

Pricing Novel Goods

Francesco Giovannoni*

Toomas Hinnosaar†

August 11, 2022

Abstract

We study a buyer-seller problem of a novel good for which the seller does not yet know the production cost. A contract can be agreed upon at either the ex-ante stage, before learning the cost, or at the ex-post stage, when both parties will incur a costly delay, but the seller knows the production cost. We show that the optimal ex-ante contract for a profit-maximizing seller is a fixed price contract with an “at-will” clause: the seller can choose to cancel the contract upon discovering her production cost. However, sometimes the seller can do better by offering a guaranteed-delivery price at the ex-ante stage and a second price at the ex-post stage if the buyer rejects the first offer. Such a “limited commitment” mechanism can raise profits, allowing the seller to make the allocation partially dependent on the cost while not requiring it to be embedded in the contract terms. Analogous results hold in a model where the buyer does not know her valuation ex-ante and offers a procurement contract to a seller.

1 Introduction

Consider two scenarios. In the first scenario, a seller faces demand for a good but is uncertain how much it will cost to produce it. For example, a pharmaceutical company (a seller) produces a vaccine for the national health authority (a buyer) to immunize a population against a new disease. The value of the vaccine of a given level of efficacy is (privately) known to the buyer. The seller is unsure of how much it will cost to produce the vaccine, a cost that will be (privately) revealed to the seller only at a later time after enough research is done. In the second scenario, a buyer is considering procuring a good but is uncertain of the benefits she will incur. For example, a defense department

*University of Bristol, francesco.giovannoni@bristol.ac.uk

†University of Nottingham and CEPR, toomas@hinnosaar.net

is considering whether to procure a new weapons system. The seller (privately) knows how much the system will cost, but the defense department is unsure of the system's future usefulness. This uncertainty may resolve over time, but waiting will also delay the delivery of the weapons.

In both scenarios, we have a buyer-seller framework with a principal, who will be ex-post privately informed of her valuation for a novel good, but who does not have that information at the ex-ante stage, whereas the counterpart is immediately (and privately) informed. We ask the question: how should this principal sell (or buy) the good? Should she contract at the ex-ante stage, when she does not yet have her information, or at the ex-post stage, when she will have the information, but at the cost of delay?

To study this problem, we consider a model where a risk-neutral seller faces a risk-neutral buyer with unit demand.¹ The buyer has a private valuation for the object, and the seller has a private cost of production. We depart from the standard setting by allowing for three stages. At time 0 (the *ex-ante* stage) the buyer already knows her valuation, while the seller will only learn her cost at time 1 (the *ex-post* stage). We assume the seller can offer the buyer a contract at the ex-ante stage or at the ex-post stage (or both). We assume that the contract can only be executed one period after it is agreed. So, payoffs are either realized at time 1 or at time 2, and we assume that both buyer and seller discount such payoffs with the same discount factor. In the vaccine example, these assumptions capture the difference between (ex-ante) contracting before the vaccine is fully developed so that development and testing are done together and deployment can be executed immediately after these are completed, and (ex-post) contracting once the vaccine is fully developed, which requires a delay in deployment.

In this context, we consider several mechanisms a profit-maximizing seller could choose. The existing literature (Myerson (1981), Riley and Samuelson (1981), and Yilankaya (1999)) shows that when both buyer and seller contract with knowledge (public or private) of their own valuation or cost, the optimal mechanism for the seller will be a take-it-or-leave-it offer to sell the good at a posted price. This price will depend on the seller's cost but will be the same regardless of whether it is the seller's private information or not. In our setting, these results suggest two natural mechanisms. First, in the *ex-ante fixed price* (EAFP) mechanism, the seller immediately makes a take-it-or-leave-it offer of a price to the buyer, using her *expected* cost to determine such price. Alternatively, in the *ex-post optimal* (EPO) mechanism, the seller does not make an offer immediately but waits for her cost to be realized and makes the take-it-or-leave-it offer using this realization to de-

¹We focus on the case where the principal is the seller, but our results apply equally to the reverse case where the principal is the buyer.

termine her optimal price.² In both cases, the seller commits to a take-it-or-leave-it offer where the good is sold if and only if the offer is accepted. The EAFP mechanism allows trade to occur sooner than the EPO mechanism but forces the seller to offer a price that only reflects her expected cost rather than her actual cost.

We show that *neither* the EAFP or the EPO mechanisms are optimal. In particular, we show that the former is always dominated by an ex-ante mechanism which gives the seller an option of backing out of the transaction altogether if her cost turns out to be too high. The seller fixes a price p , but while in the EAFP contract the seller has to deliver the good if the buyer accepts (a “specific-performance” clause), in this new contract, she can still decide to cancel the order, refunding the price (an “at-will” clause).³ We prove that this contract is the optimal full-commitment contract available to the seller at the ex-ante stage and therefore call this contract the *ex-ante optimal* (EAO) contract.

Still, while the seller can fully commit at the ex-ante stage, she cannot commit to a mechanism that depends on something she doesn’t yet know (i.e., the production cost) unless such mechanism satisfies incentive compatibility constraints on herself as well as the buyer. This opens the possibility that ex-post mechanisms may improve on the seller’s profit. As discussed above, once the seller knows her private information, she optimally offers a contract where the price embeds her valuation *without* the need for this contract to be incentive compatible for her. Indeed, we show that if the discount factor is sufficiently high, then the seller will prefer to use the optimal EPO mechanism over the optimal EAO mechanism.

More surprisingly, we also show that the seller can always do better than the EPO mechanism by using what we call the *dynamic* (D) mechanism. In this mechanism, the seller offers a price at the ex-ante stage, and if the buyer does not accept the first price, a second price is offered at the ex-post stage. If the buyer accepts the offer at the ex-ante stage, the seller commits to deliver the good at that price regardless of the realization of her cost (a specific-performance clause, as in the EAFP mechanism), but if the first price is rejected, the seller makes a take-it-or-leave-it offer at the ex-post stage with a second price (as in the EPO mechanism). A crucial point here is that seller *does not commit* ex-ante to what the price will be at the ex-post stage, nor indeed does she commit to offering a price at all. We show that the optimal D mechanism is *strictly better than* the optimal EPO mechanism: the seller will always choose a period-1 price that the buyer accepts with positive probability.⁴ But then, this leads to our main result: if the discount factor

²This is ex-post optimal because Yilankaya (1999) shows that once the seller knows her valuation ω , it is optimal for her to offer some price $p(\omega)$ which the buyer can either accept or reject.

³We adopt this terminology from Aghion, Dewatripont and Rey (1994).

⁴The EPO mechanism is a special case of the D contract in which the time 0 price is always too high for any buyer to accept, which is equivalent to waiting for time 1 to offer a contract. Thus, the

is sufficiently high, then the best mechanism for the seller amongst those we discussed is the D mechanism; conversely, for lower values of the discount factor, the EAO mechanism is the best one. To understand the trade-off, observe that the D mechanism allows the seller to make two, *sequential*, price offers. At time 0, a first price is offered which is independent of the realized cost because this is unknown at that stage. But at time 0, no commitment is made as to what the seller will do at time 1. At time 1, the cost is realized and if the buyer rejects the time 0 price, the seller will always find it optimal to set a new price that takes the cost realization into account. Thus, if the discount factor is high enough, the dynamic mechanism can be an improvement over the EAO mechanism as it allows the seller to adapt the time 1 price to the cost realization while skimming off the highest buyer types with the time 0 price; contracts where the seller commits at time 0 to a time 1 price as a function of the realized cost cannot have this feature as the seller may be tempted to lie about the realized cost, to get a better price. On the other hand, if the discount factor is sufficiently low, the EAO mechanism is optimal because the seller puts little value on what she can get at the ex-post stage and amongst ex-ante mechanisms it is better to impose an at-will clause.

Next, we discuss how our work relates to the existing literature. Section 2 describes our model and section 3 presents the benchmark mechanisms EAFP and EPO. In section 4 we introduce the EAO and D mechanisms and theorem 1 gives our main result. Section 5 discusses these results and some natural extensions. Most proofs are relegated to appendix A; in appendix B we show, through an example, how our main results continue to hold in a setting where we reverse the roles of buyer and seller. There, the buyer is the principal but does not yet know her valuation at the ex-ante stage. Finally, in appendix C we consider examples of alternative dynamic mechanisms and show that these do not improve over the mechanisms featured in theorem 1.

Related literature. As we already discussed above, the paper builds on the literature on mechanism design, starting with Myerson (1981), Riley and Samuelson (1981). There is also an important literature in dynamic mechanism design. In this literature, dynamic contracts may be optimal if buyers arrive over time (Gershkov and Moldovanu, 2009; Board and Skrzypacz, 2016) or buyers' valuations change or buyers learn about their values over time Baron and Besanko (1984); Courty and Li (2000); Battaglini (2005); Esó and Szentes (2007); Board (2008); Pavan, Segal and Toikka (2014); Garrett (2016); Ely, Garrett and Hinnosaar (2017). However, starting from Stokey (1979); Conlisk, Gerstner and Sobel (1984) this literature has emphasized that if buyers' valuations are persistent

D mechanism is, by construction, no worse for the seller than the EPO mechanism. That it is strictly better is not obvious, as the lack of commitment on the time 1 price triggers Coasian dynamics.

over time, there is no benefit in delaying the trade, i.e., intertemporal price discrimination does not help a monopolist. Our work differs from this literature in two ways: it is the principal who learns about her type over time and she does so only once (in this sense, the model is static). In such situations however, we show that if the discount factor is sufficiently high, dynamic pricing, *without commitment* to future prices, may be optimal.

The fact that the principal privately learns over time connects our setting to informed principal problems, but with some crucial differences. The literature on the informed principal problem (Myerson, 1983; Maskin and Tirole, 1990, 1992; Skreta, 2011; Mylovanov and Tröger, 2012, 2014) assumes that the principal has some valuable information and lacks commitment not to use this information in the contract design stage. For the specific setting of a seller with private cost facing a privately informed buyer, Yilankaya (1999) shows that it does not matter that the seller lacks the commitment not to use her information at the contract design stage: the optimal mechanism is a take-it-or-leave-it price that depends on her cost but is the same whether the cost is private or public information. Put another way, the seller is better off when she knows her cost, but once she does, whether this is common knowledge or not does not matter.⁵ We expand on this literature by looking at the case where the seller can either contract before she knows her cost or after she does, but in the latter case she incurs a cost of delay. Crucially, as the cost of the seller remains private information once known, we assume that the seller cannot write an ex-ante contract where she truthfully reveals her cost ex-post unless doing so is incentive compatible.

There is also a strong connection to the literature on limited commitment and Coasian dynamics. Since Coase (1972) and Bulow (1982) it has been known that a monopolist who lacks commitment power may not be able to exploit his position, as forward-looking buyers would simply wait and get a better price in the future. The literature on mechanism design with limited commitment has taken this as a starting point and studied various ways how the principal can at least partially overcome this problem Bester and Strausz (2001); Vartiainen (2013); Gerardi, Hörner and Maestri (2014); Deb and Said (2015); Doval and Skreta (2019); Fugger, Gretschno and Pollrich (2019); Liu, Mierendorff, Shi and Zhong (2019); Doval and Skreta (2022). In our model, the seller cannot be relied upon to report her future cost truthfully and thus needs to ensure that an ex-ante contract is designed in a way that is incentive-compatible not just for the buyer, but for her future self. By delaying contracting this issue is resolved, but at the cost of delayed execution. One

⁵There are other settings where this lack of commitment does not matter. Tan (1996), for example, shows that, in a procurement auction context, the buyer benefits from knowing her valuation but once this is known, the optimal mechanism is an auction with a reserve price that is the same regardless of whether this is private or public information. It is easy to see that our results would extend to this setting.

way to interpret our main findings is that in such situations, it may be in the principal’s interest to voluntarily choose a limited commitment contract (by not committing to a future price) as a way to overcome her incentive problem.⁶

Perhaps the closest paper to ours in terms of application is Schmitz (2022), which studies a similar bilateral trade problem, where a buyer wants to procure a novel object, but the seller’s costs are only revealed with delay. He shows that an at-will contract, which allows the seller to walk away if the costs turn out to be too high, may generate more social surplus than a specific-performance contract, which requires the seller to deliver an object no matter the costs. The crucial difference is that in our setting it is the principal whose private information is learned with delay, whereas in Schmitz, it is the agent. In our setting, at-will contracts always dominate static specific-performance contracts (in our terminology, EAO and EAFP mechanisms, respectively), because we focus on optimality. More importantly, we show that when the principal has delayed (private) information, a dynamic contract with limited commitment may be preferable: this is never the case when it is the agent that faces information delay.

2 Model

There is a single indivisible object. The buyer has a private valuation θ in the interval $[0, 1]$ for the object. The seller has a private cost ω in the interval $[0, 1]$ for delivering the object, but learning this cost takes time. Both have a common discount factor $\delta \in (0, 1)$. Delivering the object takes time and can only happen one period after the parties agree to trade. Both agents are risk-neutral. We assume that θ has a distribution F with an almost everywhere (a.e.) positive density f while ω has a distribution G with an a.e. positive density g .⁷ We denote with

$$\psi(\theta, \bar{\theta}) := \theta - \frac{F(\bar{\theta}) - F(\theta)}{f(\theta)}$$

the virtual valuation of F when this is truncated to the interval $[0, \bar{\theta}]$ for some $\bar{\theta} \leq 1$. We assume F is regular: $\psi(\theta, 1)$ is a strictly increasing function of ω .⁸ Finally, we denote

⁶There is also a connection to renegotiation proof contracts (Hart and Tirole, 1988; Dewatripont, 1988), which is another form of limited commitment: principals lack the ability to commit to a contract that may be ex-post Pareto dominated. Typically, this literature looks for contracts that won’t be renegotiated, whereas our dynamic mechanism can be interpreted as a sequence of distinct contracts.

⁷Our results can be generalized to a setting where θ and ω are distributed over arbitrary intervals, as long there is a non-empty intersection between such intervals. Using the $[0, 1]$ interval greatly simplifies the exposition.

⁸If $\psi(\theta, 1)$ is strictly increasing in θ , then so is $\psi(\theta, \bar{\theta})$ for any value of $\bar{\theta}$. A standard way to guarantee this is to assume that the hazard rate of F is (weakly) increasing.

with $\mathcal{G}(x) := \int_0^x G(\omega)d\omega$ the left hand integral of the cumulative G .

The seller's goal is to maximize the expected ex-ante profit (from now on, just profit). We use π^M to denote the maximized profit from mechanism (of class) M and $\Pi^M(\cdot)$ as the corresponding profit function of the choice variables of a particular mechanism (typically, but not always, this will be a price).

The timing is as follows. At time 0 (the ex-ante stage), the seller does not know the cost ω and can offer a contract specifying under which circumstances and at which price the object is delivered at time 1 (the ex-post stage). The buyer can accept or reject the contract. If a contract was offered and accepted, we say that the contract was *agreed*. At time 1, the seller learns the cost ω . If a contract was agreed at time 0, its terms are implemented (i.e., the object may be delivered and transfers paid). If a contract was not agreed, the seller can offer a new contract which the buyer can accept or reject. If this latter contract is offered and agreed, its terms are implemented at time 2 and the resulting payoffs are discounted by δ . We call a sequence of such contracts a mechanism.

We are looking for outcomes that are implementable in dominant strategies. When we look at full commitment contracts, we can focus without loss of generality on direct mechanisms, $(q_s(\theta, \omega), t_s(\theta, \omega))$, where $s \in \{0, 1\}$ specifies whether the mechanism was agreed at time 0 or 1, $q_s(\theta, \omega) \in [0, 1]$ specifies the probability that the object is delivered (when the buyer announces value θ and the seller cost ω) and $t_s(\theta, \omega) \in \mathbb{R}$ is the corresponding transfer (from the buyer to the seller if positive).

3 Benchmarks

As discussed in the introduction, Yilankaya (1999) considers the same problem we face, except that at time 0 the seller already (privately) knows the cost of the object. He shows that the optimal mechanism consists of a take-it-or-leave-it offer of a price to the buyer, a price that will, in general, depend on seller's cost. This suggests two natural solutions to the optimality problem in our case. The first entails the seller making a take-it-or-leave-it offer of a price to the buyer at time 0, using the *expected cost* instead of the realized cost. The other is for the seller to wait until time 1, learn the actual cost, and then make a take-it-or-leave-it offer of a price to the buyer, given such cost. We will call the first mechanism the Ex-ante Fixed Price mechanism (EAFP) and the second mechanism the Ex-post Optimal mechanism (EPO). We can apply the Yilankaya (1999) result, which tells us that the optimized EPO mechanism is best amongst the mechanisms available to the seller at time 1 (hence the name), but it incurs a delay and the resulting payoffs will have to be discounted.

The EAFP mechanism commits the seller to a posted price p independently of cost

or buyer's type and delivers the object to the buyer if and only if the buyer accepts the price. Therefore buyer types $\theta \geq p$ accept the offer and the expected cost of delivering the object is $\mathbb{E}_G[\omega]$. This immediately implies:

Proposition 1. *The maximal profit for the seller in the EAFP mechanism is*

$$\pi^{EAFP} = [1 - F(p^{EAFP})] (p^{EAFP} - \mathbb{E}_G[\omega]) > 0,$$

where the optimal price $p^{EAFP} \in (0, 1)$ is the unique solution to $\psi(p, 1) = \mathbb{E}_G[\omega]$.

Given the regularity of F the proof is standard, but note that $p^{EAFP} \in (0, 1)$ guarantees that π^{EAFP} is strictly positive.

In the EPO mechanism, the seller does not offer anything at time 0 and then offers the optimal mechanism at time 1, after learning the cost ω .

Proposition 2. *The maximal profit for the seller in the EPO mechanism is*

$$\pi^{EPO} = \delta \int_{\theta^*}^1 \mathcal{G}(\psi(\theta, 1)) f(\theta) d\theta,$$

where $\theta^* \in (0, 1)$ is the value of θ that solves $\psi(\theta, 1) = 0$.

As discussed, that the maximized EPO mechanism is optimal follows from Yilankaya (1999). In appendix A we provide a proof that allows us to calculate π^{EPO} explicitly.

A comparison of the two mechanisms is then an almost immediate consequence of the two results above:

Corollary 1. *There exists a value $\delta^* \in (0, 1)$ such that $\pi^{EPO} > \pi^{EAFP}$ for all $\delta > \delta^*$ and $\pi^{EPO} < \pi^{EAFP}$ for all $\delta < \delta^*$.*

Proof. In the limiting case where $\delta = 1$, the Yilankaya (1999) result implies π^{EPO} will provide the highest profits obtainable by the seller even when the cost is unknown at time 0: the seller faces no loss of profits by just waiting for the cost to be revealed. In the EPO mechanism, trade occurs if and only if the virtual valuation is no smaller than the realized cost, whereas in the EAFP mechanism (which does not depend on δ), trade occurs if and only if the virtual valuation is no smaller than the expected cost. Therefore π^{EAFP} must be strictly smaller than π^{EPO} for δ close to 1. In sum, π^{EPO} is a linearly increasing function of δ which is equal to zero when $\delta = 0$ and is higher than π^{EAFP} when $\delta = 1$ whereas π^{EAFP} is strictly positive and constant in δ . The result follows immediately. \square

Example. If, for example, we assume that both F and G are uniform distributions, then $\mathcal{G}(x) = \frac{x^2}{2}$, $\psi(\theta, 1) = 2\theta - 1$, $\theta^* = \frac{1}{2}$, and $\mathbb{E}_G[\omega] = \frac{1}{2}$, so that $p^{EAFP} = \frac{3}{4}$ and $\pi^{EAFP} = \frac{1}{16}$. For EPO mechanism, the price $p_1^*(\omega) = \frac{1+\omega}{2}$ and $\pi^{EPO} = \frac{\delta}{12}$, implying that $\delta^* = \frac{3}{4}$.

4 Results

4.1 Ex-ante Optimal Mechanisms

The EAFP and EPO are intuitive mechanisms in our setting where the seller's cost is only revealed after the first transaction opportunity presents itself. Both mechanisms only depend on the buyer's private information on which we impose incentive compatibility constraints, conditioning the mechanisms themselves either on the expected value of the cost (EAFP) or on the realized value of the cost (EPO). Both of these have drawbacks, however, as conditioning on the actual cost improves on conditioning only on the expected cost, but requires a delay.

A different option, which we turn to now, is to consider ex-ante mechanisms, but rather than restricting these mechanisms to depend just on the buyer's private information, we allow them to depend on the future realization of the seller's cost. This creates a significant difference with ex-post mechanisms such as EPO mechanisms: in these latter cases, the seller knows the cost realization when the mechanism is offered and can therefore condition the mechanism itself on such realization, while with ex-ante mechanisms, the seller cannot condition the mechanism on a realization she does not know and so she will have to impose that any such mechanism makes it incentive compatible for her to reveal her cost truthfully. Without loss of generality, we can therefore focus on direct mechanisms $(q_0(\theta, \omega), t_0(\theta, \omega))$, where q_0 is the probability of the product being delivered and t_0 is the transfer from the buyer to the seller. Therefore the problem now is

$$\pi^{EAO} = \max_{q_0, t_0} \int_0^1 \int_0^1 [t_0(\theta, \omega) - \omega q_0(\theta, \omega)] dG(\omega) dF(\theta), \quad (\text{EA})$$

subject to constraints: $q_0(\theta, \omega) \in [0, 1]$,

$$U_0(\theta, \omega) := \theta q_0(\theta, \omega) - t_0(\theta, \omega) \geq 0, \quad \forall \theta, \omega, \quad (\text{IR})$$

$$U_0(\theta, \omega) \geq \theta q_0(\theta', \omega) - t_0(\theta', \omega) \geq 0, \quad \forall \theta, \theta', \omega, \quad (\text{ICB})$$

$$V_0(\theta, \omega) := t_0(\theta, \omega) - \omega q_0(\theta, \omega) \geq t_0(\theta, \omega') - \omega q_0(\theta, \omega'), \quad \forall \theta, \omega, \omega'. \quad (\text{ICS})$$

We can replace (IR) and (ICB) constraints with monotonicity and buyer's envelope condition in the same way as above. Additionally, we can use the same approach to replace the seller's incentive compatibility constraints (ICS) by requiring that q_0 must be non-increasing in ω for all θ , and the corresponding envelope condition for the seller

$$t_0(\theta, \omega) = \omega q_0(\theta, \omega) + \int_{\omega}^1 q_0(\theta, y) dy.$$

It is clear that the problem is linear in q_0 , as multiplying all q_0 's by $\alpha > 0$ multiplies transfers by α and therefore also the profit by α .

The following lemma (proved in appendix A) shows that at least one of the maximizers is an extreme point, and any extreme point has a property that it takes values $\{0, 1\}$ almost everywhere.

Lemma 1. *There exists a maximizer $q_0(\theta, \omega)$ of the maximization problem (EA) such that $q_0(\theta, \omega) \in \{0, 1\}$ almost everywhere.*

Thus, we can focus on deterministic mechanisms. Proposition 3 below characterizes the optimal ex-ante mechanism (EAO) where, again, the proof is relegated to appendix A.

Proposition 3. *There exists a maximizer of the ex-ante optimization problem (EA) such that $q_0(\theta, \omega) = \mathbf{1}[\theta \geq p \geq \omega]$ for some $p \in [0, 1]$. In particular,*

$$\pi^{EAO} = \max_{p \in [0, 1]} \{(1 - F(p))G(p)(p - \mathbb{E}_G[\omega | \omega < p])\} = \max_{p \in [0, 1]} (1 - F(p))\mathcal{G}(p).$$

The EAO mechanism is a fixed price mechanism with price p , where trade occurs if and only if the buyer's valuation is above price p and the seller's valuation is below price p . A natural implementation is the following. The seller offers a price p to the buyer. If the buyer rejects, no trade takes place, but if he accepts, then there is also an at-will clause that allows the seller to cancel the order (returning the price) once the cost is realized at time 1. Of course, the seller will do this if and only if this cost realization turns out to be higher than p . In an EAFP mechanism, instead, the seller operates in a specific performance regime, where she cannot renege on delivering even if the costs turn out to be too high. π^{EAO} represents the maximal profits the seller can obtain by offering a mechanism at time 0, so that the EAFP mechanism cannot do better. In fact, a strict comparison holds:

Corollary 2.

$$\pi^{EAO} > \pi^{EAFP}.$$

Proof. We can write the profits for the seller in an EAFP mechanism with a generic price p as

$$\Pi^{EAFP}(p) = [1 - F(p)](p - \mathbb{E}_G[\omega]) = [1 - F(p)](p - 1 + \mathcal{G}(1)).$$

Consider the difference between the expected profits for the seller in an EAO mechanism

with generic price $p < 1$ and $\Pi^{EAFP}(p)$:

$$\begin{aligned}\Pi^{EAO}(p) - \Pi^{EAFP}(p) &= [1 - F(p)](1 + \mathcal{G}(p) - p - \mathcal{G}(1)) \\ &= [1 - F(p)] \int_p^1 (1 - G(x)) dx > 0.\end{aligned}\tag{1}$$

So, for any price p that maximizes the seller's profits in an EAFP mechanism, the corresponding profit using the same price p in an EAO mechanism is strictly higher. And, of course, a profit from the EAO mechanism with an arbitrary price p is weakly lower than the profit from the optimal EAO mechanism. This establishes the result. \square

The comparison in (1) allows us also to compare any optimal price p^{EAO} in the EAO mechanism with the optimal price p^{EAFP} in the EAFP mechanism:

Corollary 3.

$$p^{EAO} < p^{EAFP}.$$

Proof. Differentiating the expression in (1) with respect to price p shows that the difference between profits from the two mechanisms at a common price p is decreasing:

$$-f(p) \int_p^1 [1 - G(\omega)] d\omega - [1 - F(p)][1 - G(p)] < 0.$$

Therefore at the optimal price p^{EAFP} from the EAFP mechanism, the profit for the EAO mechanism is strictly decreasing, which implies that the optimal price for it is strictly lower than p^{EAFP} . \square

The intuition is instructive. At $p = 0$, the seller gets zero profits (because she can cancel the order) in the at-will mechanism, whereas she has to make a loss when she is not allowed to cancel. At the other extreme, $p = 1$, profits are zero for both mechanisms because there will be no sale. Consider now a $p \in (0, 1)$ and consider a marginal increase in the price. There is no difference between the two mechanisms in terms of the reduced probability of sale, but the distance between them on expected profits conditional on sale decreases because there is a lower probability that canceling the sale will be necessary in the at-will mechanism. Hence, there is a greater incentive to increase the price in the EAFP mechanism than in the EAO mechanism.

We can now compare the optimal EAO mechanism with the optimal EPO mechanism, and we have the following result, which we prove in appendix A:

Proposition 4. *There exists a value $\bar{\delta} \in (0, 1)$, with $\bar{\delta} > \delta^*$, such that $\pi^{EPO} > \pi^{EAO}$ for all $\delta > \bar{\delta}$ and $\pi^{EPO} < \pi^{EAO}$ for all $\delta < \bar{\delta}$.*

This result has an important implication: if δ is sufficiently close to 1, the seller would prefer to wait until the cost is realized. This is a consequence of her inability to embed the cost realization in the ex-ante mechanism and the need to consider a mechanism where she has to apply incentive compatibility constraints to her future self. We will return to this in section 5.

Example. In the example where both F and G are uniform distributions, the optimal EAO mechanism has $p^{EAO} = \frac{2}{3} < p^{EAFP} = \frac{3}{4}$, and $\pi^{EAO} = \frac{2}{27} > \frac{1}{16} = \pi^{EAFP}$. Recall that $\pi^{EPO} = \frac{\delta}{12}$ so that $\bar{\delta} = \frac{8}{9} > \frac{3}{4} = \delta^*$.

4.2 Dynamic Mechanisms

The EAO mechanism defines the highest profit the seller can achieve under full commitment when contracting at time 0. One could consider dynamic mechanisms of the type (q_0, t_0, q_1, t_1) , but it is straightforward to see that such a mechanism cannot increase profits compared to the EAO mechanism as the seller would still be facing the same set of constraints. This observation is consistent with the observation that if agents' private information does not change over time, the optimal mechanism is a static one. Still, proposition 4 suggests that, if the discount factor is sufficiently high, the seller can do better by just waiting for the cost uncertainty to be realized and proposing a mechanism at that point.

We now look at another possibility. We define a Dynamic Mechanism (D) as a mechanism where at time 0 the seller offers a price p_0 to the buyer. If the buyer agrees to purchase the good at this price, then the good is delivered by the seller at time 1, and the game ends. If the buyer does not agree, upon observing her cost realization at time 1, the seller has the opportunity to offer a second price p_1 to the buyer. If this new price is agreed upon, then the good is delivered by the seller at time 2. If the second offer is also rejected, the game ends, and both sides get zero payoffs. The crucial distinction with the possible dynamic contract we discussed above is that here the seller does not commit to the time 1 price at the ex-ante stage. She will offer her time 0 price p_0 , update her beliefs about the buyer's valuation should the price offer be rejected, and then offer p_1 at time 1, when ω is realized.

We consider Perfect Bayesian Equilibria of the game induced by the mechanism where p_0 is such that all buyer types above a threshold $\bar{\theta}(p_0) \in [0, 1]$ will accept the price while all types below the threshold will reject it. If p_0 is rejected, it is optimal for the seller at time 1 to offer a fixed price dependent on her realization of ω , just as in the EPO mechanism, but with beliefs $F(\theta|\theta \leq \bar{\theta}(p_0))$.

Let π^D be the profits obtained by the seller if she optimizes the D mechanism described above. In appendix A, we prove the following result:

Proposition 5.

$$\pi^D > \pi^{EPO}.$$

It is easy to see that the optimal EPO mechanism is a special case of the D mechanism where the first-period price is so high that all buyer types reject it. The importance of the above result lies in the fact that the inequality is strict: it is never optimal for the seller to wait until time 1 to offer a mechanism to the seller. This result is surprising because the seller does not commit to p_1 at time 0 and therefore the usual Coasian dynamics imply that p_0 be set in such a way that types $\theta \geq \bar{\theta}$ are not tempted to wait to purchase at time 1. Putting together corollary 2, proposition 4, proposition 5 and comparing the optimal EAO and D mechanisms, in appendix A we prove our main result:

Theorem 1. *Amongst all the mechanism considered (EAFP, EPO, EAO, D), there exists a value $\delta^{**} < \bar{\delta}$ such that at $\delta = \delta^{**}$ we have $\pi^D = \pi^{EAO}$ whereas*

1. π^D is the highest profit achievable by the seller for all $\delta > \delta^{**}$,
2. π^{EAO} is the highest profit achievable by the seller for all $\delta < \delta^{**}$.

Thus, the EAFP and EPO mechanisms are never optimal. We will discuss some of the implications of this result in the next section.

Example. Going back to our example where both F and G are uniform distributions, we obtain:

$$\bar{\theta}^* = \frac{4 + \delta - \sqrt{(4 - \delta)^2 - 8\delta}}{4\delta}.$$

This leads to time 0 price $p_0^* < \bar{\theta}^*$ because only sufficiently high types will be willing to accept this price and forgo the opportunity to get an even lower price with delay. As expected, we have $p_1^*(\omega) = \frac{\bar{\theta}^* + \omega}{2}$ for $\omega \in [0, \bar{\theta}^*]$. Time 0 prices are depicted in figure 1a, as a function of δ , where the dashed red line shows $\bar{\theta}^*$, the solid red line p_0^* . For ease of comparison, we also show p^{EAO} and p^{EAFP} .

Figure 1a also confirms that the optimal price in the EAO mechanism is lower than the corresponding price in the EAFP mechanism and this, in turn, is lower than $\bar{\theta}^*$. This is the relevant comparison because, as discussed above, in the D mechanism, there will be a gap between the time 0 price and the lowest buyer type willing to accept this price. The fact that $\bar{\theta}^*$ is higher than the other prices, therefore, confirms that in the D mechanism

the seller is only willing to commit to deliver the good when facing very high types. This is a general property which we prove in appendix A:⁹

Proposition 6. *Let $\bar{\theta}^*(\delta)$ be a maximizer of the D mechanism. Then, it is strictly increasing in δ , and, for $\delta > 0$,*

$$\bar{\theta}^*(\delta) > p^{EAFP}.$$

Finally, figure 1a shows that p_0^* , the optimal time 0 price in the D mechanism, is not monotone in δ . The intuition lies in the fact that when δ is relatively low, the seller values a sale at time 0 significantly more than a sale at time 1, so following a marginal increase in δ , she is willing to lower p_0^* to keep the fraction of buyers who accept the time 0 offer (who have an increasing incentive to wait for the time 1 offer, the Coasian effect) relatively high. When δ is relatively high, then the Coasian effect still applies, but the seller no longer has much of an incentive to counteract it because she values contracting at time 1 much more herself and so p_0^* increases.

In figure 1b we show the ex-ante expected profits from the optimized versions of all the mechanisms described here, again as a function of δ . As is easy to see, the figure confirms our results, with $\delta^{**} = \frac{260-8\sqrt{10}}{279} \approx 0.84 < \bar{\delta} = \frac{8}{9} \approx 0.89$.¹⁰

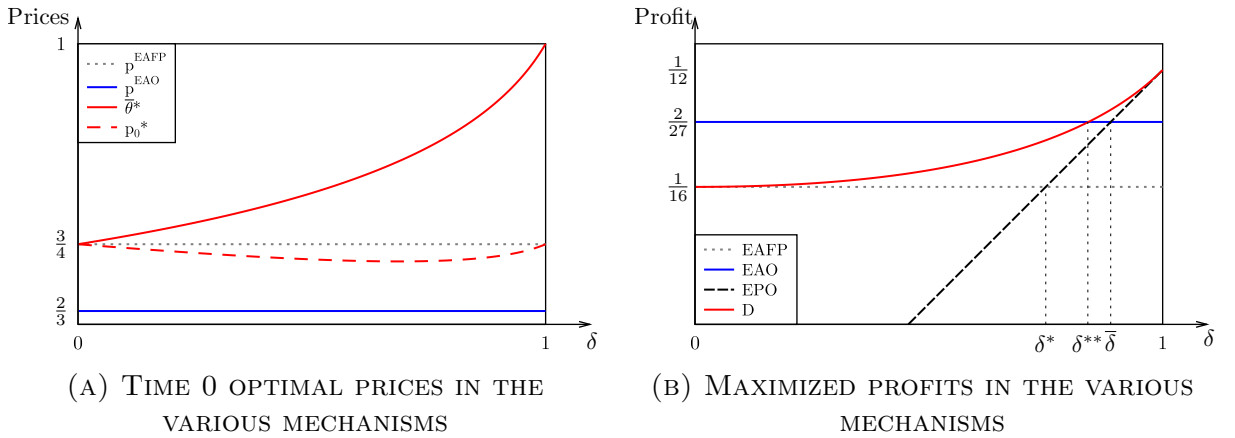


FIGURE 1: EXAMPLE WITH UNIFORM DISTRIBUTIONS

The D mechanism dominates the EPO mechanism because it allows the seller to serve very high types at time 0, types that she would very likely sell to even if she knew her cost. But it still allows the seller the flexibility to serve the buyer once her cost is revealed at time 1, just as the EPO mechanism does. The EAO mechanism allows the seller to have an agreed deal at time 0 and its at-will feature still gives the seller the flexibility to renege if the cost turns out to be too high. Compared to the D mechanism, it serves more

⁹Recall that we already proved that $p^{EAFP} > p^{EAO}$ in corollary 3.

¹⁰In this example, π^D is a strictly increasing function of δ . This is not a general property: as shown in theorem 1, π^D will eventually be strictly increasing, but it may be non-increasing for $\delta < \delta^{**}$.

buyer types at time 0 and so is preferable when δ is small, but for a high discount factor D dominates as the postponement of the contracting with low types allows the seller to relax his own incentive constraints.

5 Discussion

We study a pricing problem of a novel good, where a seller faces demand before learning the cost of production. In this setting, we ask whether a profit-maximizing seller should offer a price to the buyer before (ex-ante) or after (ex-post) she learns her cost. It turns out that the answer is: neither. Indeed, we show that the best ex-ante contract is not just a simple take-it-or-leave-it price; it must also include an at-will clause that allows the seller to renege on the agreement if the cost turns out to be too high. We assume that the seller has full commitment power to a mechanism, which implies the optimal mechanism is *static*: it is never optimal for the seller to commit to a mechanism where trade may be agreed ex-post.

Still, in theorem 1 we consider that a dynamic mechanism, where at the ex-ante stage the seller only commits to deliver the good at some price (i.e. without the at-will clause). She does not commit to anything more, but if that offer is rejected, she offers a price ex-post. We show that this will do better if the discount factor is sufficiently high. This may be puzzling at first glance because it suggests that a mechanism available to a seller with limited commitment power (in the sense of only being able to commit to spot contracts) can do better than any full commitment mechanism. It also suggests that a dynamic mechanism can be optimal in a static setting. The key to understanding this result is that ex-post, the seller is an informed principal that can embed her private information in the mechanism itself: the seller makes a take-or-leave-it offer of a price that depends on her cost realization. Thus, there is no need for incentive compatibility constraints for the seller. Ex-ante, on the other hand, the seller does not know her cost realization and therefore cannot embed it in the mechanism; she can make the mechanism depend on her cost realization, but she has to impose incentive compatibility constraints on her ex-post self because this cost will be her private information. Thus, in our setting, where the principal has private information and is seeking to maximize her payoff, the ability to fully commit ex-ante is different from (and not as good as) the ability to fully commit ex-post. Indeed, if the mechanism designer was a hypothetical external agent who wanted to maximize the seller's profits but did not know the seller's cost ex-ante or ex-post, she would not be able to do better than the optimal ex-ante mechanism because then incentive compatibility on the seller would have to be imposed at both the ex-ante and ex-post stage. The same would apply if we were interested in efficient, instead of seller-

optimal, mechanisms. This explains why for a high enough discount factor, the seller may opt to wait to contract ex-post rather than ex-ante—after all, the optimal ex-post mechanism is, from an ex-ante perspective, a limited commitment contract—the seller makes no offer (or offers something that no buyer type will accept) and waits to learn her cost realization before making a second offer.¹¹

Theorem 1 suggests that the seller should employ two strategies for dealing with the fact that ex-ante, she does not know the cost. One such strategy is offering a price that only high types will accept, leaving the option of serving lower types when the cost is realized. An alternative strategy is to commit to a price ex-ante but to include an at-will clause that allows the seller to renege the offer when the cost is too high. Can the seller do better than this? She clearly cannot do better with a mechanism at the ex-ante stage or one at the ex-post stage because we characterized the optimal ones in each case. But is there a dynamic mechanism that improves upon the mechanism D we described here? Since in these dynamic mechanisms the seller does not commit ex-ante to the ex-post price, we cannot rely on the revelation principle to give us an answer, and a fully general question is an open question in the literature.¹² However, one could conjecture that a dynamic mechanism that combines the features of the two strategies mentioned above could still improve on the Theorem 1 result. In such a mechanism, the seller does not commit to what she will do at the ex-post stage, but at the ex-ante stage, she offers a price with an at-will clause instead of one where delivery is guaranteed. After all, the D mechanism is a combination of EAFP and EPO mechanisms, and given that the EAO mechanism dominates the former, it is natural to ask whether a similar combination of EAO and EPO mechanisms can do better. We do not provide a general answer to this question, but in appendix C we analyze numerically two versions of this mechanism for the case where both F and G are uniform distributions. The analysis shows that such adjusted mechanisms are never optimal in this setting: for low values of δ , they improve on the D mechanism but not on the EAO mechanism, whereas for the remaining values of δ , they improve on the EAO mechanism, but are dominated by the D mechanism. This suggests that the two strategies available to the principal to deal with the issue of delayed information, an at-will clause or a two-price mechanism with no commitment on

¹¹Our results would also hold in a procurement setting where a utility-maximizing buyer who does not know ex-ante her valuation but only learns it ex-post faces a privately informed seller as in the weapon systems example in the introduction. In appendix B we show this for the case of uniform F and G .

¹²Doval and Skreta (2022) suggest, for example, that some rationing may improve the seller’s outcomes. But these results obtain under assumptions that are, on certain dimensions, more tractable (e.g., typically, it is assumed that the buyer has two possible types or that the seller has no costs). More importantly, it is still not obvious what the space of available mechanisms should be. For example, depending on the assumptions made regarding this, the Coase conjecture may (Doval and Skreta, 2019) or may not hold (Brzustowski, Georgiadis and Szentes, 2021; Lomys and Yamashita, 2022).

the ex-post price should be considered as substitutes, not complements.

Another issue that is worth discussing is our assumption that ω is realized at time 1 regardless of whether there was an agreed contract at time 0. Our analysis does not focus on the hold-up issues which arise when investment decisions are made before investors (in our case, the seller) know the possible payoffs, and so our setting abstracts from this. Still, it is worth asking what would happen if the seller could privately learn ex-post her (marginal cost) of production only if she decides to invest in learning it. We can show that if the seller faces multiple buyers, each with unit demand, and if the marginal costs are only learned if a certain number of buyers commit to buying at time 0, then it is still the case that the D mechanism dominates the EAO mechanism, as long as the discount factor is sufficiently high *and* there is a sufficiently large number of buyers. Intuitively, the requirement that a sufficient number of sales be obtained reduces the attractiveness of the D mechanism for a fixed number of buyers because it requires the time 0 price to be quite high so that the seller faces a new risk of not being able to reach such threshold. Still, if the number of potential buyers is large enough, then this threshold will be reached almost surely. Thus, our results are robust to settings where ex-ante investments may be needed to learn the costs.¹³

Finally, our model can be adapted to one where the good's quality is a variable of interest. Suppose that the product has quality $s > 0$ and marginal cost c , potentially both unknown and learned over time. Then the buyer's payoffs would be $\theta sx - p$, where x is the quantity and p the payment. And seller's profit is $p - cx$. Let us define quality-adjusted cost $\omega = c/s$ and quality-adjusted transfers $t = p/s$. With this relabeling, we can apply our original model in terms of payoffs. In this extended model, our results still apply if either s is commonly known (or contractible) or the private information is about quality-adjusted values θ and ω .

¹³The idea that a sufficiently high number of purchases are needed to obtain a realization of ω can be seen as a reduced-form of a model where there needs to be an investment cost for ω to be realized. Cornelli (1996) considers a setting with multiple buyers where the seller faces a fixed cost of production. Cornelli shows that the optimal mechanism for the seller assigns an object to each buyer who has a positive virtual valuation, but only if the sum of the virtual valuations is sufficient to cover the fixed cost. In other words, the optimal mechanism in the presence of a fixed cost implies the same prices as in the optimal mechanism without a fixed cost but requires that enough buyers are willing to purchase at that price regardless of cost. It is easy to see, then, that if such a fixed cost is publicly known or only privately known to the seller (ex-ante or ex-post), it does not matter. The issue of the costs being private information and being realized ex-post instead of ex-ante that we highlight here only matters when these costs affect prices. Hence, for our purposes, whether the investment cost is public or private does not matter, under the natural assumption that it is a fixed cost. So, we can think of a setting, as we do in this discussion, where the fixed cost is set at zero for simplicity, but we do need a sufficiently high number of orders for the seller to learn her cost.

References

- Aghion, P., Dewatripont, M. and Rey, P. (1994) Renegotiation Design with Unverifiable Information, *Econometrica*, **62**, 257–282.
- Baron, D. P. and Besanko, D. (1984) Regulation and Information in a Continuing Relationship, *Information Economics and Policy*, **1**, 267–302.
- Battaglini, M. (2005) Long-Term Contracting with Markovian Consumers, *American Economic Review*, **3**, 637–658.
- Bester, H. and Strausz, R. (2001) Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case, *Econometrica*, **69**, 1077–1098.
- Board, S. (2008) Durable-Goods Monopoly with Varying Demand, *Review of Economic Studies*, **75**, 391–413.
- Board, S. and Skrzypacz, A. (2016) Revenue Management with Forward-Looking Buyers, *Journal of Political Economy*, **124**, 1046–1087.
- Brzustowski, T., Georgiadis, A. and Szentes, B. (2021) Smart Contracts and the Coase Conjecture, *Manuscript*.
- Bulow, J. I. (1982) Durable-Goods Monopolists, *Journal of Political Economy*, **90**, 314–332.
- Coase, R. H. (1972) Durability and Monopoly, *Journal of Law and Economics*, **15**, 143–149.
- Conlisk, J., Gerstner, E. and Sobel, J. (1984) Cyclic Pricing by a Durable Goods Monopolist, *Quarterly Journal of Economics*, **99**, 489–505.
- Cornelli, F. (1996) Optimal Selling Procedures with Fixed Costs, *Journal of Economic Theory*, **71**, 1–30.
- Courty, P. and Li, H. (2000) Sequential Screening, *Review of Economic Studies*, **67**, 697–717.
- Deb, R. and Said, M. (2015) Dynamic Screening with Limited Commitment, *Journal of Economic Theory*, **159**, 891–928.
- Dewatripont, M. (1988) Commitment Through Renegotiation-Proof Contracts with Third Parties, *Review of Economic Studies*, **55**, 377–390.

- Doval, L. and Skreta, V. (2019) Optimal Mechanism for the Sale of a Durable Good, *Manuscript*.
- Doval, L. and Skreta, V. (2022) Mechanism Design With Limited Commitment, *Econometrica*, **90**, 1463–1500.
- Ely, J. C., Garrett, D. and Hinnosaar, T. (2017) Overbooking, *Journal of the European Economic Association*, **15**, 1258–1301.
- Esó, P. and Szentes, B. (2007) Optimal Information Disclosure in Auctions and the Handicap Auction, *Review of Economic Studies*, **74**, 705–731.
- Fugger, N., Gretschko, V. and Pollrich, M. (2019) Sequential Procurement with Limited Commitment, *Manuscript*.
- Garrett, D. F. (2016) Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values, *American Economic Review*, **106**, 3275–3299.
- Gerardi, D., Hörner, J. and Maestri, L. (2014) The Role of Commitment in Bilateral Trade, *Journal of Economic Theory*, **154**, 578–603.
- Gershkov, A. and Moldovanu, B. (2009) Dynamic Revenue Maximization with Heterogeneous Objects: A Mechanism Design Approach, *American Economic Journal: Microeconomics*, **1**, 168–198.
- Hart, O. D. and Tirole, J. (1988) Contract Renegotiation and Coasian Dynamics, *Review of Economic Studies*, **55**, 509–540.
- Liu, Q., Mierendorff, K., Shi, X. and Zhong, W. (2019) Auctions with Limited Commitment, *American Economic Review*, **109**, 876–910.
- Lomys, N. and Yamashita, T. (2022) A Mediator Approach to Mechanism Design with Limited Commitment, *Manuscript*.
- Maskin, E. and Tirole, J. (1990) The Principal-Agent Relationship with an Informed Principal: The Case of Private Values, *Econometrica*, **58**, 379–409.
- Maskin, E. and Tirole, J. (1992) The Principal-Agent Relationship with an Informed Principal, II: Common Values, *Econometrica*, **60**, 1–42.
- Myerson, R. (1981) Optimal Auction Design, *Mathematics of Operations Research*, **6**, 58–73.

- Myerson, R. B. (1983) Mechanism Design by an Informed Principal, *Econometrica*, **51**, 1767–1797.
- Mylovanov, T. and Tröger, T. (2012) Informed-Principal Problems in Environments With Generalized Private Values, *Theoretical Economics*, **7**, 465–488.
- Mylovanov, T. and Tröger, T. (2014) Mechanism Design by an Informed Principal: Private Values with Transferable Utility, *Review of Economic Studies*, **81**, 1668–1707.
- Pavan, A., Segal, I. and Toikka, J. (2014) Dynamic Mechanism Design: A Myersonian Approach, *Econometrica*, **82**, 601–653.
- Riley, J. G. and Samuelson, W. F. (1981) Optimal Auctions, *American Economic Review*, **71**, 381–392.
- Schmitz, P. W. (2022) How (Not) To Purchase Novel Goods and Services: Specific Performance Versus At-Will Contracts, *Economic Journal*, p. ueac024.
- Skreta, V. (2011) On the Informed Seller Problem: Optimal Information Disclosure, *Review of Economic Design*, **15**, 1–36.
- Stokey, N. L. (1979) Intertemporal Price Discrimination, *Quarterly Journal of Economics*, **93**, 355–371.
- Tan, G. (1996) Optimal Procurement Mechanisms for an Informed Buyer, *Canadian Journal of Economics*, **29**, 699–716.
- Vartiainen, H. (2013) Auction Design Without Commitment, *Journal of the European Economic Association*, **11**, 316–342.
- Yilankaya, O. (1999) A Note on the Seller’s Optimal Mechanism in Bilateral Trade with Two-Sided Incomplete Information, *Journal of Economic Theory*, **87**, 267–271.

A Proofs

A.1 Proof of Proposition 2

Proof. One way to interpret the Yilankaya (1999) result is that when the seller privately knows ω , she can do no better than when ω is publicly known. So, it is without loss of generality to restrict the set of direct mechanisms $(q_1(\theta, \omega), t_1(\theta, \omega))$ to those of the form $(q_1(\theta|\omega), t_1(\theta|\omega))$, where the mechanism embeds the realized cost ω and where there is

no need for incentive compatibility constraints on the seller. The rest of the argument is standard, but it allows us to calculate the seller's profits.

The seller's problem is

$$\pi^{EPO} = \delta \max_{q_1, t_1} \int_0^1 \int_0^1 (t_1(\theta|\omega) - \omega q_1(\theta|\omega)) dG(\omega) dF(\theta),$$

subject to constraints: $q_1(\theta|\omega) \in [0, 1]$,

$$U_1(\theta|\omega) := \theta q_1(\theta|\omega) - t_1(\theta|\omega) \geq 0, \quad \forall \theta, \omega, \quad (\text{IR})$$

$$U_1(\theta|\omega) \geq \theta q_1(\theta'|\omega) - t_1(\theta'|\omega) \geq 0, \quad \forall \theta, \theta', \omega. \quad (\text{ICB})$$

Using the standard steps, we can replace buyer's incentive compatibility constraint (ICB) and individual rationality constraint (IR) with (1) $U(0|\omega) = 0$, (2) monotonicity of $q_1(\theta|\omega)$ (i.e., q_1 is a weakly increasing function of θ for all ω), and (3) the envelope condition

$$t_1(\theta|\omega) = \theta q_1(\theta|\omega) - \int_0^\theta q_1(x|\omega) dx.$$

Using the envelope condition to substitute t_1 out from the optimization problem, we can rewrite the problem as

$$\pi^{EPO} = \delta \max_{q_1} \int_0^1 \int_0^1 (\psi(\theta, 1) - \omega) q_1(\theta|\omega) dG(\omega) dF(\theta).$$

The optimum is $q_1(\theta|\omega) = \mathbf{1}[\psi(\theta, 1) \geq \omega]$. This implies that buyers with low types such that $\psi(\theta, 1) < 0$ never get the product. This condition is equivalent to $\theta < \theta^*$, where θ^* solves $\psi(\theta, 1) = 0$. For the remaining types $\theta \geq \theta^*$, we have $q_1(\theta|\omega) = \mathbf{1}[\theta \geq p^*(\omega)]$ where $p^*(\omega)$ solves $\psi(p, 1) = \omega$. We can rewrite the maximum payoff for the seller as follows:

$$\begin{aligned} \pi^{EPO} &= \delta \int_{\theta^*}^1 \int_0^{\psi(\theta, 1)} (\psi(\theta, 1) - \omega) dG(\omega) dF(\theta) \\ &= \delta \int_{\theta^*}^1 \psi(\theta, 1) \int_0^{\psi(\theta, 1)} dG(\omega) dF(\theta) - \delta \int_{\theta^*}^1 \int_0^{\psi(\theta, 1)} \omega dG(\omega) dF(\theta) \\ &= \delta \int_{\theta^*}^1 (\psi(\theta, 1) - \mathbb{E}_G[\omega|\omega < \psi(\theta, 1)]) G(\psi(\theta, 1)) dF(\theta) \\ &= \delta \int_{\theta^*}^1 \left(\psi(\theta, 1) - \left(\psi(\theta, 1) - \frac{\mathcal{G}(\psi(\theta, 1))}{G(\psi(\theta, 1))} \right) \right) G(\psi(\theta, 1)) f(\theta) d\theta \\ &= \delta \int_{\theta^*}^1 \mathcal{G}(\psi(\theta, 1)) dF(\theta) \end{aligned}$$

as specified in the proposition. □

A.2 Proof of Lemma 1

Proof. Let \mathcal{H} be the set of all bounded functions $h : [0, 1]^2 \rightarrow [0, 1]$. We endow \mathcal{H} with the standard linear structure so that it is a vector space. Moreover, we endow \mathcal{H} with the norm $|h| = \int_0^1 \int_0^1 |h(x, y)| dx dy$. Let $\mathcal{M} \subset \mathcal{H}$ the set of functions such that $h \in \mathcal{M}$ satisfies (1) $h(x, y) \in [0, 1]$ for all x, y , (2) h is weakly increasing in x for all y , and (3) h is weakly decreasing in y for all x .

We claim that \mathcal{M} is convex and compact. To prove convexity, let us take any $h, k \in \mathcal{M}$ and $\lambda \in [0, 1]$. Then $\alpha h + (1 - \alpha)k$ is clearly also monotone (in the relevant directions) and in $[0, 1]$. Therefore in \mathcal{M} . The compactness of \mathcal{M} follows from a straightforward two-dimensional generalization of Helly's selection theorem.¹⁴

For notational convenience, we drop the subscript from q_0 . We apply the Extreme Point theorem: \mathcal{M} is a compact, convex subset of a normed vector space \mathcal{H} , and the objective is a continuous linear function of q . Therefore there exist some extreme points, and in particular, there exists an extreme point q^* that is a maximizer of the objective.

We next claim that q is an extreme point of \mathcal{M} if and only if $q(\theta, \omega) \in \{0, 1\}$ for almost all θ, ω . For necessity, take $q \in \mathcal{M}$ such that $q(\theta, \omega) \in \{0, 1\}$ for almost all values. Take any $\hat{q} \in \mathcal{M}$ such that $\hat{q} \neq 0$, i.e., there exists a positive measure of (θ, ω) such $\hat{q}(\theta, \omega) \neq 0$. Then for all such values, (θ, ω) , one of four cases holds:

1. $q(\theta, \omega) = 0, \hat{q}(\theta, \omega) > 0$, then $q(\theta, \omega) - \hat{q}(\theta, \omega) < 0$, so $q - \hat{q} \notin \mathcal{M}$.
2. $q(\theta, \omega) = 0, \hat{q}(\theta, \omega) < 0$, then $q(\theta, \omega) + \hat{q}(\theta, \omega) < 0$, so $q + \hat{q} \notin \mathcal{M}$.
3. $q(\theta, \omega) = 1, \hat{q}(\theta, \omega) > 0$, then $q(\theta, \omega) + \hat{q}(\theta, \omega) > 1$, so $q + \hat{q} \notin \mathcal{M}$.
4. $q(\theta, \omega) = 1, \hat{q}(\theta, \omega) < 0$, then $q(\theta, \omega) - \hat{q}(\theta, \omega) > 1$, so $q - \hat{q} \notin \mathcal{M}$.

Therefore q is an extreme point of \mathcal{M} , which proves the necessity. For sufficiency, take an extreme point $q \in \mathcal{M}$ and suppose by contradiction that there exist a positive mass of pairs (θ, ω) such that $q(\theta, \omega) \in (0, 1)$. Define a new function

$$\hat{q}(\theta, \omega) = \begin{cases} q(\theta, \omega) & \text{if } q(\theta, \omega) \leq 0.5, \\ 1 - q(\theta, \omega) & \text{if } q(\theta, \omega) \geq 0.5. \end{cases}$$

¹⁴The proof is available upon request.

This function \hat{q} is a bounded function from $[0, 1]^2$ to \mathbb{R} (it is not monotone, so not in \mathcal{M} , but this does not matter). Also, $\hat{q} \neq 0$. Now,

$$q(\theta, \omega) + \hat{q}(\theta, \omega) = \begin{cases} 2q(\theta, \omega) & \text{if } q(\theta, \omega) \leq 0.5, \\ 1 & \text{if } q(\theta, \omega) \geq 0.5, \end{cases}$$

so, $q + \hat{q}$ is a function that only takes values in $[0, 1]$, and it is weakly increasing in θ for any given ω and weakly decreasing in ω for any given θ .¹⁵ Therefore $q + \hat{q} \in \mathcal{M}$. Similarly,

$$q(\theta, \omega) - \hat{q}(\theta, \omega) = \begin{cases} 0 & \text{if } q(\theta, \omega) \leq 0.5, \\ 2q(\theta, \omega) - 1 & \text{if } q(\theta, \omega) \geq 0.5, \end{cases}$$

which is in $[0, 1]$ and weakly increasing in θ and weakly decreasing in ω , so $q - \hat{q} \in \mathcal{M}$ as well. Therefore q is not an extreme point of \mathcal{M} , which is a contradiction. \square

A.3 Proof of Proposition 3

Proof. We continue to drop the subscript from q_0 . Lemma 1 implies that we only need to consider $q(\theta, \omega) \in \{0, 1\}$ a.e. It is useful to combine (ICB) and (ICS) constraints into one as follows:

$$(\theta - \omega)q(\theta, \omega) = \int_0^\theta q(x, \omega)dx + \int_\omega^1 q(\theta, y)dy.$$

First, note that the right-hand side of the expression is always non-negative, so $q(\theta, \omega) = 0$ for all $\theta < \omega$. Therefore we can rewrite the expression as

$$(\theta - \omega)q(\theta, \omega) = \int_\omega^\theta q(x, \omega)dx + \int_\omega^\theta q(\theta, y)dy.$$

Observe that monotonicity plus the previous lemma jointly imply that there exist weakly increasing function $p_1(\omega) \in [\omega, 1]$ and a weakly decreasing function $p_2(\theta) \in [0, \theta]$

¹⁵Fix $\theta' > \theta, \omega' > \omega$, then

$$(q + \hat{q})(\theta', \omega) - (q + \hat{q})(\theta, \omega) = \begin{cases} 2q(\theta', \omega) - 2q(\theta, \omega) \geq 0 & \text{if } q(\theta, \omega) \leq q(\theta', \omega) \leq 0.5, \\ 1 - 2q(\theta, \omega) \geq 0 & \text{if } q(\theta, \omega) \leq 0.5 \leq q(\theta', \omega), \\ 1 - 1 = 0 & \text{if } 0.5 \leq q(\theta, \omega) \leq q(\theta', \omega), \end{cases}$$

$$(q + \hat{q})(\theta, \omega') - (q + \hat{q})(\theta, \omega) = \begin{cases} 2q(\theta, \omega') - 2q(\theta, \omega) \leq 0 & \text{if } q(\theta, \omega') \leq q(\theta, \omega) \leq 0.5, \\ 2q(\theta, \omega') - 1 \leq 0 & \text{if } q(\theta, \omega') \leq 0.5 \leq q(\theta, \omega), \\ 1 - 1 = 0 & 0.5 \leq \text{if } q(\theta, \omega') \leq q(\theta, \omega). \end{cases}$$

such that (except possibly for a zero-measure set):

$$q(\theta, \omega) = \begin{cases} 0 & \text{if } \theta < p_1(\omega) \\ \in [0, 1] & \text{if } \theta = p_1(\omega) \\ 1 & \text{if } \theta > p_1(\omega) \end{cases} = \begin{cases} 0 & \text{if } \omega > p_2(\theta), \\ \in [0, 1] & \text{if } \omega = p_2(\theta), \\ 1 & \text{if } \omega < p_2(\theta). \end{cases}$$

Fix $\omega^* \in (0, 1)$ such that $p^* := p_1(\omega^*) < 1$ (such value exists, otherwise profit is 0). Take any $\theta > p^*$. Then monotonicity implies

$$\theta - \omega^* = \int_{\omega^*}^{p_1(\omega^*)} 0 dx + \int_{p_1(\omega^*)}^{\theta} 1 dx + \int_{\omega^*}^{p_2(\theta)} 1 dy + \int_{p_2(\theta)}^{\theta} 0 dy = \theta - p^* + p_2(\theta) - \omega^*.$$

Therefore $p_2(\theta) = p^*$. Now, take (θ, ω) such that $\theta > p^* > \omega$. Then monotonicity implies that $q(\theta, \omega) = 1$ and therefore

$$\theta - \omega = \theta - p_1(\omega) + p^* - \omega.$$

So that $p_1(\omega) = p^*$. This proves the claim. \square

A.4 Proof of Proposition 4

Proof. We showed that π^{EAO} represents the maximized profits from an at-will posted price mechanism. Let $\Pi^{EAO}(x)$ denote profit from the same class of mechanisms but with generic variable x , so that $\pi^{EAO} = \max_x \Pi^{EAO}(x)$. We can write

$$\begin{aligned} \Pi^{EAO}(x) &= (x - \mathbb{E}_G[\omega | \omega < x])G(x)(1 - F(x)) = \left(x - \frac{xG(x) - \int_0^x G(y)dy}{G(x)} \right) G(x)(1 - F(x)) \\ &= (1 - F(x))\mathcal{G}(x) = - \int_x^1 [-f(\theta)\mathcal{G}(\theta) + (1 - F(\theta))G(\theta)] d\theta \\ &= \int_x^1 \left[\mathcal{G}(\theta) - \frac{1 - F(\theta)}{f(\theta)}G(\theta) \right] f(\theta)d\theta = \int_x^1 [\mathcal{G}(\theta) - (\theta - \psi(\theta, 1))G(\theta)] f(\theta)d\theta. \end{aligned}$$

Now, let $\Pi^{FP}(x)$ represent the continuation (i.e. not discounted) profits for the seller in the fixed price mechanism where she knows ω for a generic x (so that $\pi^{EPO} = \delta\Pi^{FP}(\theta^*)$). As shown in the proof of proposition 2, we can express

$$\Pi^{FP}(x) = \int_x^1 \mathcal{G}(\psi(\theta, 1))f(\theta)d\theta.$$

and can now write

$$\begin{aligned}\Pi^{FP}(x) - \Pi^{EAO}(x) &= \int_x^1 [\mathcal{G}(\psi(\theta, 1)) - \mathcal{G}(\theta) + [\theta - \psi(\theta, 1)]G(\theta)] f(\theta)d\theta \\ &= \int_x^1 \int_{\psi(\theta, 1)}^\theta [G(\theta) - G(y)] dy f(\theta)d\theta > 0,\end{aligned}$$

because $\psi(\theta, 1) < \theta$ for all $\theta < 1$, and for all $y < \theta$, we have $G(\theta) > G(y)$ (as g is a.e. positive, G is strictly increasing in $[0, 1]$).

Therefore, we get

$$\Pi^{FP}(\theta^*) \geq \Pi^{FP}(p^{EAO}) > \Pi^{EAO}(p^{EAO}) = \pi^{EAO},$$

where p^{EAO} is a maximizer of $\Pi^{EAO}(x)$. Note that both $\Pi^{FP}(\theta^*)$ and π^{EAO} are independent of δ . Define $\bar{\delta} = \frac{\pi^{EAO}}{\Pi^{FP}(\theta^*)}$, so that $\bar{\delta} \in (0, 1)$. Now

$$\pi^{EPO} = \delta \Pi^{FP}(\theta^*) > \pi^{EAO} \iff \delta > \bar{\delta}.$$

Finally, since $\pi^{EAO} > \pi^{EAFP}$ it must be that $\bar{\delta} > \delta^*$. □

A.5 Proof of Proposition 5

Proof. In this mechanism, if the buyer did not buy at time 0, the seller at time 1 faces a buyer $\theta \in [0, \bar{\theta}]$ with a cumulative distribution function $F(\theta|\theta \leq \bar{\theta}) = \frac{F(\theta)}{F(\bar{\theta})}$. Following the usual steps we can write the seller's expected profits at time 1, as a function of the probability of trade, and given the buyer's type is in $[0, \bar{\theta}]$ as

$$\Pi^{D^1}(\bar{\theta}) = \int_0^{\bar{\theta}} \int_0^{\bar{\theta}} \left[q_1(\theta, \omega) - \int_0^\theta q_1(y, \omega) dy - \omega q_1(\theta|\omega) \right] \frac{f(\theta)}{F(\bar{\theta})} d\theta dG(\omega),$$

so that the expected profits for the seller at time 0 in the D mechanism, given $\bar{\theta}$, are

$$\Pi^D(\bar{\theta}, \delta) = (1 - F(\bar{\theta})) \int_0^1 (p_0 - \omega) dG(\omega) + \delta F(\bar{\theta}) \Pi^{D^1}(\bar{\theta}),$$

where p_0 solves, for a fixed $\bar{\theta}$, the equation

$$\bar{\theta} - p_0 = \delta \int_0^{\bar{\theta}} \int_0^{\bar{\theta}} q_1(y, \omega) dy dG(\omega), \tag{CO}$$

in which the right-hand side represents the expected utility for the type $\bar{\theta}$ buyer at time 1 if she rejects price p_0 . Equation (CO) captures the Coasian constraint that types $\theta \in [\bar{\theta}, 1]$

prefer to accept price p_0 to the expected utility they could get if they waited for the mechanism offered at time 1. Now, since at time 1, the seller learns ω , we can apply the Yilankaya (1999) result, which tells us that

$$q_1(\theta, \omega) = \mathbf{1} \left[\omega < \theta - \frac{F(\bar{\theta}) - F(\theta)}{f(\theta)} \right] = \mathbf{1} [\omega < \psi(\theta, \bar{\theta})],$$

so that

$$p_0 = \bar{\theta} - \delta \int_0^{\bar{\theta}} \int_0^{\max(0, \psi(\theta, \bar{\theta}))} dG(\omega) dy = \bar{\theta} - \delta \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} G(\psi(\theta, \bar{\theta})) d\theta,$$

where $\theta^{**}(\bar{\theta})$ is the value of θ for which $\psi(\theta, \bar{\theta}) = 0$. If $\psi(\theta, 1)$ is strictly increasing in θ then so is $\psi(\theta, \bar{\theta})$, which implies that $\theta^{**}(\bar{\theta})$ is well-defined, and $\theta^{**}(1) = \theta^*$. Also,

$$\begin{aligned} \Pi^{D^1}(\bar{\theta})F(\bar{\theta}) &= \int_0^{\bar{\theta}} \int_0^{\bar{\theta}} \left[q_1(\theta, \omega) - \int_0^{\theta} q_1(y, \omega) dy - \omega q_1(\theta | \omega) \right] f(\theta) d\theta dG(\omega) \\ &= \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} (\theta f(\theta) - F(\bar{\theta}) + F(\theta)) G(\psi(\theta, \bar{\theta})) d\theta \\ &\quad - \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} f(\theta) [\psi(\theta, \bar{\theta}) G(\psi(\theta, \bar{\theta})) - \mathcal{G}(\psi(\theta, \bar{\theta}))] d\theta \\ &= \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} \mathcal{G}(\psi(\theta, \bar{\theta})) f(\theta) d\theta, \end{aligned}$$

Hence,

$$\Pi^D(\bar{\theta}, \delta) = (1 - F(\bar{\theta})) \left[\bar{\theta} - \delta \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} G(\psi(\theta, \bar{\theta})) d\theta - \mathbb{E}_G[\omega] \right] + \delta \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} \mathcal{G}(\psi(\theta, \bar{\theta})) f(\theta) d\theta,$$

Recalling that $\int_0^1 \omega dG(\omega) = \mathbb{E}_G[\omega] = 1 - \mathcal{G}(1)$, we now consider

$$\begin{aligned} \frac{\partial \Pi^D(\bar{\theta}, \delta)}{\partial \bar{\theta}} &= -f(\bar{\theta}) \left[\bar{\theta} - \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} G(\psi(\theta, \bar{\theta})) d\theta - 1 + \mathcal{G}(1) \right] \\ &\quad - \delta(1 - F(\bar{\theta})) \left[\frac{d \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} G(\psi(\theta, \bar{\theta}))}{d\bar{\theta}} \right] \\ &\quad + \delta f(\bar{\theta}) \left[\mathcal{G}(\psi(\bar{\theta}, \bar{\theta})) - \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} G(\psi(\theta, \bar{\theta})) \right] d\theta, \end{aligned}$$

and if we evaluate this at $\bar{\theta} = 1$ we get

$$\begin{aligned}\frac{\partial \Pi^D(1, \delta)}{\partial \bar{\theta}} &= -f(1) \left[1 - \int_{\theta^*}^1 G(\psi(\theta, 1)) d\theta - 1 + \mathcal{G}(1) \right] + \delta f(1) \left[\mathcal{G}(1) - \int_{\theta^*}^1 G(\psi(\theta, 1)) d\theta \right] \\ &= f(1) \left[\delta \int_{\theta^*}^1 G(\psi(\theta, 1)) d\theta - \mathcal{G}(1) + \delta \mathcal{G}(1) - \delta \int_{\theta^*}^1 G(\psi(\theta, 1)) d\theta \right] \\ &= f(1) \mathcal{G}(1) (\delta - 1),\end{aligned}$$

because, clearly, $\frac{d \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} G(\psi(\theta, \bar{\theta}))}{d\bar{\theta}} < \infty$ for any $\bar{\theta}$. So, for any $\delta < 1$, $\Pi^D(\bar{\theta}, \delta)$ is decreasing at $\bar{\theta} = 1$ and this implies that it is optimized for $\bar{\theta} < 1$. Since $\Pi^D(1, \delta) = \pi^{EPO}$ then this proves the result. \square

A.6 Proof of Theorem 1

Proof. We begin by comparing the optimal D with the optimal EAO mechanism. From the proof of proposition 5 it is easy to see that π^D is a continuous function of δ . There, we also show that for $\delta < 1$ we have $\pi^D > \pi^{EPO}$ while if $\delta = 1$, then $\pi^D = \pi^{EPO}$. These imply that $\pi^D > \pi^{EAO}$ for a non-empty set of values of δ . By continuity, there exists a δ^{**} which is the smallest value of δ such that $\pi^D(\delta) = \pi^{EAO}$. Now, let

$$\begin{aligned}\Pi^D(x, \delta) &= (1 - F(x)) \left[x - \delta \int_{\theta^{**}(x)}^x G(\psi(\theta, x)) d\theta - \mathbb{E}_G[\omega] \right] + \delta \int_{\theta^{**}(x)}^x \mathcal{G}(\psi(\theta, x)) f(\theta) d\theta, \\ &= \underbrace{(1 - F(x)) [x - \mathbb{E}_G[\omega]]}_{=\Pi^{EAFP}(x)} + \delta \underbrace{\left[\int_{\theta^{**}(x)}^x \mathcal{G}(\psi(\theta, x)) f(\theta) d\theta - \int_{\theta^{**}(x)}^x G(\psi(\theta, x)) d\theta \right]}_{=:\Phi(x)} \\ &= \Pi^{EAFP}(x) + \delta \Phi(x),\end{aligned}$$

and as we can define $\pi^D(\delta) = \Pi^D(x^*(\delta), \delta) = \Pi^{EAFP}(x^*(\delta)) + \delta \Phi(x^*(\delta))$, where $x^*(\delta)$ is a maximizer of $\Pi^D(x, \delta)$, then by the envelope theorem,

$$\frac{d\pi^D(\delta)}{d\delta} = \Phi(x^*(\delta)).$$

Defining $\pi^{EAFP} = \Pi^{EAFP}(x^o)$, where x^o is the maximizer of $\Pi^{EAFP}(x)$, we therefore have that

$$\pi^{EAO} = \pi^D(\delta^{**}) = \Pi^{EAFP}(x^*(\delta^{**})) + \delta^{**} \Phi(x^*(\delta^{**})) > \Pi^{EAFP}(x^o) = \pi^{EAFP}.$$

By definition, $\Pi^{EAFP}(x^o) \geq \Pi^{EAFP}(x^*(\delta^{**}))$ which implies that $\Phi(x^*(\delta^{**})) > 0$. Now, suppose $\delta > \delta^{**}$.

$$\pi^D(\delta) = \Pi^{EAFP}(x^*(\delta)) + \delta\Phi(x^*(\delta)) \geq \Pi^{EAFP}(x^*(\delta^{**})) + \delta\Phi(x^*(\delta^{**})) > \pi^D(\delta^{**}) > \pi^{EAFP},$$

which means that $\Phi(x^*(\delta)) > 0$. Hence, once $\pi^D(\delta)$ reaches π^{EAO} at δ^{**} , then it must be strictly increasing and $\pi^D(\delta) > \pi^{EAO}$ for all $\delta > \delta^{**}$. Thus, δ^{**} is unique.

Conversely, $\pi^D(0) = \pi^{EAFP} < \pi^{EAO}$ which means that there is a non-empty set of values of δ for which $\pi^D(\delta) < \pi^{EAO}$. By the construction above, this set must be the set $[0, \delta^{**})$. That $\delta^{**} < \bar{\delta}$ follows from proposition 5. Finally, corollary 2 implies that the optimal EAFP is always dominated by the optimal EAO mechanism, whereas proposition 5 implies that the optimal D mechanism always dominates the optimal EPO mechanism. \square

A.7 Proof of Proposition 6

Proof. From the proof of theorem 1 consider the function $\Phi(x)$. Define

$$\phi(x) := \frac{d\Phi(x)}{dx} = f(x) \left[\mathcal{G}(x) - (x - \psi(x, 1))G(x) + \int_{\theta^{**}(x)}^x (\psi(\theta, x) - \psi(\theta, 1))g(\psi(\theta, x))d\theta \right].$$

Note that $\phi(x)$ is a continuous function. Remember that x^o , the maximizer of $\Pi^{EAFP}(x)$, satisfies $\psi(x^o, 1) = \mathbb{E}_G[\omega]$, so

$$\begin{aligned} \frac{\phi(x^o)}{f(x^o)} &= \mathcal{G}(x^o) - (x^o - \mathbb{E}_G[\omega])G(x^o) + \int_{\theta^{**}(x^o)}^{x^o} (\psi(\theta, x^o) - \psi(\theta, 1))g(\psi(\theta, x^o))d\theta \\ &= G(x^o) \underbrace{(\mathbb{E}_G[\omega] - \mathbb{E}_G(\omega|\omega < x^o))}_{>0} + \int_{\theta^{**}(x^o)}^{x^o} \underbrace{(\psi(\theta, x^o) - \psi(\theta, 1))}_{>0} g(\psi(\theta, x^o))d\theta > 0. \end{aligned}$$

We know by proposition 5 that $x^*(\delta)$ must be an interior solution, and from proposition 1 that x^o must also be an interior solution. Thus, by the argument above, for any $x^*(\delta)$ that maximizes $\Pi^D(x, \delta)$ and for x^o we must have

$$\begin{aligned} \frac{\partial \Pi^D(x^*(\delta), \delta)}{\partial x} &= \frac{d\Pi^{EAFP}(x^*(\delta))}{dx} + \delta\phi(x^*(\delta)) = 0, \\ \frac{\partial \Pi^D(x^o, \delta)}{\partial x} &= \frac{d\Pi^{EAFP}(x^o)}{dx} + \delta\phi(x^o) = \delta\phi(x^o) > 0, \end{aligned}$$

while simple inspection shows that $\frac{d\Pi^{EAFP}(x)}{dx}$ is positive for $x < x^o$ and negative for $x > x^o$ so that x^o is unique. Also, $\phi(x^o) > 0$ implies $x^*(\delta) \neq x^o$. So, suppose, contrary to the

statement, that for some fixed δ , there is at least one maximizer $x^*(\delta)$ that is smaller than x^o . Then, given the conditions above, it must be that $\frac{d\Pi^{EAFP}(x^*(\delta))}{dx} > 0$ which implies $\phi(x^*(\delta)) < 0$. Consider now the highest such $x^*(\delta) < x^o$ and any $x \in (x^*(\delta), x^o)$. For such x it must be that

$$\begin{aligned} \frac{d\Pi^D(x, \delta)}{dx} &= \frac{d\Pi^{EAFP}(x)}{dx} + \delta\phi(x) < 0, \\ \frac{d\Pi^{EAFP}(x)}{dx} &> 0. \end{aligned}$$

The first inequality applies because continuity of $\Pi^D(x, \delta)$ would otherwise imply that $x^*(\delta)$ is not the largest maximizer smaller than x^o . The second inequality applies because $x < x^o$. Together, these imply that, for any $x \in (x^*(\delta), x^o)$, $\phi(x) < 0$. But, together with the continuity of ϕ , this contradicts that, as shown above, $\phi(x^o) > 0$. Then, there is no largest maximizer that is smaller than x^o , and so there is no maximizer that is smaller than x^o . Finally, to see that $x^*(\delta)$ is increasing in δ , notice first that if $\phi(x^o) > 0$, as shown above, then $\phi(x) > 0$ for all $x > x^o$ because the only component of $\phi(x)$ that may be negative:

$$\mathcal{G}(x) - (x - \psi(x, 1))G(x) = G(x)[\psi(x, 1) - \mathbb{E}_G(\omega | \omega < x)]$$

must in fact be positive for $x > x^o$ because $\psi(x^o, 1) = \mathbb{E}_G(\omega)$, $\psi(x, 1)$ is strictly increasing and $\mathbb{E}_G(\omega | \omega < x) \leq \mathbb{E}_G(\omega)$. From the implicit function theorem, we have that

$$\frac{dx^*(\delta)}{d\delta} = -\frac{\phi(x^*(\delta))}{\frac{\partial^2 \Pi^D(x^*(\delta), \delta)}{\partial x^2}},$$

but because $x^*(\delta)$ is an interior maximum, it must be that $\frac{\partial^2 \Pi^D(x^*(\delta), \delta)}{\partial x^2} < 0$ while, as argued above, $\phi(x^*(\delta)) > 0$ and this gives us the result. \square

B The Buyer as Principal

We consider the reverse case from that of the main text, where the mechanism designer is the buyer. We assume here that the seller knows her cost realization ω at time 0, whereas the buyer will only discover his value θ at time 1. The only modeling assumption that we need to modify is that instead of asking that the virtual valuation be increasing, here we need the virtual costs $\xi(\omega, 0) = \omega + \frac{G(\omega)}{g(\omega)}$ to be increasing. Still, we will not pursue the general analysis here as it is entirely analogous to that in the main text and will therefore just show how our results still hold via the example where both F and G are uniform

distributions. We use the letter P to indicate the optimal price in these mechanisms and u^M to indicate the buyer's (ex-ante expected) utility in mechanism M when this is maximized.

EAFP mechanism. The buyer is maximizing

$$(\mathbb{E}_F[\theta] - P)P = \left(\frac{1}{2} - P\right)P,$$

so that $P^{EAFP} = \frac{1}{4} = 1 - p^{EAFP}$; $u^{EAFP} = \frac{1}{16} = \pi^{EAFP}$.

EPO mechanism. The buyer is maximizing, at time 1

$$(\theta - P)P,$$

which is maximized by $P^{EPO} = \frac{\theta}{2}$ and gives

$$u^{EPO} = \delta \int_0^1 \left(\theta - \frac{\theta}{2}\right) \frac{\theta}{2} d\theta = \frac{\delta}{12} = \pi^{EPO}.$$

EAO mechanism. It is easy to see that the optimal mechanism being one where $q_0(\theta, \omega) = 1[\theta \geq p \geq \omega]$ as in proposition 3 still applies here and so the buyer is maximizing

$$(\mathbb{E}_F(\theta|\theta > P) - P)P(1 - P) = \frac{1}{2}(1 - P)^2P,$$

so that

$$P^{EAO} = \frac{1}{3} = 1 - p^{EAO}; u^{EAO} = \frac{2}{27} = \pi^{EAO}.$$

D mechanism. We proceed by noting that now we look at equilibria where the buyer offers a price P_0 such that a seller type ω will only be willing to accept it if and only if $\omega \leq \underline{\omega}$, where $\underline{\omega} \in [0, 1]$. Hence, at time 1, the buyer will maximize the following expression:

$$U^{D^1}(P, \theta, \underline{\omega}) = \int_{\underline{\omega}}^P \frac{\theta - P}{1 - \underline{\omega}} d\omega,$$

which gives us

$$P_1^D = \frac{\underline{\omega} + \theta}{2}; u^{D^1}(\theta, \underline{\omega}) = \frac{(\theta - \underline{\omega})^2}{4(1 - \underline{\omega})}.$$

Now, consider time 0. We must guarantee that the marginal seller is indifferent between the profits on offer in the first period ($P_0^D - \omega$) and the expected profits if she declines to

sell. These are

$$\delta \int_{\underline{\omega}}^1 \left(\frac{\underline{\omega} + \theta}{2} - \underline{\omega} \right) d\theta = \frac{\delta}{4}(1 - \underline{\omega})^2,$$

so that $P_0^D = \underline{\omega} + \frac{\delta}{4}(1 - \underline{\omega})^2$. Putting it all together, we get

$$\begin{aligned} U^D(\underline{\omega}, \delta) &= \int_0^1 \underline{\omega} \left(\theta - \underline{\omega} - \frac{\delta}{4}(1 - \underline{\omega})^2 \right) d\theta + (1 - \underline{\omega})\delta \int_{\underline{\omega}}^1 u^{D^1}(\theta, \underline{\omega}) d\theta \\ &= \underline{\omega} \left(\frac{1}{2} - \underline{\omega} - \frac{\delta}{4}(1 - \underline{\omega})^2 \right) + \frac{\delta(1 - \underline{\omega})^3}{12}. \end{aligned}$$

The expression above is maximized by

$$\underline{\omega}^* = \frac{3\delta - 4 + \sqrt{16 - 16\delta + \delta^2}}{4\delta} = 1 - \bar{\theta}^*$$

and, finally, the maximized utility is $u^D = \pi^D$.

Conclusion. We found that $\pi^M = u^M$ for all mechanisms M considered in the paper. Note that the exact equality is a consequence of the assumption that F and G are both identical symmetric distributions but that the ranking of the different mechanisms is the same as for the case when the principal is the seller holds for any general distributions that satisfy our assumptions. As should be expected, in terms of thresholds at which the agent accepts the time 0 offer, we get the reverse ranking that we had for the case when the principal was the seller:

$$P^{EAO} > P^{EAFP} > \underline{\omega}^*.$$

The intuition is simple. The buyer wants to minimize the prices she will pay for the good, so the seller will get the best price in the EAO mechanism (but this will be an at-will offer, so there will be a risk that the buyer will renege). With the EAFP mechanism, the seller gets a lower price, but now the buyer is committed to buying at that price. Finally, in the D mechanism, the buyer will commit to buying only from very low-cost buyers, leaving the opportunity to buy from higher-cost sellers only when she knows her valuation.

C At-will Dynamic Mechanisms

In this appendix, we consider the example with uniform F and G distributions and discuss whether adding an at-will element to the D mechanism might help the seller. The reason we are asking this is natural: as we showed in theorem 1, for a low discount factor, the seller prefers to use the EAO mechanism, which is an at-will posted price mechanism,

whereas for a high discount factor D mechanism performs best. In the D mechanism, the acceptance of the first-period offer ensures service with certainty; it is therefore natural to ask whether replacing this with an at-will option could perform even better.

In particular, we assume that at time 0, the seller offers a price p_0 . If the buyer accepts the offer, the seller still has a right to walk back from the offer (i.e., it is an at-will contract with a price p_0). If the buyer does not accept the offer at time 0, then the seller offers a price $p_1(\omega)$ at time 1, after learning the cost. As with the D mechanism, there is no ex-ante commitment from the seller as to what happens at time 1. Naturally, buyers with types above some threshold $\bar{\theta} \in [0, 1]$ accept the offer, but now the seller completes the transaction if and only if the cost is low enough.

The remaining question is what happens at time 1 if the buyer accepts the offer but the seller reneges? We consider two possible alternatives. In mechanism D1, the interaction ends after the seller walks back from the offer, i.e., if $\theta \geq \bar{\theta}$ and $\omega > p_0$, there is no further transaction. In mechanism D2, the seller can offer another price $\bar{p}_1(\omega)$ at time 1 after such a scenario.

C.1 At-will Dynamic Mechanisms Without a Second Opportunity (D1)

Consider time 1. If the buyer accepted the offer at time 0, then there is no more interaction here. If the buyer rejected the offer, the seller expects the type to be $\theta < \bar{\theta}$ and therefore offers a posted price $p_1(\omega) = \frac{\bar{\theta} + \omega}{2}$. This means that in this event, the seller's expected payoff is $\frac{\bar{\theta}^2}{12}$ and the expected payoff for a type θ buyer is $\frac{1}{4}(2\theta - \bar{\theta})^2$.

Let us now go back to time 0. If the buyer accepts the price p_0 , the trade occurs with probability $Pr(\omega < p_0) = G(p_0) = p_0$ and therefore the expected payoff for the buyer is $p_0(\theta - p_0)$. The marginal type $\bar{\theta}$ is indifferent between accepting and rejecting the offer if

$$p_0(\bar{\theta} - p_0) = \delta \frac{1}{4}(2\bar{\theta} - \bar{\theta})^2 = \delta \frac{\bar{\theta}^2}{4}.$$

This gives us $p_0 = \frac{\bar{\theta}}{2}(\sqrt{1 - \delta} + 1) \in (0, 1)$. The expected profit for the seller from this mechanism at a fixed $\bar{\theta}$ is, therefore

$$\begin{aligned} \Pi^{D1}(\bar{\theta}) &= Pr(\theta \geq \bar{\theta})Pr(\omega \leq p_0)(p_0 - \mathbb{E}_G[\omega | \omega \leq p_0]) + Pr(\theta < \bar{\theta})\delta \frac{\bar{\theta}^2}{12} \\ &= (1 - \bar{\theta})\frac{p_0^2}{2} + \delta \frac{\bar{\theta}^3}{12} = (1 - \bar{\theta})\frac{\left(\frac{\bar{\theta}}{2}(\sqrt{1 - \delta} + 1)\right)^2}{2} + \delta \frac{\bar{\theta}^3}{12}. \end{aligned}$$

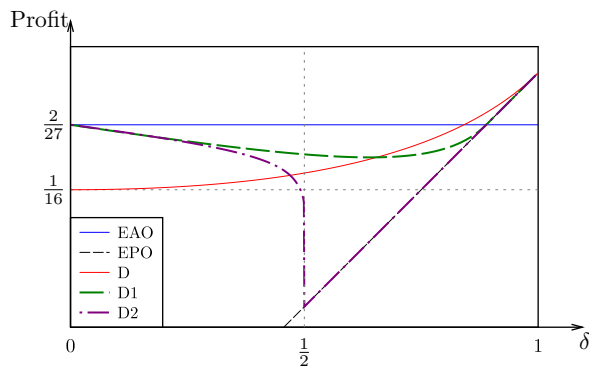


FIGURE 2: PROFITS FROM D1 AND D2 MECHANISMS COMPARED TO OTHER MECHANISMS CONSIDERED

Maximizing this expression with respect to $\bar{\theta}$ gives the solution

$$\bar{\theta}^* = \min \left\{ \frac{4\sqrt{1-\delta} - 2\delta + 4}{6\sqrt{1-\delta} - 5\delta + 6}, 1 \right\}.$$

Notice that $\bar{\theta}^* = 1$ if and only if $\delta \geq \frac{8}{9} = \bar{\delta}$, which means that in this region $\pi^{D1} = \pi^{EPO} < \pi^D$. Moreover, for all $\delta < \bar{\delta}$, the following figure 2 shows that $\pi^{D1} < \pi^{EAO}$. So, D1 is not an improvement compared to the maximum of D and EAO mechanisms.

Let us first compare the D1 mechanism with the EAO mechanism. Both offer an at-will contract at time 0, but if the buyer rejects the offer, D1 may make another offer at time 1, whereas EAO does not. This additional offer comes with a cost and a benefit. The benefit is generating some surplus and capturing some rents from a buyer whose value is relatively low. However, the cost is that the potential of getting the second offer creates information rents in the first period, pushing the price of the initial at-will contract down. The calculations show that the cost always dominates the benefit, except in the limit case where the seller prefers to give up the at-will contract completely and simply offer an optimal time 1 contract (the EPO mechanism). Therefore D1 mechanisms are always (weakly) dominated by either an EAO or an EPO mechanism (in contrast to D mechanisms that always dominate EPO mechanisms and dominate EAO mechanisms for large enough δ).

C.2 At-will Dynamic Mechanisms With a Second Opportunity (D2)

In the second scenario, the D2 mechanism, the seller offers the same set of contracts, except for the additional possibility of making another offer after a history where the buyer accepted the offer, but the seller walked away. Again, this additional offer has the

potential benefit of creating a surplus in situations where the buyer would not have been served otherwise. That is, when the buyer's type is high, and the seller's cost is high, the seller may want to "re-negotiate" with the buyer and offer the object one period later at a different price instead of just canceling the first-period contract.

For fixed thresholds $\bar{\theta}$ and $\bar{\omega}$ (in D1, we had $\bar{\omega} = p_0$), the continuation payoffs after the buyer declines the offer is the same as with the D1 mechanism. If the buyer accepts the offer, the seller learns that $\theta > \bar{\theta}$ and therefore, the posterior distribution is uniform $(\bar{\theta}, 1]$. So, if the seller walks away from the time 0 offer and offers a new price $\bar{p}_1(\omega) \geq \bar{\theta}$, the optimal price $\bar{p}_1(\omega)$ would solve the following problem

$$\max_{p \geq \bar{\theta}} \int_p^1 (p - \omega) Pr(\theta \geq p | \theta \geq \bar{\theta}) = \max_{p \geq \bar{\theta}} \int_p^1 (p - \omega) \frac{1 - p}{1 - \bar{\theta}}.$$

The maximizer is $\bar{p}_1(\omega) = \max \left\{ \bar{\theta}, \frac{1 + \omega}{2} \right\}$. Therefore the continuation profit to the seller is

$$\begin{cases} \frac{(1 - \omega)^2}{4(1 - \bar{\theta})} & \text{if } \omega > 2\bar{\theta} - 1, \\ \bar{\theta} - \omega & \text{if } \omega \leq 2\bar{\theta} - 1. \end{cases}$$

On the other hand, if the seller does not walk away and serves the buyer at a price p_0 , the profit is $p_0 - \omega$. As the marginal type of seller, $\bar{\omega}$, must be indifferent between renegeing and delivering the product, we get a condition

$$p_0 - \bar{\omega} = \begin{cases} \delta \frac{(1 - \bar{\omega})^2}{4(1 - \bar{\theta})} & \text{if } \bar{\omega} > 2\bar{\theta} - 1, \\ \delta(\bar{\theta} - \bar{\omega}) & \text{if } \bar{\omega} \leq 2\bar{\theta} - 1. \end{cases}$$

From the buyer's perspective, accepting the offer at time 0 at price p_0 brings utility $\theta - p_0$ if the seller does not renege (with probability $Pr(\omega < \bar{\omega}) = G(\bar{\omega}) = \bar{\omega}$). If the seller reneges, there is a chance to get a new offer and still get some surplus. However, this payoff is zero for the marginal type (and close to zero for types close to the marginal type), because for all realizations of $\theta \leq 2\bar{\theta} - 1$ the price offer is $\bar{p}_1(\omega) = \bar{\theta}$, leaving no surplus for $\bar{\theta}$, and for all $\theta > 2\bar{\theta} - 1$, the price offer is $\bar{p}_1(\omega) = (1 + \omega)/2 > \bar{\theta}$. Therefore the expected payoff from accepting the offer for the marginal type is $\bar{\omega}(\bar{\theta} - p_0)$. This gives us a condition

$$\bar{\omega}(\bar{\theta} - p_0) = \delta \frac{\bar{\theta}^2}{4}.$$

Combining the last two equations, we get

$$\bar{\theta} - \delta \frac{\bar{\theta}^2}{4\bar{\omega}} = p_0 = \begin{cases} \bar{\omega} + \delta \frac{(1-\bar{\omega})^2}{4(1-\bar{\theta})} & \text{if } \bar{\omega} > 2\bar{\theta} - 1, \\ \bar{\omega} + \delta(\bar{\theta} - \bar{\omega}) & \text{if } \bar{\omega} \leq 2\bar{\theta} - 1. \end{cases} \quad (2)$$

If there is an interior solution to the seller's problem, it must have $(\bar{\theta}, \bar{\omega})$ satisfying the equation above. As this equation generally does not have a closed-form solution, we proceed with the analysis numerically. That is, we compute all combinations of $(\bar{\theta}, \bar{\omega})$ for each δ that satisfy the equality and then maximize the seller's profit, which is

$$\begin{aligned} \Pi^{D2}(\bar{\theta}, \bar{\omega}) &= Pr(\theta < \bar{\theta})\delta \frac{\bar{\theta}^2}{12} + Pr(\theta \geq \bar{\theta})Pr(\omega \leq \bar{\omega})(p_0 - \mathbb{E}_G[\omega | \omega < \bar{\omega}]) \\ &\quad + Pr(\theta \geq \bar{\theta})Pr(\omega > \bar{\omega})\delta \left(\int_{\bar{\omega}}^{\max\{\bar{\omega}, 2\bar{\theta}-1\}} (\bar{\theta} - \omega) d\omega + \int_{\max\{\bar{\omega}, 2\bar{\theta}-1\}}^1 \frac{(1-\omega)^2}{4(1-\bar{\theta})} d\omega \right) \\ &= \delta \frac{\bar{\theta}^3}{12} + (1-\bar{\theta})\bar{\omega} \left(p_0 - \frac{\bar{\omega}}{2} \right) \\ &\quad + (1-\bar{\theta})(1-\bar{\omega})\delta \begin{cases} \frac{(1-\bar{\omega})^3}{12(1-\bar{\theta})} & \text{if } \bar{\omega} \geq 2\bar{\theta} - 1, \\ \frac{1}{6} \left(4\bar{\theta}^2 - 6\bar{\theta}\bar{\omega} - 2\bar{\theta} + 3\bar{\omega}^2 + 1 \right) & \text{if } \bar{\omega} < 2\bar{\theta} - 1. \end{cases} \end{aligned}$$

We find that when $\delta \geq \frac{1}{2}$, equation (2) has no solutions. This means that either $\bar{\omega} = 0$ or $\bar{\omega} = 1$. Note that in both cases, we reach the same conclusion. If $\bar{\omega} = 1$, then the seller never reneges, i.e., does not renege even at cost $\omega = 1$. This can happen only if $p_0 = 1$, which in turn means that the buyer never accepts the first-period offer. Therefore the mechanism is equivalent to the EPO mechanism. On the other hand, if $\bar{\omega} = 0$, then the seller always reneges, which means that the buyer would accept the offer only if $p_0 = 0$. Therefore the time 0 contract is a worthless agreement; at zero price, the seller offers a contract that he never fulfills. Again, this implies that we have a mechanism that is equivalent to the EPO mechanism.

On the other hand, if $\delta < \frac{1}{2}$, equation (2) not only has a solution, but has a continuum solution. We compute $\Pi^{D2}(\bar{\theta}, \bar{\omega})$ for all such combinations $(\bar{\theta}, \bar{\omega})$ and find a unique profit maximizing combination, which we report on figure 2.

As figure 2 (the line for D2) illustrates, the profit from the D2 mechanism is strictly lower than the EAO mechanism for all $\delta < \frac{1}{2}$ and then equal to the profit from the EPO mechanism for all $\delta \geq \frac{1}{2}$. This means that it is again dominated by the combination of EAO and EPO mechanisms.

Intuitively, the D2 mechanism can potentially be either better or worse than the D1 mechanism. It may be better because it initiates a contract in some situations where D1

could not. However, it may also be worse, as the seller now has an additional incentive to renege, because the value of the time 0 contract is reduced. The latter effect is relatively small when δ is very low, as renegeing would delay the delivery of the product, and with low δ , this makes renegeing unappealing. Therefore D2 is slightly better than D1 for very low values of δ . Still, it is never dominating EAO for low values of δ for the same reasons discussed above. As δ increases, the negative effect of the additional flexibility increases, making D2 worse than D1. And as δ increases further, the effect is so large that any time 0 contract would be renegeed for sure. Thus, the mechanism converges to the EPO mechanism. Once again, these results show that at-will dynamic contracts can never improve on the optimal ex-ante or ex-post mechanisms, while the D contract which has no at-will clause, will be better than both of those when δ is sufficiently high. This suggests that while a dynamic contract (with limited commitment) and at-will clauses are two ways for the seller to address the problem of receiving her private information with delay, they should be seen as substitutes rather than complements. Depending on the value of the discount factor one may be preferable to the other, but it is never optimal to design a mechanism with both features.