tau_{N-1},T]}\right) = 0 \geqslant Q\left(V^{[0,\tau_{N-1})}, X_{N-1}^{[\tau_{N-1},\tau_N)}, X_N^{[\tau_N,T]}\right)$ by (C).
- For $X_{N-1} < X_N$: $\int_{X_N}^{\alpha} Q\left(V^{[0,\tau_{N-1})}, w^{[\tau_{N-1},T]}\right) dw = 0 \geqslant \int_{X_N}^{\alpha} Q\left(V^{[0,\tau_{N-1})}, X_{N-1}^{[\tau_{N-1},\tau_N)}, w^{[\tau_N,T]}\right) dw$ 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 $\mathbb{E}\left[V_T \mathbb{1}\left(\max_{t\geqslant 0} V_t \geqslant \alpha\right)\right]$ which implies that the seller's profit is not higher than

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

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

10.4 Proofs for alternate implementations

Proof of Proposition 1. 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 $\left(\mathbb{E}[V_T|V_t = v] - p_t^n\right)^+ + p_t^r$. 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 $\Pi(\alpha^*)$. $\square$

Proof of Proposition 2. The proof is similar to Step II in the proof of Theorem 1. This problem again can be thought as a stopping problem where the buyer initially becomes a regular, this is without loss, because there is fixed payment and the buyer can immediately discontinue the service for free.

The buyer with value $V_t = v$ who switches at t to the premium service gets the following payoff:

$$\mathbb{E}[V_T|V_t = v] - p_t^s - \int_t^T dm_s^p = e^{-\lambda(T-t)}(v - \alpha^*)$$

Of course, the buyer can keep the basic service until the final date which yields:

$$\mathbb{E}[V_T|V_t = v] - \int_t^T dm_s^s ds = e^{-\lambda(T-t)}(v - \alpha^*)^+ \geqslant e^{-\lambda(T-t)}(v - \alpha^*)$$

It follows that the buyer always weakly prefers to wait, thus his value at t when $V_t = v$ is $e^{-\lambda(T-t)}(v-\alpha^*)^+$. The smallest optimal stopping time is to stop and switch at the first instance with $V_t \geqslant \alpha^*$.

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 $\Pi(\alpha^*)$. □

10.5 Proofs for extension to linear costs

Proof of Proposition 3. 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 identified in Proposition 3. 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:

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

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, for example, a non-decreasing inverse hazard ratio.

□

Proof of Proposition 4. 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 $\Pi^c(c) \leqslant \Pi^c(\alpha^c)$. □

10.6 Proofs for gains from stochastic mechanisms

Proof of Proposition 5. The proof has two steps: first we define the set of prices which implements the buyer's decision described in Proposition xxx and then we compute the seller's profit.

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

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

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:

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

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

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

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

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

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

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

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:

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

which clearly dominates both gains: from stopping 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

$$p_t^z + \left(1 - e^{-\lambda(T-t)}\right)\int_0^\delta F(v)dv$$

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

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

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

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

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

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

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 $N_T = N$:

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

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 $N \geqslant 2$:

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

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

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

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. □

10.7 Proofs for single arrival

Proof of Proposition 6. 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 expression of $G_t^B(v)$ and $W_t^B(v)$:

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

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

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:

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

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

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

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 $\Pi^s(\alpha)$:

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

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 can not be higher than the maximal surplus net the buyer's expected payoff, that is

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

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 $\mathbb{E}\left[V_T 1\left(V_0 \geqslant \alpha + \beta\right)\right]$, 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 7. 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 ob 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 $\mathbb{E}[V_T|V_t = x_0]$, 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

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

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:

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

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 6 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 6 this allocation can be implemented by dynamic pricing.

□

References

W. J. Adams and J. L. Yellen. Commodity bundling and the burden of monopoly. *The Quarterly Journal of Economics*, 90(3):475–498, 1976.

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.

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.

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

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.

Kautilya. *Arthshastra*. Penguin classics, 1992.

D. Krähmer and R. Strausz. Dynamic mechanism design. In *An Introduction to the Theory of Mechanism Design by Tilman Börgers*, chapter 11, pages 204–234. Oxford University Press, 2015.

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

R. P. McAfee, J. McMillan, and M. D. Whinston. Multiproduct monopoly, commodity bundling, and correlation of values. *The Quarterly Journal of Economics*, 104(2):371–383, 1989.

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.

A. C. Pigou. *The economics of welfare.* London: Macmillan, 1920.

J.-C. Rochet and P. Choné. Ironing, sweeping, and multidimensional screening. *Econometrica*, 66(4):783–826, 1998.

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

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

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

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