

Pricing Novel Goods*

Francesco Giovannoni[†] Toomas Hinnosaar[‡]

August 2024

Abstract

We study a bilateral trade problem where a principal has private information that is revealed with delay, such as a seller who does not yet know her production cost. Postponing the contracting process incurs a costly delay, while early contracting with limited information can create incentive issues, as the principal might misrepresent private information that will be revealed later. We show that the optimal mechanism can effectively address these challenges by leveraging the sequential nature of the problem. The optimal mechanism is a menu of two-part tariffs, where the variable part is determined by the principal's incentives and the fixed part by the agent's incentives. As two-part tariffs might be impractical in some applications, we also study price mechanisms. We show that the optimal price mechanism often entails trade at both the ex-ante and ex-post stages. Dynamic price mechanisms can lower the cost of delay by transacting with high-type agents early and relax the incentive constraints by postponing contracts with lower-type agents. We also generalize our analysis to costly learning and study ex-post efficiency in our context.

*We are grateful to Rossella Argenziano, Laura Doval, Duarte Gonçalves, Johannes Hörner, Andreas Kleiner, Alessandro Pavan, Ludvig Sinander, Vasiliki Skreta, and seminar participants at the University of Bristol, University of Nottingham, University of Essex, Queens University Belfast, University of Copenhagen, University of Manchester, University of Leicester, University of Bath, Cardiff University, Durham Economic Theory Conference, Birmingham Economic Theory Workshop, SAET (Paris), Junior Workshop in Economic Theory (Bonn), Lancaster Game Theory Conference, QMUL/City Workshop in Economic Theory (London), CEPS PolEcon Workshop (Paris), Conference on Mechanism and Institution Design (Budapest) for their comments and suggestions. A one-page abstract of an earlier version of this paper was published in Proceedings of the 24th ACM Conference on Economics and Computation (EC '23) <https://doi.org/10.1145/3580507.3597694>.

[†]University of Bristol, francesco.giovannoni@bristol.ac.uk

[‡]University of Nottingham and CEPR, toomas@hinnosaar.net

1 Introduction

Many economic transactions are initiated by the buyer, and at the first possibility of contracting, the seller often does not yet know all the details, such as production costs. In this paper, we ask whether it is better for the seller to contract early or to wait until all relevant information is available. For example, consider a bilateral trade between a customer and a monopolistic service provider, such as a plumber or electrician. The buyer knows privately their willingness to pay for the repair and approaches the seller to get the job done. At this point, the seller does not yet know what needs to be done to fix the problem. One natural solution is to inspect the issue first and then, once the details are clear, provide a quote that the buyer can approve or reject. This mechanism has one weakness—since the work cannot begin until the uncertainty is resolved—trade is delayed. Another approach is to offer a contract where the buyer agrees to cover the contractor’s costs and fees, whatever they turn out to be, up to a specified limit, allowing work to start immediately. This type of contract avoids delays but creates potential incentive issues for the seller—after learning what the buyer is willing to pay, the seller may be incentivized to inflate the reported costs.

We model such situations using a standard bilateral trade framework, with three key assumptions: First, delaying the contract until the seller’s private information is revealed incurs costs for both parties due to delay. Second, the seller cannot commit to truthfully reporting her private information. Third, the seller cannot ignore the information provided by the buyer. In scenarios where these assumptions do not hold, the optimal mechanism is straightforward: a take-it-or-leave-it offer based on the seller’s type. However, under the conditions we consider, this proposed mechanism is not implementable.¹

The trade-offs we study here are relevant in significant public policy contexts. For instance, consider a pharmaceutical company acting as a monopolistic seller developing a new cure for some disease. In this scenario, a national health authority (the buyer) privately knows the value of this product for its population, but the seller is initially uncertain about the production costs, which will only become clear after extensive research. Delaying the contract until this information is available can be costly because early contracting would enable the health authority to organize approval and delivery much sooner. Although this paper focuses on cases where the principal is a seller and the agent is a buyer, our results also apply in reverse situations, provided the principal is the party that learns their private information with a delay. For example, consider a defense department

¹If there were no cost of delay, both parties could simply wait until the seller’s type is realized. If the seller could commit to truthfully revealing her private information once it is known, the optimal mechanism could be replicated with a type-contingent menu of mechanisms. Finally, if the seller could ignore the buyer’s information, she would have no incentive to misrepresent her type once it is realized.

as the buyer, contemplating the purchase of a new weapons system from a seller who privately knows the production costs. The defense department faces uncertainty regarding the future usefulness of the system. Over time, this uncertainty may resolve, but waiting for this information could result in delayed delivery of the system.

Our first main result shows that an optimal mechanism can effectively address the challenges posed by asymmetric information and timing in bilateral trade with information that arrives over time. Specifically, we show that a menu of two-part tariffs is an optimal mechanism under our assumptions. These two-part tariffs consist of a fixed fee and an *at-will* delivery price. The buyer chooses a tariff, and the fixed fee is paid upfront. The seller commits to stick to the delivery price chosen by the buyer but reserves the right to refuse to trade upon learning her own type. In equilibrium, the seller will agree to trade when the price component of the tariff chosen by the buyer turns out to be higher than the seller's type, ensuring that the seller has no incentive to misreport her type. At the same time, it also implies that higher prices come with a higher probability of trade, which creates a potential for screening. However, if the menu only consisted of delivery prices, the buyer would not want to choose optimal prices from the seller's perspective. Thus, the tariff adds a specific fixed fee to each price, which can be positive or negative, and this fee adjusts expected payments to restore incentive compatibility for the buyer and optimality for the seller. In equilibrium, intermediate buyer types choose tariffs where the price is low, resulting in ex-post information rents that are too high, necessitating an additional charge to mitigate them. Conversely, the seller wants high types to choose very high delivery prices, which requires an ex-ante subsidy to provide the necessary informational rents. This combination leads to a novel form of non-monotonicity in price discrimination because as the delivery fee (i.e., per-unit price) increases, the corresponding fixed fee initially rises and then falls. This finding contrasts sharply with typical nonlinear pricing models, where a higher quantity (or probability of trade) is generally associated with a higher fixed fee.

Our second main result studies price mechanisms that do not allow the use of two-part tariffs. This analysis is motivated by the observation that optimal two-part tariff menus, especially those with the non-monotone structure we discussed, are rarely observed in practice. We begin by examining optimal static price mechanisms, where contracting occurs either before or after type information becomes available, without dynamic screening. When the discount factor is sufficiently high, the optimal static price mechanism involves waiting until the seller's type is revealed and then offering a take-it-or-leave-it price that will depend on that type. However, when the discount factor is low, the optimal price mechanism consists of a menu of offers made before the seller's type is revealed. This optimal menu consists of (possibly many) at-will prices and a guaranteed-delivery price. In

equilibrium, the highest buyer types choose guaranteed delivery at a high price, while the at-will menu allows for further screening of intermediate buyer types. While these are the optimal static price contracts, we also show that when the discount factor is large enough, the optimal price mechanism becomes dynamic, offering early delivery to high types to minimize delay costs and late delivery to low types to mitigate incentive problems.

In the final part of our paper, we extend our model in two directions. First, we consider the possibility that the seller only learns her type if she pays a cost. This could occur, for example, if the good has never been produced before, requiring the seller to invest in order to learn the production cost, or if the good is available but the seller needs to perform research or hire an outside expert to evaluate its worth. We show that our main result still applies: if learning the cost is necessary in order to trade, the seller either chooses not to trade (if the cost is too high) or uses the two-part tariff mechanism but trades with a smaller set of buyers. If learning the cost is optional, she will offer to trade with guaranteed delivery when the cost is high. However, if the cost is sufficiently low, she will learn her type when the buyer type is intermediate and offer a two-part tariff, whereas if the buyer type is sufficiently high, she will offer a guaranteed delivery option. These high buyer types are willing to pay enough, so learning her type is not worth it for the seller.

Finally, we ask whether ex-post efficiency is implementable in our setting. In comparison with the classic Myerson and Satterthwaite (1983) setting, we have the same constraints for the buyer, while the seller faces weaker participation constraints (she learns her type with delay) but stronger incentive compatibility constraints (she can update her information about the buyer's type before revealing her own). We show that in this setting, the former effect more than compensates for the latter, and ex-post efficiency is implementable. This is again achieved with two-part tariffs, but this time the seller always subsidizes the buyer regardless of whether trade occurs.

Related literature. The paper builds on the literature on seller profit maximization in mechanism design, starting with Myerson (1981) and Riley and Samuelson (1981). Here, specifically, we look at the questions in the context of bilateral bargaining. Williams (1987) was the first to look at profit maximization in a bilateral bargaining context where both sides face Bayesian incentive compatibility and interim participation constraints. It found that the optimal allocation required that the seller's type be smaller than the buyer's virtual valuation. In this paper, we show that if the seller learns her type after the buyer, then that same allocation is an upper bound on what the seller can achieve in our context. More interestingly, we show that the same allocation can be obtained, but two-part tariffs are required. Myerson and Satterthwaite (1983) is the seminal work

in bilateral bargaining and shows that in the same environment as in Williams (1987) ex-post efficiency cannot be implemented. In one of our extensions, we show that in our setting two-part tariffs allow for the implementation of ex-post efficient trade.

The sequential way buyer and seller learn their type also introduces the possibility of considering both static and dynamic mechanisms. 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). Importantly for our results, starting from Stokey (1979); Conlisk, Gerstner and Sobel (1984) this literature has emphasized that if buyers' valuations are persistent 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 (making it, by definition, persistent). In such situations however, we show that if we focus on price mechanisms and the discount factor is sufficiently high, dynamic contracting can 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 private information at the contract design stage and cannot commit not to use this information. For the specific bilateral bargaining setting we are considering, however, Yilankaya (1999) shows that the Williams (1987) result, which assumes that mechanism is designed by a third party that wishes to maximize the seller's profit continues to hold when the informed seller designs the mechanism. It could be argued that we expand on this literature by looking at a "not-yet informed" principal problem, where the issue is not that the principal's choice of a mechanism itself may leak information to the agent but that the principal may not be relied on to report her private information, once learned, truthfully.

There is also a 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 (2024); Fugger, Gretschko and Pollrich (2019); Liu, Mierendorff, Shi and Zhong (2019); Doval and Skreta (2022). In these models, the optimal contract is typ-

ically dynamic with a decreasing sequence of offers. In our model, with two-part tariffs the seller can achieve optimality by trading immediately, but when one focuses on pricing contracts, then as long as the buyer and seller value the future enough, dynamic contracts can improve on static ones. This is because a dynamic contract allows the seller to discriminate between high types—who trade immediately—and intermediate types, who trade later, once the seller’s cost is known. The key here is that the seller does not have to learn exactly the buyer’s type at time 0, only whether this type is high or not. This allows the seller to dampen her future’s self incentives when trading with intermediate types.²

Closely related to this point about commitment, there is also a connection to the literature on credibility in mechanisms (Akbarpour and Li, 2020). Just as in that literature, the principal can commit to the outcome function, but she cannot commit to revealing her private information truthfully. The difference is that in our setting, such information is her future type, whereas credibility concerns itself with the possibility that the principal may lie about the reports that participants in the mechanism may have sent her.

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 may generate more social surplus than a guaranteed delivery contract. The crucial difference is that in our setting it is the principal whose private information is learned with delay, whereas in Schmitz (2022), it is the agent’s. The implications are very different: we show that the seller’s incentive issues force the optimal contract to be a two-part tariff. If we limit ourselves to pricing, then optimal ex-ante price contracts always involve menus of offers, which have both an at-will and a guaranteed delivery component. This is also true in dynamic contracts.

The rest of the paper is as follows. In section 2, we introduce the model, interpret it, and show how it can be reduced to a simpler model. In section 3, we study general mechanisms, starting from benchmarks that provide an upper bound for our analysis, and then show how these outcomes could be achieved in our setting. In section 4, we study a special class of simpler mechanisms we call price mechanisms. We first characterize the optimal static price mechanisms and then show that in some situations dynamic mechanisms can do strictly better than any static mechanism. In section 5, we study two types of extensions to our analysis: costly learning and efficiency. Finally, section 6 concludes. In the appendix A, we provide proofs and additional technical details skipped

²There 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.

in the text.

2 Model

2.1 Setup

Environment. We consider a standard bilateral trade problem, where the seller has a single unit of an indivisible good available for sale to a buyer.³ The buyer has a private type $\theta \in \Theta = [0, 1]$, which represents his valuation for the object, while the seller has a private type $\omega \in \Omega = [0, 1]$ that represents her valuation of the object.⁴ There are two time periods $\tau \in \{0, 1\}$: trade can be agreed at time 0 or at time 1, but if it is agreed at time 1, payoffs for both players have a common discount factor $\delta \in (0, 1)$.

We assume that it is common knowledge that θ has a distribution F with density $f(\theta) > 0$ while ω has a distribution G with density $g(\omega) > 0$. We assume these distributions are independent and satisfy the usual regularity conditions for both distributions. In particular, the regularity conditions guarantee that both the virtual valuation $\psi(\theta) = \theta - \frac{1-F(\theta)}{f(\theta)}$ and the virtual cost $\phi(\omega) = \omega + \frac{G(\omega)}{g(\omega)}$ are strictly increasing. For some of our results, we further assume that these are differentiable functions. As is standard, the notation $\mathbf{1}[A]$ represents the indicator function that is equal to 1 if A is true and 0 otherwise. Further, we will often use the notation $\mathcal{G}(x)$ as a shorthand for $\int_0^x G(y)dy$ and the notation $\psi^+(\theta)$ as a shorthand for $\max[\psi(\theta), 0]$.

Timing. The timing unfolds as follows. To begin with, the seller commits to a mechanism $\mathcal{M} = (\mathcal{B}_0, \mathcal{B}_1, \mathcal{S}, q_0, q_1, t_0, t_1)$, which we describe below.

At time 0, the buyer privately learns his valuation and chooses whether to participate in mechanism \mathcal{M} . If the buyer chooses not to participate, both agents get a payoff 0. If the buyer participates, he sends a public and contractible message $b_0 \in \mathcal{B}_0$. The seller observes this message and privately learns the cost ω . Only after this, the seller sends a public and contractible message $s \in \mathcal{S}$. Then, the seller delivers the goods to the buyer with probability $q_0(b_0, s)$ and the buyer pays $t_0(b_0, s)$ to the seller.⁵

At time 1, having observed b_0 and s , the buyer sends another public and contractible message $b_1 \in \mathcal{B}_1$. Then, the seller delivers the good to the buyer with probability $q_1(b_0, b_1, s)$, and the buyer pays $t_1(b_0, b_1, s)$ to the seller.

³It is equivalent to a procurement problem with appropriate relabeling.

⁴Our results can be (with appropriate adjustments) 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 simplifies the exposition.

⁵Transfers could be negative, which, of course, means that the buyer receives a payment from the seller.

Payoffs. The buyer's utility after messages (b_0, b_1, s) and with valuation θ is

$$u(b_0, b_1, s|\theta) = \theta q(b_0, b_1, s) - t(b_0, b_1, s),$$

where $q(b_0, b_1, s) = q_0(b_0, s) + \delta q_1(b_0, b_1, s)$ is the discounted allocation rule and $t(b_0, b_1, s) = t_0(b_0, s) + \delta t_1(b_0, b_1, s)$ is the discounted transfer rule. Similarly, the seller's profit is

$$v(b_0, b_1, s|\omega) = t(b_0, b_1, s) - \omega q(b_0, b_1, s).$$

Strategies and beliefs. A (behavioral mixed) strategy profile consists of $(\alpha, \beta_0, \sigma, \beta_1)$, where $\alpha(\theta)$ is the probability that the buyer participates in the mechanism, and (conditional on buyer participation) $\beta_0(\theta)$ is the probability distribution over the buyer's message space \mathcal{B}_0 at time 0, $\sigma(\omega, b_0)$ is the probability distribution over the seller's message space \mathcal{S} at time 0, and $\beta_1(\theta, b_0, s)$ is the probability distribution over the buyer's message space \mathcal{B}_1 at time 1.

The only relevant belief updating happens when the seller observes b_0 and forms an updated belief about the buyer's type θ . This is because at time 0, when the buyer makes his decisions, the mechanism has revealed no information about the seller's type, so the beliefs are equal to posterior G , and at time 1 when the buyer chooses the second message to send, the information about seller's type is irrelevant, as the buyer's payoff depends only on message s and not the seller's type ω directly. Given this, with a slight abuse of notation, we denote the seller's updated beliefs by $F(\theta|b_0)$.

Equilibrium conditions. We are looking for a perfect Bayesian equilibrium (PBE) of this game. That is, both buyer and seller behave optimally at each decision node, conditional on messages exchanged up to that point, and beliefs are updated according to Bayes' rule whenever possible. A pure strategy PBE has the following four equilibrium conditions⁶

$$\int_0^1 u(\beta_0(\theta), \beta_1(\theta, \beta_0(\theta)), \sigma(\omega, \beta_0(\theta))), \sigma(\omega, \beta_0(\theta))|\theta) dG(\omega) \geq 0, \quad \forall \theta \quad (\text{IRB}_0)$$

$$\beta_0(\theta) \in \operatorname{argmax}_{b_0 \in \mathcal{B}_0} \int_0^1 u(b_0, \beta_1(\theta, b_0), \sigma(\omega, b_0)), \sigma(\omega, b_0)|\theta) dG(\omega), \quad \forall \theta \quad (\text{ICB}_0)$$

$$\sigma(\omega, b_0) \in \operatorname{argmax}_{s \in \mathcal{S}} \int_0^1 v(b_0, \beta_1(\theta, b_0, s), s|\omega) dF(\theta|b_0), \quad \forall \omega, \forall b_0 \in \mathcal{B}_0 \quad (\text{ICS}_0)$$

$$\beta_1(\theta, b_0, s) \in \operatorname{argmax}_{b_1 \in \mathcal{B}_1} u(b_0, b_1, s|\theta), \quad \forall \theta, \forall b_0 \in \mathcal{B}_0, \forall s \in \mathcal{S} \quad (\text{ICB}_1)$$

⁶Mixed strategy PBEs are defined analogously, but the notation is somewhat more cumbersome, which is why we postpone these to appendix A.1.1.

Additionally, we need to impose the feasibility constraints:

$$q_0(b_0, s), q_1(b_0, b_1, s) \geq 0, q_0(b_0, s) + q_1(b_0, b_1, s) \leq 1, \quad \forall b_0 \in \mathcal{B}_0, b_1 \in \mathcal{B}_1, s \in \mathcal{S} \quad (\text{F})$$

Maximization problem. The seller's goal is to maximize her expected ex-ante profit (from now on, just profit) by choosing a mechanism \mathcal{M} . That is, the maximization problem is⁷

$$\max_{\mathcal{M}} \int_0^1 \int_0^1 v(\beta_0(\theta), \beta_1(\theta, \beta_0(\theta)), \sigma(\omega, \beta_0(\theta)), \sigma(\omega, \beta_0(\theta)) | \omega) dG(\omega) dF(\theta). \quad (1)$$

subject to (F), (IRB₀), (ICB₀), (ICS₀), and (ICB₁).

2.2 Interpretation

The model captures the dynamic interaction between the buyer and the seller without imposing additional commitment assumptions. First, consider a class of mechanisms where the seller opts not to engage with the buyer at the final stage (time 1). This is achieved by setting $q_1 \equiv t_1 = 0$ and making $\mathcal{B}_1 = \{0\}$. The transaction, if it happens, occurs at time 0. Any information the buyer communicates to the mechanism through messaging (message b_0) is received before the seller sends her message s . Thus, agreement must be reached before the seller's uncertainty is realized. Importantly, there is a separate set of (ICS₀) constraints for each b_0 , i.e., the mechanism must account for incentives of the seller, but there is no discounting.

Conversely, the seller can also choose to only trade at time 1 by setting $q_0 \equiv t_0 \equiv 0$ and $\mathcal{B}_0 = \{0\}$. Here, at time 0, the seller learns the cost ω and sends the message s without learning anything about the buyer's type. Thus, the set of (ICS₀) constraints is reduced to one interim constraint, and the remaining constraints on the seller won't be binding.⁸ The issue is, however, that all payoffs in this scenario are discounted.

2.3 Reduced-Form Mechanisms

In many mechanism design problems, the revelation principle is crucial because it allows the analysis to focus only on direct mechanisms where agents report their types. In our model, the revelation principle does not apply because the information that the buyer reveals to the mechanism has a dual role: it enters the outcome function (as is standard),

⁷Again, we write it for the pure strategies here and postpone the mixed strategy formulation to appendix A.1.1.

⁸This follows from Williams (1987). A detailed discussion is in section 3.1

but it also may change the seller's beliefs and, hence, her incentives. Nevertheless, we can simplify the problem using a similar idea. Lemma 1 shows that we can focus on reduced-form mechanisms, where both agents report only their types. To account for the dual role of information, we introduce a garbling device, which transfers information about the buyer's type to the seller and also influences early allocation.

Specifically, a reduced-form mechanism $\mathcal{M}^R = (\mathcal{B}, \mu, q_0, q_1, t_0, t_1)$, is such that:

1. The buyer observes value $\theta \in \Theta$ and reports $\theta' \in \Theta$ to the garbling device.
2. The garbling device outputs a public message $b \in \mathcal{B}$, with probability $\mu(b|\theta')$. This is common knowledge.⁹
3. The seller observes message b and upon learning the cost $\omega \in \Omega$, reports $\omega' \in \Omega$ to the mechanism.
4. Ex-ante transaction: with probability $q_0(b, \omega')$, the object is transferred to the buyer, and the latter pays $t_0(b, \omega')$ to the seller.
5. Ex-post transaction: with probability $q_1(b, \omega', \theta')$ the object is transferred to the buyer, and the latter pays $t_1(b, \omega', \theta')$ to the seller. The payoffs from this transaction are discounted by $\delta \in (0, 1)$.

The seller's problem in the reduced-form environment is

$$\max_{M \in \mathcal{M}^R} \int_0^1 \int_0^1 \mathbb{E}_{\mu(b|\theta)} [t(b, \omega, \theta) - \omega q(b, \omega, \theta)] dG(\omega) dF(\theta) \quad (2)$$

subject to $q_0, q_1 \in [0, 1], q_0 + q_1 \leq 1$,

$$U(\theta|\theta) \geq 0, \forall \theta, \quad (\text{IRB})$$

$$U(\theta|\theta) \geq U(\theta'|\theta), \forall \theta, \theta', \quad (\text{ICB})$$

$$V(\omega|\omega, b) \geq V(\omega'|\omega, b), \forall b, \omega, \omega', \quad (\text{ICS})$$

⁹ \mathcal{B} may be countable or uncountable so that $\mu(b|\theta)$ may be both a probability mass function and a (truncated) density.

where

$$\begin{aligned}
U(\theta'|\theta) &= \int_0^1 \mathbb{E}_{\mu(b|\theta')} [\theta q(b, \omega, \theta') - t(b, \omega, \theta')] dG(\omega), \\
q(b, \omega', \theta') &= q_0(b, \omega') + \delta q_1(b, \omega', \theta'), \\
t(b, \omega', \theta') &= t_0(b, \omega') + \delta t_1(b, \omega', \theta'), \\
V(\omega'|\omega, b) &= \int_0^1 (t(b, \omega', \theta') - \omega q(b, \omega', \theta)) dF(\theta|b).
\end{aligned}$$

Remember that $F(\theta|b)$ is the posterior belief about the buyer's type θ after observing message b . Note also that we can assume without loss of generality that each $b \in \mathcal{B}$ is meaningful, i.e., there exists θ such that $\mu(b|\theta) > 0$.

Lemma 1. *For any mechanism \mathcal{M} and any equilibrium in it, there exists a reduced-form mechanism \mathcal{M}^R where a truthful equilibrium exists and guarantees both players the same equilibrium payoffs.*

The proof is in appendix A.1.1. Two remarks are in order. First, note that Lemma 1 does not say that for any mechanism and truthful equilibrium in the reduced problem (2), there exists a mechanism and an equilibrium in (1) that is payoff equivalent. This is because, in the reduced problem, the buyer only makes one announcement, leading to weaker constraints than in the general problem, where the buyer selects two distinct messages. However, this is still helpful if the mechanisms we identify through the reduced problem (2) are feasible in the general problem (1) so that if they are optimal for the former, they will be optimal for the latter. Second, the reduced-form mechanism design problem includes a dynamic mechanism design problem (the choice of q_0, q_1, t_0, t_1) as well as an information design problem (the choice of \mathcal{B} and μ). These are not, however, independent of each other as the choice of information structure (\mathcal{B}, μ) affects the outcome functions.

3 General Mechanisms

In our reduced-form model, we make three crucial assumptions:

1. The seller cannot commit to revealing her type truthfully.
2. The garbling device's message is publicly observed.
3. If a trade occurs at time 1, both the buyer and seller's payoffs are discounted by $\delta \in (0, 1)$.

As implied by Lemma 1, these same assumptions hold for the general model.¹⁰

In this section, we first consider changing the first assumption to a setting where the seller *can commit* to revealing her type. This constitutes a natural benchmark and gives us an upper bound for the seller’s profit among all possible mechanisms. We also show that this upper bound can be easily obtained if the seller could not observe any information from the garbling device and the implementation of the mechanism was delegated to a third party. Finally, we show that if $\delta = 1$, it would also be easy to achieve the upper bound by simply waiting for the seller to learn her cost and making trade possible only at time 1.

In section 3.2, we then show that, even when there is discounting and the seller cannot commit to both revealing her type and ignoring the information provided on the buyer’s type, the benchmark allocation and profit for the seller can still be achieved. This is done through a mechanism where the buyer chooses from a menu of two-part tariffs. For each element of the menu, the first payment is a price for the good with an “at-will” clause: the seller will deliver the good if she later learns that her type is below such price. The second payment is either a fee (if positive) or a subsidy (if negative) and is given regardless of whether trade occurs. We discuss this mechanism and argue that price mechanisms, where money changes hands only if there is trade, may be more realistic.

Before we begin the analysis itself, we argue that it is without loss to focus on a special class of reduced-form mechanisms.

Definition 1. *In a fully-revealing mechanism, $\mathcal{B} = \Theta$ and $\mu(\theta|\theta) = 1$.*

In other words, in a fully-revealing mechanism, the buyer publicly announces her type. It is easy to see that in the benchmark case, where the seller commits to reveal her type truthfully, the seller has no role to play and none of the issues that the sequential nature of our problem arise. Given that, we can use the revelation principle and just focus on direct mechanisms, which in this case are equivalent to fully-revealing mechanisms where trade can only happen at time 0: $q_1 = t_1 = 0$.

When we then turn to analyzing the case where the seller does play a role, the focus on fully revealing mechanisms is, in principle, with loss. However, proposition 2 shows that the benchmark utility and profits for the seller can be obtained via a fully-revealing mechanism.

As we shall see in section 4, when we impose that mechanisms are prices, fully revealing mechanism can sometimes be improved upon.

¹⁰In the general model, it is not the garbling device’s message that is observed, but that of the buyer directly. The key point is that, in order to implement the mechanism, the seller must observe the variables that determine the outcome.

3.1 Benchmarks

S commits to reveal her type truthfully. As discussed, above, when the seller plays no role, then we can apply the revelation principle. Thus, the problem becomes:

$$\max_{q_0, t_0} \int_0^1 \int_0^1 [t_0(\theta, \omega) - \omega q_0(\theta, \omega)] dF(\theta) dG(\omega) \quad (3)$$

subject to

$$\begin{aligned} \int_0^1 [\theta q_0(\theta, \omega) - t_0(\theta, \omega)] dG(\omega) &\geq 0, \forall \theta \\ \int_0^1 [\theta q_0(\theta, \omega) - t_0(\theta, \omega)] dG(\omega) &\geq \int_0^1 [\theta q_0(\theta', \omega) - t_0(\theta', \omega)] dG(\omega), \forall \theta, \theta' \end{aligned}$$

We then have:

Proposition 1. *The solution to 3 is the allocation:*

$$q_0(\theta, \omega) = \mathbf{1}[\omega \leq \psi(\theta)],$$

which can be implemented with transfers

$$t_0(\theta, \omega) = \psi^{-1}(\omega) \mathbf{1}[\omega \leq \psi(\theta)],$$

The expected utility for the seller from this solution is

$$V^B = \int_0^1 \int_0^1 (\psi(\theta) - \omega) \mathbf{1}[\psi(\theta) \geq \omega] dF(\theta) dG(\omega) = \int_{\psi^{-1}(0)}^1 \mathcal{G}(\psi(\theta)) dF(\theta)$$

The proof of this result is standard and, in fact, the result itself is implicit in Williams (1987) as discussed below. We state it here explicitly for clarity. The proof is relegated to appendix A.2.1. It is also immediate to check that this mechanism is also feasible in the general environment.

We refer to the allocation obtained here as the *benchmark* allocation and to its corresponding expected utility for the seller as the *benchmark* profit for the seller. One thing we wish to emphasize at this point is that while the transfers $t_0(\theta, \omega)$ described above are optimal, they are not unique.¹¹ Still, it is important to note that in the benchmark case, a natural implementation of the benchmark mechanism is one where the seller offers the buyer a price menu $\{\psi^{-1}(\omega)\}_{\omega \in \Omega}$ before ω is revealed. The buyer reports the highest price

¹¹This is because the buyer's constraints only put restrictions on the buyer's expected transfers.

p she is willing to accept; once ω is revealed, trade occurs at a price $\psi^{-1}(\omega)$ if this is no greater than p , while no trade or payments are made otherwise.

This is clearly not implementable if the seller cannot commit to truthfully revealing her type. Any maximal price p chosen by the buyer effectively reveals her type to the seller. The latter, upon learning ω , has an incentive, if $\psi^{-1}(\omega) < p$, to report instead that her type is ω' , where $\psi^{-1}(\omega') = p$. We discuss below how our assumptions rule out two immediate solutions to this problem: the problem would not exist if the seller could commit to ignoring the information about the buyer she learns or if there was no discounting, allowing the seller to simply wait for her type ω to be revealed and then propose a single price $\psi^{-1}(\omega)$.

S does not observe the garbling device's message. As we've shown above, removing incentive compatibility constraints for the seller leads to an upper bound for our problem. Now, we show that this upper bound could be reached if the seller did not observe the information that the garbling device provides and implementation was delegated to a third party. Formally, we are assuming here that problem (2) holds in its entirety, but we impose that for any b , $F(\theta|b) = F(\theta)$. It is easy to see that with this restriction, the seller's expected profit when reporting ω' does not depend on b :

$$\int_0^1 \mathbb{E}_{\mu(b|\theta)} [t(b, \omega', \theta) - \omega q(b, \omega', \theta)] dF(\theta)$$

and so, slightly abusing notation, we can rewrite (ICS) as

$$\int_0^1 [\tilde{t}(\omega, \theta) - \omega \tilde{q}(\omega, \theta)] dF(\theta) \geq \int_0^1 [\tilde{t}(\omega', \theta) - \omega \tilde{q}(\omega', \theta)] dF(\theta), \forall \omega, \omega'. \quad (\text{ICS}^\emptyset)$$

where $\tilde{q}(\theta, \omega) = \mathbb{E}_{\mu(b|\theta)} q(b, \omega, \theta)$ and analogously for \tilde{t} . But then we can apply to the constraint above, to the buyer's constraints, and to the objective function, the same construction just as we did in the case where the seller can commit to the truth. Hence, we end up with problem (3) plus the incentive compatibility constraint for the seller (ICS[∅]) described above where \tilde{q} and \tilde{t} are replaced by q_0 and t_0 .

Williams (1987) shows that the benchmark solution satisfies the (ICS[∅]) constraint, making V^B achievable.¹² Thus, it is not the existence of an incentive compatibility constraint for the seller, per se, that is problematic. Rather, the problem arises if the seller announces her private information after updating her beliefs about the buyer's valuation.

¹²In fact, Williams (1987) shows that we obtain the same solution if we impose (ICS[∅]) and interim individual rationality for the seller, so Proposition 1 implies that these constraints are, in fact, slack when we are maximizing the seller's expected utility.

Thus, a solution to the implementation of the benchmark mechanism would be a device that effectively implies that buyer and seller could agree to trade at time 0 without the latter knowing it.¹³

No discounting. Suppose that $\delta = 1$, so that waiting for ω to realize is without loss for either player. We now show that the benchmark payoff can be achieved. Since

$$q_0(b, \omega) + q_1(b, \omega, \theta) = q(b, \omega, \theta)$$

(and, analogously for t), then we can wlog focus on mechanisms where $q_0 = t_0 = 0$ because (q_1, t_1) can implement anything that (q_0, t_0) can.

Now, set $\mathcal{B} = \{b\}$ so that it must be that $\mu(b|\theta) = 1$ for all θ and therefore $F(\theta|b) = F(\theta)$. Thus, effectively, $\mathcal{M}^R = (q_1(\theta, \omega), t_1(\theta, \omega))$ since both \mathcal{B} and μ are degenerate and problem (2) becomes again problem (3) plus the (ICS^o) constraint, the only difference being that it applies to (q_1, t_1) instead of (q_0, t_0) . Hence, the optimal mechanism is the benchmark mechanism where trade happens at time 1 and this is without loss as there is no discounting.

When there is no discounting, the seller can, in essence, commit to not learning anything about the buyer's type by imposing that trade can only occur at time 1. Then, as we discussed in the previous case, she can implement the benchmark mechanism because it is incentive compatible for her future self (who still has prior beliefs on θ) to truthfully report her type. Of course, this works, and it achieves the benchmark payoff for the seller, because there is no loss in waiting, but for any $\delta < 1$, waiting would lead the seller to a strictly lower expected utility.

3.2 Optimal Mechanisms

We now show that even without violating assumptions 1)-3), the benchmark allocation rule—and therefore the benchmark profit for the seller—can be achieved. This can be done in a fully-revealing mechanism with $q_1 = t_1 = 0$, as in the benchmark case, but cannot be done with a menu of prices.

Given that we are still looking at fully revealing mechanisms with $q_1 = t_1 = 0$, but we are otherwise keeping all of our assumptions, the seller's problem 2 becomes

$$\max_{q_0, t_0} \int_0^1 \int_0^1 [t_0(\theta, \omega) - \omega q_0(\theta, \omega)] dG(\omega) dF(\theta) \quad (4)$$

¹³As our examples in the introduction show, this is not a credible possibility in many realistic cases.

subject to $q_0 \in [0, 1]$,

$$\int_0^1 [\theta q_0(\theta, \omega) - t_0(\theta, \omega)] dG(\omega) \geq 0, \forall \theta \quad (\text{IRB})$$

$$\int_0^1 [\theta q_0(\theta, \omega) - t_0(\theta, \omega)] dG(\omega) \geq \int_0^1 [\theta q_0(\theta', \omega) - t_0(\theta', \omega)] dG(\omega), \forall \theta, \theta' \quad (\text{ICB})$$

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

Problem 4 has a natural interpretation. The seller asks the buyer to report her type. Since the latter makes this report before observing the seller's report, she faces Bayesian IC and IR constraints. Having observed the buyer's type, and upon learning her own valuation, the seller reports. Since she knows the buyer's type, her IC constraint is in dominant strategies. Thus, the asymmetry in timing in our setting leads to a problem where the two sides' constraints have different informational requirements.¹⁴

The following result describes our optimal mechanism:

Proposition 2. *The solution to problem (4) is a mechanism that implements the benchmark allocation:*

$$q_0(\theta, \omega) = \mathbf{1}[\omega \leq \psi(\theta)],$$

with transfers

$$t_0(\theta, \omega) = \lambda(\theta) + \psi(\theta)\mathbf{1}[\omega \leq \psi(\theta)],$$

where

$$\lambda(\theta) = (\theta - \psi(\theta))G(\psi(\theta)) - \int_0^\theta G(\psi(y))dy,$$

whenever $\theta \geq \psi^{-1}(0)$ and $\lambda(\theta) = 0$ otherwise.

Proof. As is standard, IC and IR for the buyer imply that $Q_0^B(\theta) = \int_0^1 q_0(\theta, \omega)dG(\omega)$ is weakly increasing and that

$$T_0^B(\theta) = \int_0^1 t_0(\theta, \omega)dG(\omega) = \theta Q_0^B(\theta) - \int_0^\theta Q_0^B(y)dy$$

with $T^B(0) = 0$.

¹⁴In certain settings Bayesian and dominant strategy implementation are outcome equivalent. We implicitly impose ex-post balancedness as is standard in bilateral bargaining problems, and it is well known (Manelli and Vincent (2010)) that where ex-post balancedness is required, this outcome equivalence does not necessarily hold.

Now, given $v(\theta, \omega) = t_0(\theta, \omega) - \omega q_0(\theta, \omega)$ then IC for the seller implies that $q_0(\theta, \omega)$ be weakly decreasing in ω for any θ and that

$$t_0(\theta, \omega) = v(\theta, 1) + \int_{\omega}^1 q_0(\theta, x) dx + \omega q_0(\theta, \omega).$$

We denote

$$\lambda(\theta) = v(\theta, 1),$$

then we have

$$T_0^S(\theta) = \int_0^1 t_0(\omega, \theta) dG(\omega) = \lambda(\theta) + \int_0^1 \left[\int_{\omega}^1 q_0(x, \theta) dx + \omega q_0(\omega, \theta) \right] dG(\omega).$$

Suppose $q_0(\omega, \theta) = \mathbf{1}[\omega \leq \psi(\theta)]$. We show below that this allocation is implementable in this problem and since it is the optimal allocation for the relaxed problem without seller constraints, it must also be the optimal allocation here. Recalling that $\psi^+(\theta) = \max\{\psi(\theta), 0\}$, we have

$$\begin{aligned} T_0^S(\theta) &= \lambda(\theta) + \int_0^1 \left[\int_{\omega}^1 \mathbf{1}[x \leq \psi(\theta)] dx + \omega \mathbf{1}[\omega \leq \psi(\theta)] \right] dG(\omega) \\ &= \lambda(\theta) + \int_0^1 \mathbf{1}[\omega \leq \psi(\theta)] G(\omega) d\omega + \int_0^1 \omega \mathbf{1}[\omega \leq \psi(\theta)] dG(\omega) \\ &= \lambda(\theta) + \mathcal{G}(\psi^+(\theta)) + \psi^+(\theta) G(\psi^+(\theta)) - \mathcal{G}(\psi^+(\theta)) \\ &= \lambda(\theta) + \psi^+(\theta) G(\psi^+(\theta)) \end{aligned}$$

while

$$\begin{aligned} T_0^B(\theta) &= \theta Q_0^B(\theta) - \int_0^{\theta} Q_0^B(y) dy \\ &= \theta \int_0^1 \mathbf{1}[\omega \leq \psi(\theta)] dG(\omega) - \int_0^{\theta} \int_0^1 \mathbf{1}[\omega \leq \psi(y)] dG(\omega) dy \\ &= \theta G(\psi^+(\theta)) - \int_0^{\theta} G(\psi^+(y)) dy. \end{aligned}$$

Of course, we need $T_0^S(\theta) = T_0^B(\theta)$ which then implies

$$\lambda(\theta) = (\theta - \psi^+(\theta)) G(\psi^+(\theta)) - \int_0^{\theta} G(\psi^+(y)) dy$$

Given the definition of $\psi^+(\theta)$, this is the $\lambda(\theta)$ function we were looking for. It is immediate to check that $q_0(\omega, \theta) = \mathbf{1}[\omega \leq \psi(\theta)]$ is weakly decreasing in ω for all θ , that $Q_0^B(\theta) =$

$G(\psi^+(\theta))$ is weakly increasing in θ , and that type $\theta = 0$ has expected utility equal to zero. Finally,

$$t_0(\theta, \omega) = \lambda(\theta) + \int_{\omega}^1 \mathbf{1}[x \leq \psi(\theta)] dx + \omega \mathbf{1}[\omega \leq \psi(\theta)] = \lambda(\theta) + \psi(\theta) \mathbf{1}[\omega \leq \psi(\theta)]$$

□

Proposition 2 tells us that the benchmark allocation is implementable even with the additional dominant-strategy IC constraint for the seller, but this requires very specific transfers, which require payments even when trade does not take place. The corollary below collects some properties of these payments:

Corollary 1. *Let $\theta^* = \psi^{-1}(0)$. Then $\lambda(\theta)$ is everywhere continuous and a differentiable function for $\theta > \theta^*$. Further:*

1. $\int_0^1 \lambda(\theta) dF(\theta) = 0$,
2. $\lambda'_{\theta}(\theta^*) > 0$,
3. $\lambda(1) < 0$.

The proof is in Appendix A.2.2. The first property confirms that this mechanism achieves the benchmark expected utility for the seller since the expected utility in this mechanism is

$$\begin{aligned} & \int_0^1 \int_0^1 [t_0(\theta, \omega) - \omega q_0(\theta, \omega)] dG(\omega) dF(\theta) \\ &= \int_0^1 \int_0^1 [\lambda(\theta) + (\psi(\theta) - \omega) \mathbf{1}[\psi(\theta) \geq \omega]] dG(\omega) dF(\theta) \\ &= V^B + \int_0^1 \int_0^1 \lambda(\theta) dG(\omega) dF(\theta) = V^B \end{aligned}$$

In the mechanism, buyer types for which the virtual valuation is non-positive never trade and don't pay any transfers. The second and third properties imply that buyers whose virtual value is just above zero, in equilibrium make a positive payment $\lambda(\theta)$ to the seller and agree to pay a further price $\psi(\theta)$ for the good if the seller agrees to deliver the good once her type is revealed. If the seller decides not to deliver the good, they don't pay this additional price. In equilibrium, the seller will agree to deliver the good if $\omega \leq \psi(\theta)$. Types who are sufficiently close to 1, agree to the same deal, *but* their $\lambda(\theta)$ is negative so that for these types, $\lambda(\theta)$ represents a subsidy they receive from the seller regardless of

whether the latter delivers or not.¹⁵

To get an intuition for this mechanism, note that the transfers we had in the benchmark mechanism are a function of the seller's type. In a setting where the seller cannot commit to truthfully revealing her type, such a mechanism will not be implementable as the seller will sometimes have the incentive to misreport her type to achieve a better price. So, the mechanism makes the transfers a function of the buyer's type instead. In particular, the seller proposes to the buyer that he chooses a price p he is willing to pay for the good, with the proviso that the seller will only deliver the good if her type turns out to be no greater than p . This is clearly incentive compatible for the seller. In addition, if each buyer type θ chooses $p(\theta) = \psi(\theta)$ whenever $\psi(\theta) \geq 0$ and proposes no price otherwise, then we would implement the benchmark allocation. The problem is that $p(\theta) = \psi(\theta)$ is not incentive compatible for the buyer. For example, since $\psi(1) = 1$, type $\theta = 1$ would commit to pay 1 for a certain purchase of a good that is worth 1 to him, resulting in a total utility of zero. He is clearly better off choosing a price $p < 1$ for an expected utility of $(1 - p)G(p) > 0$.

If we allowed buyer types to choose prices according to an incentive-compatible price schedule $p(\theta) \neq \psi(\theta)$, but still insisted that the seller would deliver the good only if $\psi(\theta) \geq \omega$ we'd by definition achieve incentive compatibility for the buyer, and this would respect the benchmark allocation. But it's easy to see that we'd lose again incentive compatibility for the seller. This is because whenever $p(\theta) > \omega > \psi(\theta)$ the seller would have an incentive to underreport her type and make the trade happen and, conversely, whenever $\psi(\theta) > \omega > p(\theta)$, the seller would have an incentive to overreport her type and stop the trade.

What the optimal mechanism does is to fix the problem by making $p(\theta) = \psi(\theta)$ so that we have incentive compatibility for the seller and implementation of the benchmark allocation, but also adds a payment $\lambda(\theta)$ that is independent of the allocation but changes the total net payments that the buyer has to pay in order to achieve incentive compatibility for him. Going back to the $\theta = 1$ example, the mechanism asks this buyer type to pay a price of 1 in exchange for a certain delivery of a good that he values 1, but as the third property of Corollary 1 emphasizes, he also get a subsidy from the seller to compensate. In equilibrium, therefore, type 1 gets no utility from trade as such, but her informational rents are given by this subsidy. For other types, making it incentive compatible for a certain type to choose a price equal to his virtual valuation requires a smaller subsidy or even the payment of a fee.¹⁶

¹⁵When F and G are uniform distributions, we can show that there exists a θ^{**} such that $\lambda(\theta) > 0$ for all $\theta \in (\theta^*, \theta^{**})$ and $\lambda(\theta) < 0$ for all $\theta \in (\theta^{**}, 1]$ but for generic distribution, we can only characterize $\lambda(\theta)$'s behavior in the neighborhoods of θ^* and 1.

¹⁶In other words, when $\lambda(\theta) > 0$ the seller can effectively charge more than the virtual valuation and

The discussion above suggests an immediate, indirect implementation: the seller offers a menu of payment pairs $\{(p, \Lambda(p))\}_{p \in [0,1]}$, where $\Lambda(p) = \lambda(\psi^{-1}(p))$. Figure 1 illustrates this menu in the case of uniform F and G distributions. Each option in the menu is an *at-will* contract plus a fixed transfer. If the buyer chooses a particular price p , the buyer commits to paying that price, but the seller retains the right to not deliver upon learning her type. If she does not deliver, then p is not paid. The buyer also accepts an up-front fee or subsidy equal to $\Lambda(p)$ which is paid whether trade occurs or not. This is therefore a menu of *two-part tariffs*, with the unusual feature that for some components of the menu, the fixed component is a fee while for others it is a subsidy.

The fact that this seems unusual suggests that there may be other factors that make it unpractical. It is simple to observe, for instance, that those buyer types who pay a fee may end up not getting the good which immediately implies that this mechanism is not ex-post individually rational for the buyer; this may be a concern in some environments. It may also be that the mechanism is too complex and relies excessively on fine details of the problem: the natural implementation of the benchmark mechanism only requires buyers to think about the most they are willing to pay, while this mechanism asks them to choose from a menu of complicated pricing options. For example, it may be that $p'' > p' > p$ but that $\Lambda(p') > \Lambda(p) > \Lambda(p'')$ and so a buyer would find it difficult to compare the three options.¹⁷ Figure 1 illustrates this comparison.




FIGURE 1: OPTIMAL MENU OF TWO-PART TARIFFS IN THE CASE OF UNIFORM DISTRIBUTIONS.

still preserve incentive compatibility with this buyer type, while when $\lambda(\theta) < 0$, the virtual valuation is too high a price to be incentive compatible and the seller needs to subsidize the buyer.

¹⁷Recall that we know that $\lambda(\theta)$ is positive and increasing in a neighborhood of θ^* and negative in a neighborhood of $\theta = 1$ so it cannot be monotone.

For this reason, it is worthwhile to look at a simpler class of mechanisms, one where buyers only make a payment if trade takes place. We study such *price* mechanisms in the next section.

4 Price Mechanisms

4.1 Definitions

In the previous section, we introduced the idea that a simple mechanism would have the property that buyers only make a payment if trade takes place. We begin by defining this notion formally. Next, we show that this property is not enough. In particular, we show in Example 1 that one can construct a mechanism that satisfies this property and “almost” implements the benchmark allocation, but this mechanism involves randomization and is significantly less realistic than even two-part tariffs.

Definition 2. *A reduced-form mechanism \mathcal{M}^R satisfies the no payment without trade (NPWT) property if for all $b \in \mathcal{B}, \theta \in \Theta$, and $\omega \in \Omega$,*

$$q_0(b, \omega) = 0 \Rightarrow t_0(b, \omega) = 0 \text{ and } q_1(b, \omega, \theta) = 0 \Rightarrow t_1(b, \omega, \theta) = 0.$$

While the (NPWT) property eliminates the possibility of a literal use of two-part tariffs, it still allows for mechanisms where introducing randomization through the garbling device can approximately deliver the benchmark outcome.¹⁸

Example 1. Consider a mechanism where the garbling device (\mathcal{B}, μ) is such that $\mathcal{B} = \Theta \cup \emptyset$, $\mu(\emptyset|\theta) = \varepsilon$, $\mu(\theta'|\theta) = 1 - \varepsilon$ when $\theta' = \theta$, and $\mu(\theta'|\theta) = 0$ when $\theta' \in \Theta \setminus \theta$. Let the allocation rule be $q_0(\emptyset, \omega) = 1$ and $q_0(\theta, \omega) = \mathbf{1}[\omega \leq \psi(\theta)]$ for all $\theta \in \Theta$, and $q_1(b, \omega, \theta) \equiv 0$. In other words, the message \emptyset says “deliver the object with certainty,” and other messages lead to the benchmark allocation rule. Let $t_1(b, \omega, \theta) \equiv 0$, while $t_0(\theta, \omega) = \max\{0, \psi(\theta)\}$ and $t_0(\emptyset, \omega) = \lambda^\varepsilon(\theta)/\varepsilon$, where $\lambda^\varepsilon(\theta)$ is computed so that incentive compatibility for the buyer is satisfied, i.e.,

$$\lambda^\varepsilon(\theta) = (1 - \varepsilon) \left(\theta G(\psi(\theta)) - \int_0^\theta G(\psi(x)) dx - \int_0^{\psi(\theta)} \psi(\theta) dG(\omega) \right).$$

This mechanism satisfies all constraints and the (NPWT) property by construction. Moreover, as $\varepsilon \rightarrow 0$, this mechanism converges to the benchmark mechanism and achieves profit

¹⁸A similar example can be constructed with a garbling device that is deterministic, but the allocation and transfer rules are stochastic.

arbitrarily close to it.

The mechanism achieves this by introducing a probability ε that trade takes place regardless of the underlying types. It is easy to see that as this probability approaches zero, the payment required by the buyer $\lambda^\varepsilon(\theta)/\varepsilon$ becomes arbitrarily large. This is neither simpler nor more realistic than two-part tariffs. Therefore, we focus on a class of simpler deterministic mechanisms that satisfy the (NPWT) property.

Definition 3. *A reduced-form mechanism \mathcal{M}^R is deterministic if*

1. *It has a deterministic garbling device, i.e., there exists a function $m : \Theta \rightarrow \mathcal{B}$ such that $\mu(b|\theta) = \mathbf{1}[b = m(\theta)]$ for all $b \in \mathcal{B}$, $\theta \in \Theta$, and*
2. *It has a deterministic allocation rule, i.e., $q_0(b, \omega) \in \{0, 1\}$ and $q_1(b, \omega, \theta) \in \{0, 1\}$ for all $b \in \mathcal{B}$, $\theta \in \Theta$, $\omega \in \Omega$.*

Note that with the deterministic garbling device, the first argument of the q_1 and t_1 functions is redundant. Therefore, with slight abuse of notation, we replace $q_1(m(\theta), \omega, \theta)$ with $q_1(\theta, \omega)$ and $t_1(m(\theta), \omega, \theta)$ with $t_1(\theta, \omega)$.

Definition 4. *A price mechanism is a deterministic reduced-form mechanism \mathcal{M}^R that satisfies the (NPWT) property. We denote with \mathcal{M}^P the class of such mechanisms.*

4.2 Static Price Mechanisms

Static price mechanisms are simple and realistic types of mechanisms that provide useful benchmarks. By a static mechanism, we mean that trade only happens at a specific time (if at all), so there is no room for dynamic pricing. Specifically, we use the following definitions:

Definition 5. *A price mechanism is*

1. *An ex-post price mechanism if $q_0(b, \omega) = t_0(b, \omega) = 0$ for all $b \in \mathcal{B}$, $\omega \in \Omega$. We denote with $\mathcal{M}^{EPP} \subset \mathcal{M}^P$ the class of ex-post price mechanisms.*
2. *An ex-ante price mechanism if $q_1(\theta, \omega) = t_1(\theta, \omega) = 0$ for all $b \in \mathcal{B}$, $\theta \in \Theta$, $\omega \in \Omega$. We denote with $\mathcal{M}^{EAP} \subset \mathcal{M}^P$ the class of ex-ante price mechanisms.*
3. *A static price mechanism if it is either an ex-post or ex-ante price mechanism.*
4. *A dynamic price mechanism if it is a price mechanism that is not static.*

Intuitively, ex-post price mechanisms are those where trade can only be agreed at time 1, after the cost uncertainty has already been realized, whereas ex-ante price mechanisms are those where trade can only be agreed at time 0, before the cost information is revealed. In this subsection, we characterize the optimal ex-post and ex-ante price mechanisms, and the one of the two that guarantees higher profit is the optimal static price mechanism.

4.2.1 Ex-Post Price Mechanisms

We first consider ex-post price mechanisms. In this class of mechanisms, the seller's maximization problem is

$$\max_{\mathcal{M}^{EPP}} \delta \int_0^1 \int_0^1 [t_1(\theta, \omega) - \omega q_1(\theta, \omega)] dF(\theta) dG(\omega) \quad (5)$$

subject to

$$\begin{aligned} \delta \int_0^1 [\theta q_1(\theta, \omega) - t_1(\theta, \omega)] dG(\omega) &\geq 0, \forall \theta \\ \delta \int_0^1 [\theta q_1(\theta, \omega) - t_1(\theta, \omega)] dG(\omega) &\geq \delta \int_0^1 [\theta q_1(\theta, \omega) - t_1(\theta', \omega)] dG(\omega), \forall \theta, \theta' \\ \delta \int_0^1 [t_1(\theta, \omega) - \omega q_1(\theta, \omega)] dF(\theta) &\geq \delta \int_0^1 [t_1(\theta, \omega') - \omega q_1(\theta, \omega')] dF(\theta), \forall \omega, \omega' \end{aligned}$$

This is exactly the same problem as (3) in section 3.1, where the seller commits to ignoring the report by the buyer, with one crucial difference—everything happens at time 1 instead of time 0 and therefore all payoffs are multiplied by δ . The solution to that problem was clearly a price mechanism. Therefore the optimal mechanism within this class, which we call ex-post optimal (EPO) price mechanism, is such that $q_1(\theta, \omega) = \mathbf{1}[\psi(\theta) \geq \omega]$, and $t_1(\theta, \omega) = \psi^{-1}(\omega) \mathbf{1}[\psi(\theta) \geq \omega]$. The only difference is that the profit is discounted by $\delta \in (0, 1)$, i.e., it gives profit $V^{EPO} = \delta V^B$.

4.2.2 Ex-Ante Price Mechanisms

Let us now consider the class of ex-ante price mechanisms. The seller's maximization problem is now

$$\max_{\mathcal{M}^{EAP}} \int_0^1 \int_0^1 [t_0(m(\theta), \omega) - \omega q_0(m(\theta), \omega)] dG(\omega) dF(\theta) \quad (6)$$

$$\int_0^1 [\theta q_0(m(\theta), \omega) - t_0(m(\theta), \omega)] dG(\omega) \geq 0, \forall \theta$$

$$\int_0^1 [\theta q_0(m(\theta), \omega) - t_0(m(\theta), \omega)] dG(\omega) \geq \int_0^1 [\theta q_0(m(\theta'), \omega) - t_0(m(\theta'), \omega)] dG(\omega), \forall \theta, \theta'$$

$$t_0(b, \omega) - \omega q_0(b, \omega) \geq t_0(b, \omega') - \omega q_0(b, \omega'), \forall b, \omega, \omega'$$

In problem (5), the seller only had prior information and, as discussed in section 3.1, having no updated information slackens completely the seller's incentive compatibility problem. When we looked at general mechanisms, we also saw that if we allowed for two-part tariffs, then a fully-revealing mechanisms could still implement the benchmark outcome. Here, we are only looking at ex-ante price mechanisms so it is not obvious whether it is optimal for the seller to allow her future self to update her beliefs about the buyer's type or not.

The answer is in the lemma below:

Lemma 2. *The solution to problem (6) and the solution to problem (4)—when this is limited to mechanisms in \mathcal{M}^{EAP} —coincide.*

The proof is in Appendix A.3.1.

Thus, it turns out that within the class of ex-ante price mechanisms, fully revealing mechanisms are without loss of generality. This result follows entirely from the assumption that the garbling device is deterministic. It provides a helpful simplification, as it eliminates the information design aspect of the problem.

Menus of offers. It is useful to define a particular class of simple mechanisms. As discussed, an at-will offer is a price offer that gives the seller the opportunity to walk away if the cost is too high. In this particular context, it is a price \hat{p}_0 such that if B accepts this, then the trade happens at price \hat{p}_0 if and only if S learns that the cost $\omega \leq \hat{p}_0$. Another type of offer is a guaranteed delivery price, where the seller delivers the object no matter the cost. A menu of offers is a collection of at-will and guaranteed delivery offers. Specifically, each menu of offers is described by a function $p_0 : \Theta \rightarrow \mathbb{R}$ and a constant $P_0 \geq 0$, such that each $p_0(\theta)$ is an at-will offer, and P_0 is the guaranteed delivery price.¹⁹ It is straightforward to verify that menus of offers can be implemented as ex-ante price mechanisms. The following establishes this optimality—a menu of offers that satisfies particular conditions is indeed a solution to problem (6).

Proposition 3. *Problem (6) is solved by a menu of offers, described by $p_0(\theta) \in [0, \theta]$ and $P_0 \geq 0$, such that:*

¹⁹The assumption that there is only one guaranteed delivery price is without loss of generality, as the buyer would choose only the lowest price.

1. Function $p_0(\theta)$ is weakly increasing and satisfies the following equation for each $\theta \in \Theta$,

$$(\theta - p_0(\theta)) G(p_0(\theta)) = \int_0^\theta G(p_0(x)) dx. \quad (7)$$

2. There exists a marginal buyer type $\tilde{\theta} < 1$ so that buyers with types $\theta > \tilde{\theta}$ take the guaranteed delivery option at price P_0 .
3. The guaranteed delivery price satisfies

$$P_0 = p_0(\tilde{\theta})G(p_0(\tilde{\theta})) + \tilde{\theta}[1 - G(p_0(\tilde{\theta}))]. \quad (8)$$

The proof is in appendix A.3.2.

While the previous result does not fully determine the optimal mechanism, it puts a specific structure on it and within this class the problem can be solved numerically for specific distributions F and G . The following corollary describes some properties of any weakly increasing $p_0(\theta)$ function that satisfies constraint (7).

Corollary 2. $p_0(\theta)$ from proposition 3 satisfies the following properties:

1. If $p_0(\theta)$ is differentiable at θ then either $p'_0(\theta) = 0$ or $p'_0(\theta) = \phi^{-1}(\theta)$, where $\phi(x)$ is the virtual cost.
2. If there is a discontinuous jump at θ from \underline{p} to $\bar{p} > \underline{p}$, then $\bar{p} < \phi^{-1}(\theta) < \underline{p}$ and

$$\frac{\bar{p}G(\bar{p}) - \underline{p}G(\underline{p})}{G(\bar{p}) - G(\underline{p})} = \theta.$$

It is interesting to compare this with the case where both buyer and seller face Bayesian incentive compatibility or dominant strategy incentive compatibility constraints. In the case of Bayesian incentive compatibility, as previously discussed, the solution was originally provided by Williams (1987) and it is the benchmark mechanism, where $q_0(\theta, \omega) = \mathbf{1}[\psi(\theta) \geq \omega]$ whereas here we have $q_0(\theta, \omega) = \mathbf{1}[p_0(\theta) \geq \omega]$ for all $\theta < \tilde{\theta}$ and $q_0(\theta, \omega) = 1$ for all $\theta \geq \tilde{\theta}$. Therefore it is clear that the ex-ante optimal price mechanism cannot achieve the benchmark profits. The following corollary formalizes this observation.

Corollary 3. *The profit from the optimal ex-ante price mechanism is strictly lower than V^B .*

The proof is in Appendix A.3.3.

On the other hand, when both buyer and seller face dominant-strategy incentive compatibility constraints, Hagerty and Rogerson (1987) showed that the optimal solution is

an at-will offer, but the seller offers a unique such price that all buyer types must either accept or reject. The result above shows such a mechanism cannot be an optimal ex-ante price mechanism because it does not feature a guaranteed delivery offer.

It is natural to ask, whether the optimal ex-ante price mechanism can sometimes be a single guaranteed-delivery offer with no at-will price $p_0(\theta)$. The answer is no.

Corollary 4. *The profit from the optimal ex-ante price mechanism is strictly higher than the profit obtainable from a mechanism in which the buyer is offered only a guaranteed delivery price P_0 .*

The proof is in Appendix A.3.4. In it, we show that a single guaranteed delivery offer would give the seller a profit that would be strictly lower than the profit from a single at-will offer (with an optimally chosen price). In turn, Proposition 3 shows that such a single at-will offer is strictly dominated by a menu of offers, which includes a guaranteed delivery offer. Therefore, the optimal ex-ante price mechanism is always a menu of offers, which features at least one at-will offer and one guaranteed delivery offer.

4.2.3 Optimal Static Mechanisms

Combining the results from the previous two subsections, we can now state the following corollary from previous results.

Corollary 5. *There exists $\bar{\delta} \in (0, 1)$ such that the optimal static price mechanism is*

1. *the EAO mechanism with profit $V^{EAO} < V^B$ when $\delta \leq \bar{\delta}$,*
2. *the EPO mechanism with profit $V^{EPO} = \delta V^B$ when $\delta \geq \bar{\delta}$.*

The proof follows directly from the results above.

The following figure summarizes the results from this section when both F and G are uniform distributions. The green line depicts the benchmark profit V^B , which is the highest possible profit when the seller can use arbitrary mechanisms, including two-part tariffs. The increasing black line shows $V^{EPO} = \delta V^B$, which is the profit from the optimal ex-post mechanism (which turns out to be a price mechanism). Clearly, as $\delta \rightarrow 1$, it converges to V^B , but for any $\delta < 1$, the profit is strictly lower. The blue solid line shows the optimal ex-ante price mechanism discussed above. It is a combination of at-will and guaranteed delivery offers. The profit from ex-ante mechanisms does not depend on the discount factor δ . Therefore, naturally, if the discount factor is low, $V^{EAO} > V^{EPO}$, and if the discount factor is close to one, then $V^{EAO} < V^{EPO}$. The maximum of these two profits is the profit from the optimal static mechanism, $V^{static} = \max\{V^{EAO}, V^{EPO}\}$. Finally, the figure also depicts two other natural mechanisms. The purple line shows the

highest profit achievable by a single at-will offer, which turns out to be the best the seller can do using a menu of at-will offers. Proposition 3 shows that this is strictly below V^{EAO} . The grey line shows the highest profit from a single guaranteed delivery offer. In Corollary 4 and its proof, we show that this is strictly below the profit achievable only with at-will offers, and both are strictly below V^{EAO} .

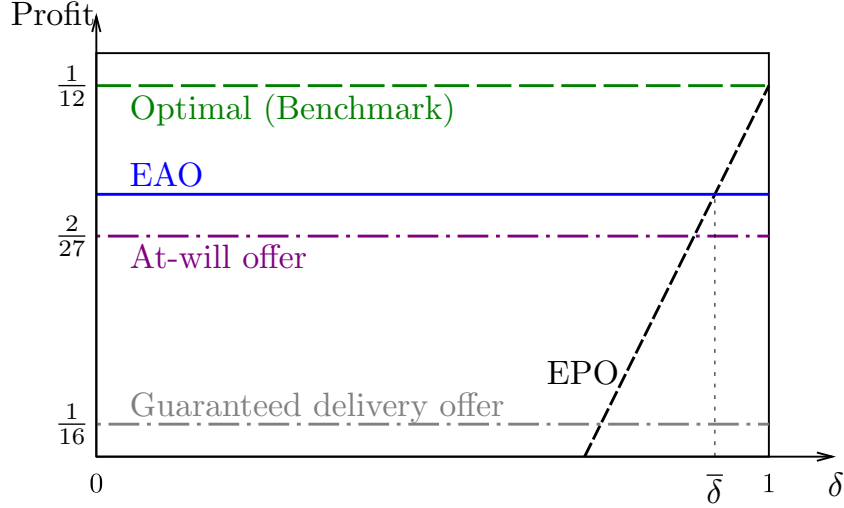


FIGURE 2: OPTIMAL STATIC PRICE MECHANISMS FOR UNIFORM DISTRIBUTIONS.

4.3 Dynamic Price Mechanisms

4.3.1 D-mechanisms

In this section, we consider dynamic price mechanisms. We show that for a sufficiently large discount factor δ , the optimal price mechanism is a dynamic price mechanism. To prove this, it is sufficient to find a particular dynamic price mechanism that gives strictly higher profit than any of the static price mechanisms, meaning strictly higher profit than $V^{static} = \max\{V^{EAO}, V^{EPO}\}$.

Let us define the following simple dynamic mechanism.

Definition 6. *D-mechanism with reserve r is a sequence of guaranteed delivery offers: $(p_0, p_1(\omega))$, such that*

1. *If the buyer accepts the first offer p_0 , then trade happens with certainty at time 0 at price p_0 .*
2. *If the buyer rejects this offer, the seller makes a new offer $p_1(\omega) \geq r$, which is again a guaranteed delivery offer, but for time 1.*

3. If the buyer rejects the second offer as well, then trade does not happen.

In terms of a reduced-form mechanism, this mechanism has a garbling device with binary signal ($\mathcal{B} = \{0, 1\}$), such that $b \in \{0, 1\}$ tells the seller whether the buyer took the first offer or not. For example, if $m(\theta) = \mathbf{1}[\theta \geq \bar{\theta}]$ for some threshold $\bar{\theta}$, then $q_0(b, \omega) = b$, $t_0(b, \omega) = bp_0$, $q_1(\theta, \omega) = \mathbf{1}[p_1(\omega) \leq \theta < \bar{\theta}]$, and $t_1(\theta, \omega) = p_1(\omega)\mathbf{1}[p_1(\omega) \leq \theta < \bar{\theta}]$.

The following proposition shows that such mechanism has its merit.

Proposition 4. *Profit from the D-mechanism with $r = 0$ is strictly higher than the profit from the optimal ex-post price mechanism for any $\delta < 1$.*

The proof is in Appendix A.3.5

Using this proposition, we can immediately conclude for sufficiently high δ , the optimal price mechanism is dynamic.

Corollary 6. *With $\delta \geq \bar{\delta}$, the maximal profit from a D-mechanism with reserve r is strictly higher than the profit from the optimal static price mechanism.*

Proof. Let $V^D(r)$ be the profit from the D-mechanism with any reserve r , and let $r^D = \operatorname{argmax} V^D(r)$. Then clearly the maximal profit from D-mechanism with reserve $V^D(r^D) \geq V^D(0)$. Proposition 4 shows that $V^D(0) > V^{EPO}$ for all δ , and Corollary 5 shows that $V^{static} = V^{EPO}$ for all $\delta \geq \bar{\delta}$. Therefore, for all $\delta \geq \bar{\delta}$, $V^D(r^D) \geq V^D(0) > V^{EPO} = V^{static}$. \square

The following figure illustrates these results in the case of uniform F and G distributions.²⁰

²⁰Calculations for $V^D(r^D)$ are available upon request.




FIGURE 3: D-MECHANISMS COMPARED WITH STATIC PRICE MECHANISMS FOR UNIFORM DISTRIBUTIONS.

4.3.2 Additional Requirements for Price Mechanisms

We end this section with an example that illustrates that even the requirements we imposed on price mechanism, i.e., mechanisms that are deterministic and satisfy (NPWT) property, are not sufficient to guarantee that the seller cannot approximate a two-part tariff closely if we allow for dynamic price mechanisms.

Example 2. Consider the following fully revealing price mechanism. Let $q_0(\theta, \omega) = \mathbf{1}[\psi(\theta) \geq \omega]$, i.e., the ex-ante allocation rules are the optimal one, and let $q_1(\theta, \omega) = \mathbf{1}[0 \leq \psi(\theta) < \omega]$, i.e., if the buyer has a positive virtual surplus and was not served at time 0, she will be served at time 1 with certainty. This is certainly a suboptimal allocation, but when δ is small, the distortion from serving the buyer with certainty at time 1 is small. Now, to make this allocation rule incentive compatible for both parties, we set $t_0(\theta, \omega) = q_0(\theta, \omega)p_0(\theta)$ and $t_1(\theta, \omega) = q_1(\theta, \omega)p_1(\theta)$, where $p_0(\theta)$ and $p_1(\theta)$ are prices determined below. This mechanism is a price mechanism by construction.

We begin with the seller's incentives. Obviously when $\psi(\theta) < 0$, there is no trade, so no transfers. Let's fix $\psi(\theta) \geq 0$. Seller's profit when announcing ω' when the true values are θ, ω is now

$$v(\omega'|\theta, \omega) = \begin{cases} p_0(\theta) - \omega & \text{if } \omega' \leq \psi(\theta) \\ \delta[p_1(\theta) - \omega] & \text{if } \omega' > \psi(\theta) \end{cases}$$

The seller type $\omega = \psi(\theta)$ must be indifferent between the two options, which gives us

$$p_0(\theta) = \delta p_1(\theta) + (1 - \delta)\psi(\theta).$$

Now we see that $p_1(\theta)$ plays a role quite similar to that of the fixed payment in the two-part tariff. The transfer includes $\delta p_1(\theta)$ no matter what the seller reports, but any report inducing the early trade means additional payment $(1 - \delta)\psi(\theta)$. As such report comes with additional cost $(1 - \delta)\omega$ (early trade instead of late), it is incentive compatible for the seller to report the type truthfully. To determine $p_1(\theta)$ we look at the buyer's incentives. Notice that the expected allocation rule for buyer who reports θ' such that $\psi(\theta) \geq 0$ is

$$Q(\theta') = \int_0^1 [q_0(\theta', \omega) + \delta q_1(\theta', \omega)] dG(\omega) = G(\psi(\theta')) + \delta[1 - G(\psi(\theta'))]$$

Specifically, the expected transfers need to satisfy

$$T(\theta) = \theta Q(\theta) - \int_{\theta^*}^{\theta} Q(x) dx.$$

On the other hand, we can use p_0 and p_1 to compute T to get

$$T(\theta) = \int_0^{\psi(\theta)} p_0(\theta) dG(\omega) + \int_{\psi(\theta)}^1 \delta p_1(\theta) dG(\omega) = \delta p_1(\theta) + (1 - \delta)\psi(\theta)G(\psi(\theta)).$$

Combining the two expressions for $T(\theta)$, we get

$$\begin{aligned} \delta p_1(\theta) &= \theta (G(\psi(\theta)) + \delta[1 - G(\psi(\theta))]) - (1 - \delta)\psi(\theta)G(\psi(\theta)) \\ &\quad - \int_{\theta^*}^{\theta} (G(\psi(x)) + \delta[1 - G(\psi(x))]) dx. \end{aligned}$$

Finally, note that

$$\lim_{\delta \rightarrow 0} \delta p_1(\theta) = \theta G(\psi(\theta)) - \psi(\theta)G(\psi(\theta)) - \int_{\theta^*}^{\theta} G(\psi(x)) dx = \lambda(\theta),$$

where $\lambda(\theta)$ is the fixed component of the two-part tariff in the optimal (general) mechanism. Similarly, $\lim_{\delta \rightarrow 0} p_0(\theta) = \psi(\theta)$. Therefore, as $\delta \rightarrow 0$, the mechanism gets closer and closer to the optimal (general) mechanism.

This example illustrates two facts. First, the class of price mechanisms that we are studying is rich enough that the seller has many levers to achieve profits close to the benchmark profit, even when we put some natural restrictions on the mechanisms. Second, as in Example 1, such mechanisms are typically not more realistic or simpler than the original two-part tariffs, because they require, when δ becomes arbitrarily small, unboundedly

large transfers by the buyer.

To avoid mechanisms such as these, that rely on the dynamic nature of the problem, but do so in an unrealistic way, it seems natural to restrict the class of mechanisms further. We call these ex-post individually rational price mechanisms. Informally, they require not only that transactions happens deterministically and that the money changes hands only when the good is delivered, but also that at the moment the transaction happens, the buyer must not be unhappy with the transaction.

Definition 7. *An ex-post individually rational price mechanism is a price mechanism such that $t_0(b, \omega) \leq \theta$ and $t_1(\theta, \omega) \leq \theta$ for all b, θ, ω . We denote with $\mathcal{M}^{EPIRP} \subset \mathcal{M}^P$ the class of ex-post individually rational price mechanisms.*

5 Extensions

In this section, we consider two natural questions that arise from our setting. Up to this point, we've assumed that the seller learns her type with delay, but there is no cost of learning. We call this the case of *free learning*. In some settings, however, it may make more sense to assume that to learn her type, the seller might have to pay a cost. For example, if we interpret the seller's type as a production cost, then the seller might need to make an investment in research and development to learn what such production cost is. In these cases, the seller cannot provide the good unless that investment is made. We call this the *necessary learning* case. In other situations, the good might be already available to the seller, but she may not be able to attribute it a precise valuation, so that learning the cost requires hiring an expert or expanding some effort to learn her type. In these cases, the seller can still provide the good even if she decides not to spend to learn her valuation. We call this the *optional learning* case.

We study both these costly learning models in section 5.1. Our analysis focuses again on optimal general mechanisms and shows, as is intuitive, that if the learning costs are too high, the seller either does not offer a sale (in the case of necessary learning) or offers a price with guaranteed delivery because she chooses not to learn her type (in the case of optional learning). If, however, these costs are low enough, then the optimal mechanisms still feature the two-part tariff property of the case where learning is free. This implies that our results are qualitatively robust to the introduction of learning costs.

An important difference between the two cases is that in the case of necessary learning (as when learning is free), the seller only offers prices that are contingent on her type realization: whatever price the buyer accepts, trade is not guaranteed because this price needs to be above the seller's type. In the case of optional learning, the seller decides to

learn her type only when the buyer type is not too high or too low and offers these buyer types a price contingent on her type realization. However, with optional learning, when the seller learns that the buyer's type is sufficiently high, she will prefer to offer a price with guaranteed delivery (because her type is likely to be below this price) over learning her cost.

In subsection 5.2 we imagine that the mechanism designer's objective is ex-post efficiency rather than maximizing the seller's profit. As is well-known, Myerson and Satterthwaite (1983) show that ex-post efficiency cannot be guaranteed when both buyer and seller face Bayesian incentive compatibility constraints and interim individual rationality constraints. In our setting, the natural counterpart is to ask the same question when the buyer faces Bayesian incentive compatibility and interim individual rationality constraints while the seller faces dominant strategy incentive compatibility constraints. In addition, since we are taking the view of the seller as an agent rather than the principal, we impose ex-ante individual rationality for her or, at least, that the seller wishes to participate upon learning the buyer's type. Our result shows that, contrary to the Myerson and Satterthwaite (1983), a mechanism that guarantees ex-post efficiency *does exist*. This suggests that the slackening of the individual rationality constraint for the seller more than compensates for the hardening of her incentive compatibility constraint. Further, the mechanism that implements the ex-post efficient allocation also relies on two-part tariffs, except that now the buyer always receives a subsidy from the seller.²¹

5.1 Costly Learning

As discussed above, we distinguish between the case where in order to deliver the good, the seller must pay a cost to learn ω from the case where the good can be delivered even if the seller does not pay the cost to learn ω . We begin with the former case and again work throughout with fully-revealing mechanisms.²²

5.1.1 Necessary learning

We follow the approach used in subsections 3.1 and 3.2, of characterizing first the benchmark allocation under the assumption that the seller reveals her type truthfully (if she decided to pay to learn her cost) and then show that this allocation can be also obtained in a setting where the seller cannot commit to tell the truth.

²¹If, in our setting, we were maximizing the buyer's expected utility, then the benchmark allocation could be implemented *without* a two-part tariff.

²²Clearly the presence of learning costs, does not change the argument for working with these.

Benchmark mechanism. We assume that, in order to learn ω , S must pay a cost $c > 0$. This decision is made *after* the seller observes the buyer's valuation. We assume this particular timing because it seems realistic to allow the seller to decide not to learn her cost if she learns that with a particular buyer, a trade won't take place regardless of the seller's type realization. The proposition below characterizes the benchmark mechanism:

Proposition 5. *In the benchmark case of necessary learning with a cost $c > 0$, the optimal allocation is:*

$$q_0(\theta, \omega) = \mathbf{1}[c \leq \mathcal{G}(1)] \cdot \mathbf{1}[\max(\omega, \mathcal{G}^{-1}(c)) \leq \psi(\theta)]$$

and the respective transfers are

$$q_0(\theta, \omega) = \mathbf{1}[c \leq \mathcal{G}(1)] \cdot \mathbf{1}[\max(\omega, \mathcal{G}^{-1}(c)) \leq \psi(\theta)] \cdot \psi^{-1}(\max(\omega, \mathcal{G}^{-1}(c))).$$

The seller payoff is zero if $c > \mathcal{G}(1)$, while if $c \leq \mathcal{G}(1)$ we have

$$V^{BNC}(c) = \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^1 \left[-c + \int_0^{\psi(\theta)} (\psi(\theta) - \omega) dG(\omega) \right] dF(\theta).$$

The proof is in Appendix A.4.1.

The case where $c \geq \mathcal{G}(1)$ is very intuitive: if the cost of learning is too high for any possible type of buyer, the seller does not expect her gains from trade to cover the cost of learning and so, from an ex-ante perspective, it is optimal to never trade. If $c < \mathcal{G}(1)$, then some trade is possible, but the menu of prices is different from the benchmark where $c = 0$. In particular, she commits to selling at $\psi^{-1}(\omega)$ as long as her type turns out to be no smaller than $\mathcal{G}^{-1}(c)$. If her type is smaller than $\mathcal{G}^{-1}(c)$, then the price is fixed at $\psi^{-1}(\mathcal{G}^{-1}(c))$ because the seller must also take the cost of learning into account. Since the decision on whether to pay c or not is made after the seller learns the buyer type, the price schedule the seller chooses is the same as the one she chooses when $c = 0$, but it is truncated at $\psi^{-1}(\mathcal{G}^{-1}(c))$. Finally, since both $\psi^{-1}(x)$ and $\mathcal{G}^{-1}(c)$, are strictly increasing functions with $\psi^{-1}(1) = 1$ and $\mathcal{G}^{-1}(0) = 0$, it is easy to see that $V^{BNC}(c)$ is a decreasing function of c with $V^{BNC}(0) = V^B$ and $V^{BNC}(\mathcal{G}(1)) = 0$.²³

Optimal Mechanism. If we now move to the case where the seller cannot commit to reveal her type truthfully, the first observation is that if $c > \mathcal{G}(1)$, then in the benchmark case, it is optimal for the seller not to trade at all, and this clearly remains the case here. If $c \leq \mathcal{G}(1)$, on the other hand, the analogy with the case where learning is free that we

²³Note that trade occurs if $\theta \geq \max(\mathcal{G}^{-1}(c), \omega)$ while c is paid if $\theta \geq \mathcal{G}^{-1}(c)$.

noted in the benchmark case continues, and the optimal mechanism is a two-part tariff, where the set of buyer types the seller wishes to potentially trade with is truncated below, and the λ function is adjusted accordingly:

Proposition 6. *Suppose $c \leq \mathcal{G}(1)$. In the case of necessary learning when the seller cannot commit to truthfully revealing her type, there is an optimal mechanism that implements the benchmark allocation:*

$$q_0(\theta, \omega) = \mathbf{1}[\max(\omega, \mathcal{G}^{-1}(c)) \leq \psi(\theta)]$$

with transfers

$$t_0(\theta, \omega) = \widehat{\lambda}(\theta) + \psi(\theta)\mathbf{1}[\max(\omega, \mathcal{G}^{-1}(c)) \leq \psi(\theta)],$$

where

$$\widehat{\lambda}(\theta) = (\theta - \psi(\theta))G(\psi(\theta)) - \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^{\theta} G(\psi(y))dy,$$

whenever $\theta \geq \psi^{-1}(\mathcal{G}^{-1}(c))$ and $\widehat{\lambda}(\theta) = 0$ otherwise.

This result is proven following the same steps as in the proof of proposition 2. Following such steps one just needs to be clear that $\widehat{\lambda}(\theta) \neq 0$ only applies for $\theta \geq \psi^{-1}(\mathcal{G}^{-1}(c))$. Finally, that $\int_0^1 \widehat{\lambda}(\theta)dF(\theta) = 0$ can be shown using the procedure in the proof of corollary 1.

The mechanism, as in that setting, asks each buyer type for which trade is possible to pay a price equal to their virtual valuation, but since $\psi^{-1}(\mathcal{G}^{-1}(c)) > \psi^{-1}(0)$, the set of buyer types that is allowed to trade is limited from those that had a positive virtual valuation to those whose virtual valuation is higher than $\mathcal{G}^{-1}(c)$. These are the types for which it is worthwhile for the seller to invest c in order to give trade a chance. Also, as in the free learning setting, those buyer types that have a positive chance of trade need to be given the incentive to accept to trade under the conditions of the benchmark allocation and that requires fee/subsidy $\widehat{\lambda}(\theta)$. This latter function still maintains differentiability and the property that:

$$\int_0^1 \widehat{\lambda}(\theta)dF(\theta) = \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^1 \widehat{\lambda}(\theta)dF(\theta) = 0$$

so that, as one would expect, the expected utility for the seller in the optimal mechanism is the benchmark expected utility $V^{BNC}(c)$. It also remains true that $\widehat{\lambda}(1) < 0$ but because of the truncation we have $\widehat{\lambda}(\psi^{-1}(\mathcal{G}^{-1}(c))) > 0$.²⁴ Finally, the immediate indirect

²⁴The proof of these claims about the $\widehat{\lambda}(\theta)$ function is immediate and is available upon request.

implementation is once again a menu of payment pairs, but with a lower bound $\mathcal{G}^{-1}(c)$ on the prices that the buyer is allowed to offer: $\{(p, \widehat{\Lambda}(p))\}_{p \in [\mathcal{G}^{-1}(c), 1]}$, where $\widehat{\Lambda}(p) = \widehat{\lambda}(\psi^{-1}(p))$. As in the free learning case, for each menu $\widehat{\Lambda}(p)$ represents a fee/subsidy that is paid regardless of trade, while trade occurs at price p if the seller agrees to deliver.

5.1.2 Optional learning

We now look at the case where the seller still pays a cost $c > 0$ to learn her type, but where this is not necessary for trade: if she does not pay c , the seller can still provide the good to the buyer, but won't know what her ex-post utility is. In this case, the rationale for the seller to learn her type is not to be able to trade but to be able to tailor the allocation rule to both θ and ω , otherwise, the allocation must solely depend on the former.

Benchmark mechanism. We begin again with the benchmark case and provide a counterpart to propositions 1 (the free learning case) and 5 (the costly and necessary learning case):

Proposition 7. *In the benchmark case of optional learning with a cost $c > 0$, the seller will not pay the cost c whenever $c > \mathcal{G}(\mathbb{E}(\omega))$. In that case, the optimal allocation is:*

$$q_0(\theta, \omega) = \mathbf{1}[\mathbb{E}(\omega) \leq \psi(\theta)],$$

where $\mathbb{E}(\omega)$ is the expected value of ω . This can be implemented by transfers

$$t_0(\theta, \omega) = \psi^{-1}(\mathbb{E}(\omega)) \cdot \mathbf{1}[\mathbb{E}(\omega) \leq \psi(\theta)].$$

If $c \leq \mathcal{G}(\mathbb{E}(\omega))$ then the seller will pay the cost if and only if $\theta \in [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]$, where

$$\chi(\theta) = \mathbb{E}(\omega) + \mathcal{G}(\psi(\theta)) - \psi(\theta),$$

so that $\psi^{-1}(\mathcal{G}^{-1}(c)) \geq \psi^{-1}(0)$ and $\chi^{-1}(c) \leq 1$. The optimal allocation is

$$q_0(\theta, \omega) = \mathbf{1}[\omega \leq \psi(\theta)] \cdot \mathbf{1}[\theta \in [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]] + \mathbf{1}[\theta > \chi^{-1}(c)].$$

This can be implemented by transfers

$$t_0(\theta, \omega) = \rho(\theta) \cdot \mathbf{1}[\omega \leq \psi(\theta)] \cdot \mathbf{1}[\theta \in [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]] + \bar{\rho} \cdot \mathbf{1}[\theta > \chi^{-1}(c)],$$

where

$$\rho(\theta) = \left(\theta - \frac{\int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^{\theta} G(\psi(y)) dy}{G(\psi(\theta))} \right) \text{ and } \bar{p} = \left(\chi^{-1}(c) - \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^{\chi^{-1}(c)} G(\psi(y)) dy \right).$$

The expected utility for the seller from this mechanism if $c > \mathcal{G}(\mathbb{E}(\omega))$ is

$$V^{BOC}(c) = \int_{\psi^{-1}(\mathbb{E}(\omega))}^1 \int_0^1 (\psi(\theta) - \omega) dG(\omega) dF(\theta)$$

while if $c \leq \mathcal{G}(\mathbb{E}(\omega))$ then

$$\begin{aligned} V^{BOC}(c) &= \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^{\chi^{-1}(c)} \left(-c + \int_0^{\psi(\theta)} (\psi(\theta) - \omega) dG(\omega) \right) dF(\theta) \\ &+ \int_{\chi^{-1}(c)}^1 \int_0^1 (\psi(\theta) - \omega) dG(\omega) dF(\theta). \end{aligned}$$

To parse this result, it is immediately easy to see that if $c > \mathcal{G}(\mathbb{E}(\omega))$, then the seller never finds it optimal to learn her type, whatever the buyer's type she faces. Therefore, she will sell the good if the buyer's virtual valuation is above her expected type (or, alternatively, at price $\psi^{-1}(\mathbb{E}(\omega))$). Since the seller does not learn her type, the allocation cannot depend on it, and delivery will be guaranteed. Obviously, this also means that $V^{BOC}(c)$ is positive and constant in c for all $c > \mathcal{G}(\mathbb{E}(\omega))$.

If $c \leq \mathcal{G}(\mathbb{E}(\omega))$, on the other hand, things are more complicated. First note that the seller observes θ as reported by the buyer, and given this buyer type, she can decide i) whether there is an opportunity to trade with such type and ii) whether she should pay c to learn her ω in order to make trade contingent on her type being smaller than the buyer's or whether she should not pay c but commit to trade for sure with this buyer. The function $\mathcal{G}(\psi(\theta))$ represents the expected utility that the seller can expect from trading with type θ if she pays c and learns her type. Hence, if $\theta \geq \psi^{-1}(\mathcal{G}^{-1}(c))$, the seller will prefer to pay c , learn her own type and have a positive probability of trading with θ over not trading with θ and not learning her type. The function $\chi(\theta)$, instead, represents the benefit that the seller has in learning ω and then trade with θ if $\omega \leq \theta$ over not learning ω and trading with θ for sure. Thus, if $\theta \leq \chi^{-1}(c)$ then the seller prefers to pay c and leave herself the opportunity to not trade with θ if it turns out that $\theta < \omega$. This means that, having observed the buyer's type, the seller will not invest in learning her cost if the buyer's type is too low ($\theta < \psi^{-1}(\mathcal{G}^{-1}(c))$) or too high ($\theta > \chi^{-1}(c)$). In the first case, this is because the seller does not wish to trade with such buyer type. In the second case, instead, the buyer type is so high that the seller prefers to sell with guaranteed delivery

because the price she gets is sufficiently high that the possibility that her type turns out to be above such price (where she would benefit from not delivering) is outweighed by the cost of learning the type itself.²⁵

For intermediate values of the buyer's type, the seller will decide to learn her cost. For such buyer types, the price that she obtains is not too high. Thus, she prefers to pay the cost of learning her type because this allows her to cancel the trade—if unprofitable—over not paying the cost and having to guarantee delivery. In the proof we show that $\psi^{-1}(\mathcal{G}^{-1}(c))$ is a strictly increasing function of c with $\psi^{-1}(\mathcal{G}^{-1}(0)) = \psi^{-1}(0)$ and $\psi^{-1}(\mathcal{G}^{-1}(\mathcal{G}(\mathbb{E}(\omega)))) = \psi^{-1}(\mathbb{E}(\omega))$ while $\chi^{-1}(c)$ is a strictly decreasing function of c with $\chi^{-1}(0) = 1$ and $\chi^{-1}(\mathcal{G}(\mathbb{E}(\omega))) = \psi^{-1}(\mathbb{E}(\omega))$ so that the interval $[\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]$ becomes the interval $[\psi^{-1}(0), 1]$ when $c = 0$. This is intuitive: if the cost of learning her type approaches zero, the seller always learns whenever a buyer's type is one where she may wish to trade, and so the optimal allocation converges to that under free learning. Conversely, the interval disappears when $c = \mathcal{G}(\mathbb{E}(\omega))$ because the cost of learning her type is higher than any possible benefit she might get, and she will either not trade or trade with guaranteed delivery.

As for the transfers, note that $\rho(\theta)$ is an increasing function with $\rho(\chi^{-1}(c)) < \bar{\rho}$. This is important because $\rho(\theta)$ is a price that the buyer agrees to pay, but trade will only occur if the seller's type is smaller than $\rho(\theta)$, whereas $\bar{\rho}$ is a price that guarantees delivery. So, to be incentive compatible for buyer of type $\chi^{-1}(c)$ to be indifferent between $\rho(\chi^{-1}(c))$ and $\bar{\rho}$, the latter must be sufficiently higher than the former.

Finally, it is easy to see that $V^{BOC}(c) > V^{BNC}(c)$ for any $c > 0$. This is quite obvious as the under necessary learning the seller has two options for any possible buyer's type: give him a positive probability of trading and pay c or not give him any chance of trade. In optional learning, the seller always also has the possibility of trading with a particular buyer type and not pay c .

Optimal Mechanism. As in the case with necessary learning, if we consider the case where the seller cannot commit to revealing her type truthfully under optional learning, the analogy with the case where learning is free is evident. Note first, however, that if $c > \mathcal{G}(\mathbb{E}(\omega))$, then the benchmark allocation is *independent* of the seller's type, simply because this is unknown. Hence, whether the seller can commit to reveal something she doesn't know is irrelevant and the optimal mechanism is the benchmark mechanism. Thus, the issue is only relevant when $c \leq \mathcal{G}(\mathbb{E}(\omega))$ so that there is an interval $[\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]$

²⁵In the proof we show, in particular, that $\chi^{-1}(c) \geq \psi^{-1}(\mathbb{E}(\omega))$ so that whenever $\theta > \chi^{-1}(c)$, automatically $\theta > \psi^{-1}(\mathbb{E}(\omega))$ which is the condition that guarantees that the seller wishes to trade with guaranteed delivery with this buyer type.

of buyer types for which the allocation depends on ω .

Proposition 8. *Suppose $c \leq \mathcal{G}(\mathbb{E}(\omega))$. In the case of optional learning, when the seller cannot commit to truthfully revealing her type, there is an optimal mechanism that implements the benchmark allocation:*

$$q_0(\theta, \omega) = \mathbf{1}[\omega \leq \psi(\theta)] \cdot \mathbf{1}[\theta \in [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]] + \mathbf{1}[\theta > \chi^{-1}(c)].$$

Transfers for this mechanism are:

$$t_0(\theta, \omega) = \psi(\theta) \cdot \mathbf{1}[\omega \leq \psi(\theta)] + \bar{\lambda}(\theta) \cdot \mathbf{1}[\theta \in [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]] + \bar{p} \cdot \mathbf{1}[\theta > \chi^{-1}(c)],$$

where \bar{p} is defined in proposition 7 while

$$\bar{\lambda}(\theta) = (\theta - \psi(\theta))G(\psi(\theta)) - \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^{\theta} G(\psi(y))dy$$

for all $\theta \in [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]$ and $\bar{\lambda}(\theta) = 0$ otherwise.

The difference between the benchmark and the optimal mechanisms is only on the transfers over the interval $[\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]$. That these transfers are dominant strategy incentive compatible for the seller follows immediately using the same argument as in proposition 2. The only other task is to verify that the transfers also equal expected transfers for the buyer, which again is the same as in proposition 2.

The benchmark mechanism and the optimal mechanism are the same except that when $\theta \in [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]$, the price is $\psi(\theta)$ to make sure the benchmark allocation is incentive compatible for the seller. Correspondingly, there is the additional fee or subsidy $\bar{\lambda}(\theta)$ that is paid regardless of trade that serves the purpose of guaranteeing incentive compatibility for buyer types in this region. There is, therefore, at least for these buyer types, a two-part tariff. Note that the truncation of the fee/subsidy function here is from two sides: if $\theta < \psi^{-1}(\mathcal{G}^{-1}(c))$ then there is no possibility of trade, and no payments are made, whereas if $\theta > \chi^{-1}(c)$, the seller chooses a price that is incentive compatible for the buyer without needing to make sure it is incentive compatible for her future self. Thus, a two-part tariff is no longer needed. One of the consequences is that while with both free learning and necessary learning (when the cost was not excessive), the very highest types always were offered an allocation contingent on the seller's type and always received a subsidy, here the highest seller types are offered a price with guaranteed delivery. This means that since $\chi^{-1}(c)$ is decreasing in c , eventually it will become sufficiently small that $\bar{\lambda}(\theta)$ will be positive for all $\theta \in [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]$. In other words, if c is small but not

too small, all types for which the seller decides to pay c and learn her type, will have to pay a positive fee in the optimal mechanism.

Finally, a natural indirect implementation of this mechanism is the following. When $c > \mathcal{G}(\mathbb{E}(\omega))$ then, as already discussed in the benchmark case, the seller offers the buyer a price $\psi^{-1}(\mathbb{E}(\omega))$ with guaranteed delivery. When $c \leq \mathcal{G}(\mathbb{E}(\omega))$, the seller offers the buyer *both* the price \bar{p} with guaranteed delivery *and* a menu of price-fee or subsidy payments with a lower bound $\mathcal{G}^{-1}(c)$ *and* an upper bound $\psi(\chi^{-1}(c))$ on the prices that the buyer is allowed to choose: $\{(p, \bar{\Lambda}(p))\}_{p \in [\mathcal{G}^{-1}(c), \psi(\chi^{-1}(c))]}$, where $\bar{\Lambda}(p) = \bar{\lambda}(\psi^{-1}(p))$. As for the previous mechanisms, $\bar{\Lambda}(p)$ is paid regardless of trade, while the prices do not guarantee delivery: they are contingent on the seller's willingness to deliver.²⁶

This implementation is similar to the optimal ex-ante price mechanism in the sense that it offers a menu that includes guaranteed delivery options as well as at-will options. The difference is that in the case of price mechanisms, the at-will elements are distorted due to incentive constraints (due to the inability to use two-part tariffs), whereas in the case of costly learning, the distortion comes from learning costs.

5.2 Efficiency

Our focus in the paper is profit maximization for the seller when she learns her information with delay. Our analysis shows that, in general, a mechanism exists that implements the same allocation as in the benchmark cases where either the seller can commit to revealing her information truthfully (eliminating incentive compatibility issues on the seller's side) or when she can commit to ignoring the information revealed by the buyer (equivalent to facing interim incentive compatibility constraints for her future self). However, this mechanism requires payments even when the trade does not occur.

Therefore, it is natural to ask what results an analogous analysis would produce when the objective is ex-post efficiency. This is particularly interesting because in this case, as shown by Myerson and Satterthwaite (1983), we cannot have ex-post efficiency when both buyer and seller face Bayesian incentive compatibility and interim individual rationality constraints, while it can be implemented if those constraints are applied only to one side. In this subsection, as we did in section 3.2, we modify the constraints the seller faces by imposing dominant strategy incentive compatibility. We will also not impose interim individual rationality since this is meaningless when the seller learns with delay. Instead, it makes sense to ask whether she is willing to participate at the ex-ante stage or, more stringently, at a “reverse” interim stage, where she has learned the buyer's type but not yet her own. As we show below, contrary to the Myerson and Satterthwaite (1983) setting,

²⁶Because of this, $\psi(\chi^{-1}(c)) < \bar{p}$.

ex-post efficiency is implementable by a two-part tariff mechanism, regardless of which of the two individual rationality constraints for the seller discussed above we impose:

Proposition 9. *Under the incentive and individual rationality constraints posed by our setting, the ex-post efficient allocation $q_0(\theta, \omega) = \mathbf{1}[\omega \leq \theta]$ is implementable by a mechanism where transfers are:*

$$t_0(\theta, \omega) = \theta \mathbf{1}[\omega \leq \theta] - \mathcal{G}(\theta).$$

The proof is in appendix A.4.3. An immediate observation one can make following proposition 9 is that while the delay with which the seller learns her type makes us impose a stronger constraint on her incentive compatibility than in Myerson and Satterthwaite (1983), the same structure allows us to weaken the seller’s individual rationality in a way that more than compensates and allows for the implementation of ex-post efficiency. Not surprisingly, given this observation, the mechanism requires the seller to provide each buyer type with a subsidy to compensate the latter for choosing a price equal to his type. Thus, all buyer types get a surplus equal to $\mathcal{G}(\theta)$ and in appendix A.4.3, we show that the “reverse” interim individual rationality is tight: upon learning the buyer’s valuation, the seller’s expected utility is zero.

This result also shows how the seller’s constraint in dominant strategies still forces the mechanism to pinpoint exactly what transfers are needed. This means that two-part tariffs are needed, although the fixed component is always a subsidy in this case.

Obviously, we found this mechanism because the informational environment is very different from that of Myerson and Satterthwaite (1983): if the seller knew her type before deciding whether to participate in the mechanism, then high ω types would not want to participate. Type $\omega = 1$, for example, would expect no gains from trade whatever θ is but would still have to pay the subsidy. The fact that the seller has to agree *before* she learns her type is crucial for implementing this mechanism.

Finally, the mechanism itself suggests an obvious indirect implementation: the buyer can pick any price $p \in [0, 1]$. For any such price, the buyer gets a subsidy $\mathcal{G}(p)$ from the seller and receives the good at price p if the seller, upon learning her cost, agrees to deliver at that price.

6 Discussion

We studied bilateral trade in situations where the seller wishes to maximize her profit but learns her type with delay. The mechanism must therefore account for the seller’s own incentive to take advantage of the information about the buyer she learns. We identify

a potential trade-off between contracting earlier with limited information, which creates stronger incentive problems, and postponing the contract to avoid those issues, though at the cost of delay.

Our first main result shows that an optimal mechanism can effectively address these challenges. The optimal mechanism involves a menu of two-part tariffs. Traditionally, two-part tariffs are used in standard contracting situations to incentivize the agent to choose the desired action, with the fixed component capturing the surplus ex-ante. For instance, in a typical textbook model involving an agent with a known downward-sloping demand for a good, a monopolist can capture the full consumer surplus by charging the marginal cost as a marginal price and a fixed fee equal to the induced consumer surplus. In our model, however, the value of two-part tariffs lies in their ability to ensure that the seller's own future incentives are correctly aligned through delivery fees, while the fixed component controls the buyer's incentives to select the appropriate menu item. Consequently, our two-part tariffs exhibit a unique feature: unlike the classic literature, the fixed fee must sometimes be positive (for intermediate buyer types) and sometimes negative (for high buyer types). This is driven by the need to adjust informational rents: intermediate types are offered a relatively low delivery price, resulting in high ex-post information rents. To reduce these information rents, an additional charge is necessary. Conversely, high types are charged a high delivery price, so providing them with the necessary information rents requires an ex-ante subsidy. Our extensions to costly learning demonstrate that this result is robust to changes in the model.

Our second main result focuses on price mechanisms. This analysis is motivated by the observation that the menus of two-part tariffs we identify as optimal are not very common in practice, possibly because the implementation of the two-part tariffs may be too complex. Therefore, we study a more restricted class of mechanisms to find more realistic options. We first characterize the optimal static price mechanisms, where the contract occurs either before or after the cost information is available, without any dynamic screening. If the discount factor is sufficiently high, the optimal static mechanism waits until the cost information is known and then makes a cost-dependent take-it-or-leave-it offer. However, if the discount factor is low, the optimal ex-ante price mechanism involves a menu of offers, including one or more at-will contracts and one guaranteed-delivery offer. The highest buyer types opt for guaranteed delivery at a higher price, while at-will contracts enable the seller to further screen intermediate-type buyers. Our second main result shows that the optimal price mechanism becomes dynamic if the discount factor is large enough. This dynamic mechanism balances trade-offs by offering early delivery to high types, thus reducing delay costs, and late delivery to low types, thus reducing incentive problems.

We extend our analysis of general mechanisms in two ways: we first consider the possibility that learning her type is costly for the buyer. We distinguish the case where the seller needs to learn her type in order to trade from the case where she can still trade even if she does not learn. Our results show that the analysis for the case where learning is free, is essentially robust in that when learning the cost is necessary, the seller either decides not to trade at all (if the cost is too high), or uses a two-part tariff. When learning is optional, the seller either chooses to trade at a fixed price without learning the cost (if the cost is too high) or chooses to trade in the same way with high buyer types, but learns the cost and uses a two-part tariff with intermediate buyer types.

Finally, we also look at the issue of implementation of ex-post efficiency. Our buyer faces the same constraints as in the classic Myerson and Satterthwaite (1983) setting, while the seller faces stronger incentive and weaker participation constraints. We show that, again using a two-part tariff, *there is* a mechanism that implements ex-post efficiency. In this case, however, the fixed component is always a subsidy: here the mechanism is such that the seller is always forced to give up all the surplus to the buyer.

A key limitation of our analysis is that while we demonstrated that dynamic price mechanisms can sometimes be optimal, we did not explicitly characterize what these optimal mechanisms look like. Identifying them remains an open question that requires further exploration. More broadly, this work represents an initial step towards a larger agenda in mechanism and information design, focusing on models where the principal actively participates in the designed game because she will have information, but with delay. While similar concepts, such as credible auctions and renegotiation-proof games, have been explored in the literature, there is still substantial work to be done.

References

- Akbarpour, M. and Li, S. (2020) Credible Auctions: A Trilemma, *Econometrica*, **88**, 425–467.
- 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.
- 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.
- 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. (2022) Mechanism Design With Limited Commitment, *Econometrica*, **90**, 1463–1500.
- Doval, L. and Skreta, V. (2024) Optimal Mechanism for the Sale of a Durable Good, *Theoretical Economics*, **Forthcoming**.
- 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.
- Hagerty, K. M. and Rogerson, W. P. (1987) Robust Trading Mechanisms, *Journal of Economic Theory*, **42**, 94–107.
- 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.
- Manelli, A. M. and Vincent, D. R. (2010) Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model, *Econometrica*, **78**, 1905–1938.
- 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.
- Myerson, R. B. and Satterthwaite, M. A. (1983) Efficient Mechanisms for Bilateral Trading, *Journal of Economic Theory*, **29**, 265–281.
- 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*, pp. 2563–2577.
- 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.
- Vartiainen, H. (2013) Auction Design Without Commitment, *Journal of the European Economic Association*, **11**, 316–342.
- Williams, S. R. (1987) Efficient Performance in Two Agent Bargaining, *Journal of Economic Theory*, **41**, 154–172.
- 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 Proofs for Section 2

A.1.1 Proof of Lemma 1

First, let us introduce the notation for mixed strategies.

- $\beta_0(\theta)$ is a random variable, with distribution $B_0(b_0|\theta)$, $\text{supp } \beta_0(\theta) \subset \mathcal{B}_0$ is its support and $\mathbb{E}_{\beta_0(\theta)}[\cdot]$ is the expectation over \mathcal{B}_0 with respect to β_0 for a given θ .²⁷
- Similarly, $\text{supp } \sigma(\omega, b_0)$, $\Sigma(s|\omega, b_0)$, $\mathbb{E}_{\sigma(\omega, b_0)}[\cdot]$, for each ω and $b_0 \in \mathcal{B}_0$.
- Similarly, $\text{supp } \beta_1(\theta, b_0, s)$, $B_1(b_1|\theta, b_0, s)$, $\mathbb{E}_{\beta_1(\theta, b_0, s)}[\cdot]$, for each θ, b_0, s .

Fix a mechanism and an equilibrium:

- $\rho^{B^1}(\theta|b_0, s) = \mathbb{E}_{\beta_1(\theta, b_0, s)}q(b_0, s, b_1)$ is the ex-post expected net present value of the quantity, conditional on history (b_0, s) and buyer type θ and $\tau^{B^1}(\theta|b_0, s)$ is defined analogously.
- $\rho^S(\theta, \omega|b_0) = \mathbb{E}_{\sigma(\omega, b_0)}\rho^{B^1}(\theta|b_0, s)$ and $\tau^S(\theta, \omega|b_0)$ is defined analogously.

²⁷With continuous distributions, $\mathbb{E}_{\beta_0(\theta)}[f(b_0)] = \int_{\mathcal{B}_0} f(b_0)dB_0(b_0|\theta)$, and an equivalent of this with mass points.

- $\rho^{B0}(\theta, \omega) = \mathbb{E}_{\beta_0(\theta)} \rho^S(\theta, \omega | b_0)$ and an analogous $\tau^{B0}(\theta, \omega)$.

In other words, these are allocation and transfer functions (NPV of each), conditional on choices made before each decision node, fixing type profiles θ, ω , and assuming equilibrium behavior starting from period t . These objects depend on true values (θ, ω) , but can be defined in the same way for situations where some type mimics the behavior of another type. In other words, for a fixed mechanism and an equilibrium, these objects are defined for B, who behaves “as if” he is type θ , and S, who behaves as if being “ ω ”.

Next, we introduce the notation for the equilibrium payoffs at each history, assuming equilibrium behavior starting from this point onwards:

- $\nu^{B1}(\theta | b_0, s) = \theta \rho^{B1}(\theta | b_0, s) - \tau^{B1}(\theta | b_0, s)$ is the utility after history (b_0, s) for the buyer of type θ .
- $\nu^S(\omega | b_0) = \int_{\Theta} [\tau^S(\theta, \omega | b_0) - \omega \rho^S(\theta, \omega | b_0)] dF(\theta | b_0)$ is the profit after history b_0 for the seller of type ω .
- $\nu^{B0}(\theta) = \int_{\Omega} [\theta q^0(\theta, \omega) - t^0(\theta, \omega)] dG(\omega)$ is the utility of the buyer of type θ before any choices.

With this notation, we can compactly define the equilibrium in mixed strategies²⁸ (B_0, Σ, B_1) is an equilibrium in (mixed) strategies, if

$$\nu^{B0}(\theta) \geq 0, \forall \theta, \quad (\text{IRB}_0)$$

$$\nu^{B0}(\theta) \geq \int_{\Omega} [\theta \rho^S(\theta, \omega | b_0) - \tau^S(\theta, \omega | b_0)] dG(\omega), \forall \theta, \quad (\text{ICB}_0)$$

$$\nu^S(\omega | b_0) \geq \int_{\Theta} [\tau^{B1}(\theta | b_0, s) - \omega \rho^{B1}(\theta | b_0, s)] dF(\theta | b_0), \forall b_0, \omega, \quad (\text{ICS}_0)$$

$$\nu^{B1}(\theta | b_0, s) \geq \theta q(b_0, s, b_1) - t(b_0, s, b_1), \forall b_0, s, \omega, \quad (\text{ICB}_1)$$

where (ICB_0) , (ICS_0) , and (ICB_1) hold as equalities whenever the corresponding value is on the support. For example, (ICB_0) is an equality for all $b_0 \in \text{supp } \beta_0(\theta)$.

Proof of Lemma 1. We prove the result in three steps. The first, lemma 1.B1, shows that the buyer’s last message space \mathcal{B}_1 can be without loss of generality replaced by type space Θ . The second, lemma 1.S shows similarly that the seller’s message space \mathcal{S} can be replaced by type space Ω . Finally, lemma 1.B0 constructs a reduced-form mechanism corresponding to a general mechanism using a garbling device.

²⁸These conditions are direct generalizations of the corresponding conditions for pure strategies in the main text, we use the same labels.

Lemma 1.B1. *For any mechanism \mathcal{M} and an equilibrium (B_0, Σ, B_1) , there exists an alternative mechanism $\widehat{\mathcal{M}}$ and a corresponding equilibrium $(\widehat{B}_0, \widehat{\Sigma}, \widehat{B}_1)$, that is payoff-equivalent to the original equilibrium in \mathcal{M} , and*

1. $\widehat{\mathcal{B}}_0 = \mathcal{B}_0, \widehat{\mathcal{S}} = \mathcal{S}, \widehat{\mathcal{B}}_1 = \Theta,$
2. $\widehat{B}_0 = B_0, \widehat{\Sigma} = \Sigma,$ and $\beta_1(\theta, b_0, s) = \theta.$

Proof. Take any mechanism $\mathcal{M} = (\mathcal{B}_0, \mathcal{S}, \mathcal{B}_1, q_0, t_0, q_1, t_1)$ and an equilibrium (B_0, Σ, B_1) . We construct a new mechanism: $\widehat{\mathcal{M}}$ such that it mimics the same mechanism but with type announcement in the last step. That is, the message spaces are $\widehat{\mathcal{B}}_0 = \mathcal{B}_0, \widehat{\mathcal{S}} = \mathcal{S},$ and $\widehat{\mathcal{B}}_1 = \Theta,$ and $\widehat{q}_0(b_0, s) = q_0(b_0, s), \widehat{t}_0(b_0, s) = t_0(b_0, s), \widehat{q}_1(b_0, s, \theta) = \mathbb{E}_{\beta_1(\theta, b_0, s)} q_1(b_0, s, b_1),$ and $\widehat{t}_1(b_0, s, \theta) = \mathbb{E}_{\beta_1(\theta, b_0, s)} t_1(b_0, s, b_1).$

We claim that the following strategy profile $(\widehat{B}_0, \widehat{\Sigma}, \widehat{B}_1)$ is an equilibrium in $\widehat{\mathcal{M}}$:

$$\widehat{B}_0(\theta) = B_0(\theta), \quad \widehat{\Sigma}(\omega, b_0) = \Sigma(\omega, b_0), \quad \widehat{\beta}_1(\theta, b_0, s) = \theta$$

and it gives the same expected payoffs for both players.

Using the same notation as above, but adding a hat for the objects associated to mechanism $\widehat{\mathcal{M}}$ and denoting the type declaration by θ' , we get that

$$\begin{aligned} \widehat{q}(b_0, s, \theta') &= \widehat{q}_0(b_0, s) + \delta \widehat{q}_1(b_0, s, b_1) = \mathbb{E}_{\beta_1(\theta', b_0, s)} q(b_0, s, b_1), \\ \widehat{\rho}^{B^1}(\theta|b_0, s) &= \mathbb{E}_{\widehat{\beta}_1(\theta, b_0, s)} \widehat{q}(b_0, s, \theta') = \widehat{q}(b_0, s, \theta) = \mathbb{E}_{\beta_1(\theta, b_0, s)} q(b_0, s, b_1) = \rho^{B^1}(\theta|b_0, s), \\ \widehat{\rho}^S(\theta, \omega|b_0) &= \mathbb{E}_{\widehat{\sigma}(\omega, b_0)} \widehat{\rho}^{B^1}(\theta|b_0, s) = \mathbb{E}_{\sigma(\omega, b_0)} \rho^{B^1}(\theta|b_0, s) = \rho^S(\theta, \omega|b_0), \\ \widehat{\rho}^{B^0}(\theta, \omega) &= \mathbb{E}_{\widehat{\beta}_0(\theta)} \widehat{\rho}^S(\theta, \omega|b_0) = \mathbb{E}_{\beta_0(\theta)} \rho^S(\theta, \omega|b_0) = \rho^{B^0}(\theta, \omega). \end{aligned}$$

Transfers have the same property. Therefore also the payoffs $\widehat{\nu}^{B^1} = \nu^{B^1}, \widehat{\nu}^S = \nu^S,$ and $\widehat{\nu}^{B^0} = \nu^{B^0}.$ This establishes the payoff-equivalence. Finally, we need to verify that the new strategy profile satisfies the equilibrium conditions (IRB₀), (ICB₀), (ICS₀), and (ICB₁). The first three are immediate, as they only depend on objects that are unchanged. For (ICB₁)

$$\begin{aligned} \widehat{\nu}^{B^1}(\theta|b_0, s) &\geq \theta \widehat{q}(b_0, s, \theta') - \widehat{t}(b_0, s, \theta'), \text{ with “=” for } \theta' = \theta \iff \\ \nu^{B^1}(\theta|b_0, s) &\geq \mathbb{E}_{\beta_1(\theta, b_0, s)} [\theta q(b_0, s, b_1) - t(b_0, s, b_1)], \text{ with “=” for } \theta' = \theta. \end{aligned}$$

As \mathcal{M} satisfies (ICB₁) for all b_0 , we have that for all $b_1 \in \mathcal{B}_1,$

$$\nu^{B^1}(\theta|b_0, s) \geq \theta q(b_0, s, b_1) - t(b_0, s, b_1), \quad (9)$$

so the inequality must also hold in expectation. Moreover, for each $b_1 \in \text{supp } \beta_1(\theta, b_0, s)$, Equation (9) holds as an equality, therefore taking an expectation over $\beta_1(\theta, b_0, s)$ gives also an equality. This establishes that (ICB₁) is still satisfied. \square

Lemma 1.S. *For any mechanism \mathcal{M} and an equilibrium (B_0, Σ, B_1) , such that $\mathcal{B}_1 = \Theta$ and $\beta_1(\theta, b_0, s) = \theta$, there exists an alternative mechanism $\widehat{\mathcal{M}}$ and a corresponding equilibrium $(\widehat{B}_0, \widehat{\Sigma}, \widehat{B}_1)$, that is payoff-equivalent to the original equilibrium in \mathcal{M} , and*

1. $\widehat{\mathcal{B}}_0 = \mathcal{B}_0, \widehat{\mathcal{S}} = \Omega, \widehat{\mathcal{B}}_1 = \Theta,$
2. $\widehat{B}_0 = B_0, \widehat{\sigma}(\omega, b_0) = \omega,$ and $\beta_1(\theta, b_0, s) = \theta.$

Proof. Take a mechanism \mathcal{M} and an equilibrium (B_0, Σ, B_1) satisfying the conditions in the lemma. We construct a new mechanism: $\widehat{\mathcal{M}}$ such that it mimics the same mechanism, but with type announcement by S. The message spaces are $\widehat{\mathcal{B}}_0 = \mathcal{B}_0, \widehat{\mathcal{B}}_1 = \mathcal{B}_1 = \Theta,$ and $\widehat{\mathcal{S}} = \Omega,$ and $\widehat{q}_0(b_0, s) = \mathbb{E}_{\sigma(\omega, b_0)} q_0(b_0, s), \widehat{t}_0(b_0, s) = \mathbb{E}_{\sigma(\omega, b_0)} t_0(b_0, s), \widehat{q}_1(b_0, s, \theta) = \mathbb{E}_{\sigma(\omega, b_0)} q_1(b_0, s, b_1),$ and $\widehat{t}_1(b_0, s, \theta) = \mathbb{E}_{\sigma(\omega, b_0)} t_1(b_0, s, \theta).$

We claim that the following strategy profile $(\widehat{B}_0, \widehat{\Sigma}, \widehat{B}_1)$ is an equilibrium in $\widehat{\mathcal{M}}$:

$$\widehat{B}_0(\theta) = B_0(\theta), \quad \widehat{\sigma}(\omega, b_0) = \omega, \quad \beta_1(\theta, b_0, s) = \beta_1(\theta, b_0, s) = \theta$$

and it gives the same expected payoffs for both players.

Using the same notation as above, and again using ω' and θ' as type announcements, we now have

$$\begin{aligned} \widehat{q}(b_0, \omega', \theta') &= \widehat{q}_0(b_0, \omega') + \delta q_1(b_0, \omega', \theta') = \mathbb{E}_{\sigma(\omega', b_0)} q(b_0, s, \theta'), \\ \widehat{\rho}^{B^1}(\theta|b_0, \omega') &= \mathbb{E}_{\widehat{\beta}_1(\theta, b_0, \omega')} \widehat{q}(b_0, \omega', b_1) = \widehat{q}(b_0, \omega', \theta) = \mathbb{E}_{\sigma(\omega', b_0)} q(b_0, s, \theta) = \mathbb{E}_{\sigma(\omega', b_0)} \rho^{B^1}(\theta|b_0, s), \\ \widehat{\rho}^S(\theta, \omega|b_0) &= \mathbb{E}_{\widehat{\sigma}(\omega, b_0)} \widehat{\rho}^{B^1}(\theta|b_0, \omega') = \widehat{\rho}^{B^1}(\theta|b_0, \omega) = \mathbb{E}_{\sigma(\omega', b_0)} \rho^{B^1}(\theta|b_0, s) = \rho^S(\theta, \omega|b_0), \\ \widehat{\rho}^{B^0}(\theta, \omega) &= \mathbb{E}_{\widehat{\beta}_0(\theta)} \widehat{\rho}^S(\theta, \omega) = \mathbb{E}_{\beta_0(\theta)} \rho^S(\theta, \omega|b_0) = \rho^{B^0}(\theta, \omega). \end{aligned}$$

Again, similar properties hold for transfers. Therefore, payoffs satisfy $\widehat{\nu}^S = \nu^S$ and $\widehat{\nu}^{B^0} = \nu^{B^0}$, which immediately establish the payoff equivalence. Also, (IRB₀) and (ICB₀) hold. As \widehat{q} and $\widehat{\rho}^{B^1}$ now differ, we still have to verify (ICS₀) and (ICB₁) in the new mechanism.

Let us consider (ICS₀) first.

$$\begin{aligned} \widehat{\nu}^S(\omega|b_0) &\geq \int_{\Theta} [\widehat{\tau}^{B^1}(\theta|b_0, \omega') - \omega \widehat{\rho}^{B^1}(\theta|b_0, \omega')] dF(\theta|b_0), \text{ with “=” for } \omega' = \omega \iff \\ \widehat{\nu}^S(\omega|b_0) &\geq \mathbb{E}_{\sigma(\omega', b_0)} \int_{\Theta} [\tau^S(\theta|b_0, s) - \omega \rho^S(\theta|b_0, s)] dF(\theta|b_0), \text{ with “=” for } \omega' = \omega. \end{aligned}$$

Note that by the (ICS₀) in the original mechanism and equilibrium, this inequality holds for each s separately, therefore it also holds in expectation. Moreover, the inequality holds as equality for each $s \in \text{supp } \sigma(\omega, b_0)$, so when $\omega' = \omega$, the condition indeed holds as equality. This establishes the (ICS₀) condition.

Finally, consider (ICB₁).

$$\begin{aligned} \widehat{\nu}^{B_1}(\theta|b_0, \omega') &\geq \theta \widehat{q}(b_0, \omega', \theta') - \widehat{t}(b_0, \omega', \theta'), \text{ with “=” for } \theta' = \theta \iff \\ \mathbb{E}_{\sigma(\omega', b_0)} \nu^{B_1}(\theta|b_0, s) &\geq \mathbb{E}_{\sigma(\omega', b_0)} [\theta q(b_0, s, \theta') - t(b_0, s, \theta')], \text{ with “=” for } \theta' = \theta. \end{aligned}$$

Again, by (ICB₁), this equation holds for each s separately, so it also holds in expectation. And again, it holds as an equality when $\theta' = \theta$ for each s separately, therefore it also holds as an equality in expectation. \square

We have shown that we can set $\mathcal{B}_1 = \Theta$ and $\mathcal{S} = \Omega$. Setting $\mathcal{B}_0 = \Theta$ would not be without loss, because then the mechanism would reveal information about the buyer's type to the seller, i.e., if $\beta_0(\theta') = \beta_0(\theta) = b_0$, then b_0 does not distinguish θ' and θ , but reporting the type would distinguish. Therefore the (ICS₀) constraint would include an expectation over a different distribution. This is the reason we need to introduce a garbling device to the mechanism so that the seller only learns a realized b_0 and nothing else from the first message.

Lemma 1.B0. *For any mechanism \mathcal{M} and an equilibrium (B_0, Σ, B_1) , such that $\mathcal{B}_1 = \Theta, \mathcal{S} = \Omega, \beta_1(\theta, b_0, s) = \theta$, and $\sigma(\omega, b_0) = \omega$, there exists a reduced-form mechanism \mathcal{M}^R and truthtelling equilibrium, that is payoff-equivalent to the original equilibrium in \mathcal{M} .*

Proof. We construct the mechanism $\mathcal{M}^R = (\mathcal{B}, \mu, q_0^D, t_0^D, q_1^D, t_1^D)$ as follows: $\mathcal{B} = \mathcal{B}_0$, $\mu(b_0|\theta') = B_0(b_0|\theta')$, and $q_i^D \equiv q_i, t_i^D \equiv t_i$.

Clearly if both the buyer and the seller reveal their types truthfully, the mechanism \mathcal{M}^R ensures the same on-path payoffs as \mathcal{M} in the original equilibrium. We therefore to verify that truthtelling is an equilibrium.

As the mechanism starting from a particular realization of b_0 is unchanged, revealing ω is still optimal for the seller. Moreover, by construction, the on-path payoff for B is unchanged, so (IRB₀) is also satisfied.

We, therefore, only need to ensure that the does not have profitable deviations in terms of b_0 . For this, note that with \mathcal{M} , B has three types of potential deviations when choosing b_0 .

1. Deviations that consistently mimic the behavior of another type. That is, some type θ behaves exactly as another type θ' , i.e., $\beta'_0(\theta) = \beta_0(\theta')$ for some $\theta' \neq \theta$ and

$\beta'_1(\theta, b_0, s) = \theta'$ for all $b_0 \in \text{supp } \beta_0(\theta')$ and all s . The new constraint (ICB) ensures that such deviations are not profitable.

2. Inconsistent deviations, i.e., mimic some other type θ' at time 0 and report another type $\theta'' \neq \theta'$ at time 1. Such deviations are off-path. Therefore we can replace the original mechanism \mathcal{M} with another mechanism \mathcal{M}' such that any inconsistent deviations are punished harshly enough so that they are never optimal and the equilibrium payoffs are unchanged. Thus we can assume that incentive constraints for such deviations are not binding.
3. Finally, deviations to report b_0 that no type would choose in the original equilibrium. Again, such deviations are off-path and thus can be easily avoided, by replacing the original mechanism with a new one that excludes such unused messages from \mathcal{B}_0 .

□

This completes the proof of lemma 1.

□

A.2 Proofs for Section 3

A.2.1 Proof of Proposition 1

Proof. Write the following problem

$$\max_{q_0, t_0} \int_0^1 \int_0^1 [t_0(\omega, \theta) - \omega q_0(\omega, \theta)] dG(\omega) dF(\theta)$$

subject to

$$\int_0^1 [\theta q_0(\omega, \theta) - t_0(\omega, \theta)] dG(\omega) \geq 0 \text{ for all } \theta \in \Theta$$

$$\int_0^1 [\theta q_0(\omega, \theta) - t_0(\omega, \theta)] dG(\omega) \geq \int_0^1 [\theta q_0(\omega, \theta') - t_0(\omega, \theta')] dG(\omega) \text{ for all } \theta, \theta' \in \Theta$$

Define

$$Q_0(\theta) = \int_0^1 q_0(\omega, \theta) dG(\omega)$$

$$T_0(\theta) = \int_0^1 t_0(\omega, \theta) dG(\omega)$$

$$U_0(\theta) = \theta Q_0(\theta) - T_0(\theta)$$

then the buyer's constraints become, respectively

$$\begin{aligned}\theta Q_0(\theta) - T_0(\theta) &\geq 0 \text{ for all } \theta \in \Theta \\ \theta Q_0(\theta) - T_0(\theta) &\geq \theta Q_0(\theta') - T_0(\theta') \text{ for all } \theta, \theta' \in \Theta\end{aligned}$$

These constraints can be replaced by

$$\begin{aligned}U_0(0) = 0 &\Leftrightarrow T_0(0) = 0 \\ Q_0(\theta) &\text{ is weakly increasing in } \theta \\ \frac{dU_0(\theta)}{d\theta} = Q_0(\theta) &\Leftrightarrow U_0(\theta) = \int_0^\theta Q_0(y) dy \Leftrightarrow T_0(\theta) = \theta Q_0(\theta) - \int_0^\theta Q_0(y) dy\end{aligned}$$

Now

$$\begin{aligned}&\int_0^1 \int_0^1 [t_0(\omega, \theta) - \omega q_0(\omega, \theta)] dG(\omega) dF(\theta) \\ &= \int_0^1 T_0(\theta) dF(\theta) - \int_0^1 \omega q_0(\omega, \theta) dG(\omega) dF(\theta) \\ &= \int_0^1 \left[\theta Q_0(\theta) - \int_0^\theta Q_0(y) dy \right] dF(\theta) - \int_0^1 \int_0^1 \omega q_0(\omega, \theta) dG(\omega) dF(\theta) \\ &= \int_0^1 \int_0^1 (\theta - \omega) q_0(\theta, \omega) dG(\omega) dF(\theta) - \int_0^1 \int_0^\theta \int_0^1 q_0(y, \omega) dG(\omega) dy dF(\theta) \\ &= \int_0^1 \int_0^1 (\theta - \omega) q_0(\theta, \omega) dG(\omega) dF(\theta) - \int_0^1 \int_0^1 \int_y^1 q_0(y, \omega) dF(\theta) dG(\omega) dy \\ &= \int_0^1 \int_0^1 (\theta - \omega) q_0(\theta, \omega) dG(\omega) dF(\theta) - \int_0^1 \int_0^1 q_0(\theta, \omega) \left(\frac{1 - F(\theta)}{f(\theta)} \right) dG(\omega) dF(\theta) \\ &= \int_0^1 \int_0^1 (\psi(\theta) - \omega) q_0(\theta, \omega) dG(\omega) dF(\theta),\end{aligned}$$

which is maximized by

$$q_0(\theta, \omega) = \mathbf{1}[\omega \leq \psi(\theta)]$$

and we can check that for $\theta \geq \psi^{-1}(0)$, $Q_0(\theta) = G(\psi(\theta))$ while for $\theta < \psi^{-1}(0)$ we $Q_0(\theta) = 0$ so that $Q_0(\theta)$ is weakly increasing. Now, for $\theta \geq \psi^{-1}(0)$

$$T_0(\theta) = \theta G(\psi^+(\theta)) - \int_0^\theta G(\psi^+(y)) dy$$

We can check that if $t_0(\theta, \omega) = \psi^{-1}(\omega) \mathbf{1}[\omega \leq \psi(\theta)]$ then whenever $\theta \geq \psi^{-1}(0)$ where

$$\begin{aligned} \int_0^1 t_0(\theta, \omega) dG(\omega) &= \int_0^1 \psi^{-1}(\omega) \mathbf{1}[\omega \leq \psi(\theta)] dG(\omega) \\ &= \int_0^{\psi(\theta)} \psi^{-1}(\omega) dG(\omega) = \theta G(\psi(\theta)) - \int_0^{\psi(\theta)} G(\omega) d\psi^{-1}(\omega) \\ &= \theta G(\psi(\theta)) - \int_0^\theta G(\psi(y)) dy \end{aligned}$$

where the last integral obtains by defining $y = \psi^{-1}(\omega)$. If $\theta < \psi^{-1}(0)$ then $t_0(\theta, \omega) = T_0(\theta) = 0$ for all ω and the proof is complete. \square

A.2.2 Proof of Corollary 1

Proof. That $\lambda(\theta)$ is differentiable for $\theta > \theta^*$ follows immediately from inspection. Note that $\lambda(\theta) = 0$ for all $\theta \leq \theta^*$ so the function is continuous and has a well defined right derivative at θ^* . For the first property we have

$$\begin{aligned} \int_0^1 \lambda(\theta) dF(\theta) &= \int_0^1 (\theta - \psi^+(\theta)) G(\psi^+(\theta)) dF(\theta) - \int_0^1 \int_0^\theta G(\psi^+(y)) dy dF(\theta) \\ &= \int_0^1 (\theta - \psi^+(\theta)) G(\psi^+(\theta)) dF(\theta) - \int_0^1 \int_y^1 G(\psi^+(y)) dF(\theta) dy \\ &= \int_0^1 (\theta - \psi^+(\theta)) G(\psi^+(\theta)) dF(\theta) - \int_0^1 G(\psi^+(\theta)) \left(\frac{1 - F(\theta)}{f(\theta)} \right) dF(\theta) \\ &= \int_0^1 \left(\theta - \frac{1 - F(\theta)}{f(\theta)} - \psi^+(\theta) \right) G(\psi^+(\theta)) dF(\theta) \\ &= \int_{\theta^*}^1 \left(\theta - \frac{1 - F(\theta)}{f(\theta)} - \psi(\theta) \right) G(\psi(\theta)) dF(\theta) = 0 \end{aligned}$$

For the second property, we slightly abuse notation by writing $\lambda'_\theta(\theta^*)$ for the right-derivative at θ^* and write:

$$\begin{aligned} \lambda'_\theta(\theta^*) &= G(\psi(\theta^*)) \left(1 - \psi'_\theta(\theta^*) \right) + (\theta^* - \psi(\theta^*)) g(\psi(\theta^*)) \psi'_\theta(\theta^*) - G(\psi(\theta^*)) \\ &= \theta^* g(\psi(\theta^*)) \psi'_\theta(\theta^*) = \theta^* g(0) \psi'_\theta(\theta^*) > 0 \end{aligned}$$

since ψ is strictly increasing. Finally,

$$\lambda(1) = (1 - 1) G(1) - \int_0^1 G(\psi^+(y)) dy = - \int_{\theta^*}^1 G(\psi(y)) dy < 0$$

\square

A.3 Proofs for Section 4

A.3.1 Proof of Lemma 2

Proof. It is obvious that (4) is a special case of (6) with a fully revealing garbling device $m(\theta) = \theta$, therefore the solution of (4) is within the constraint set of (6) and the optimum must be weakly worse. For the other direction, note that when we denote $\bar{q}_0(\theta, \omega) = q_0(m(\theta), \omega)$ and $\bar{t}_0(\theta, \omega) = t_0(m(\theta), \omega)$, then we can rewrite (6) as

$$\max_{(\bar{q}_0, \bar{t}_0)} \int_0^1 \int_0^1 [\bar{t}_0(\theta, \omega) - \omega \bar{q}_0(\theta, \omega)] dG(\omega) dF(\theta) \quad (10)$$

$$\begin{aligned} \int_0^1 [\theta \bar{q}_0(\theta, \omega) - \bar{t}_0(\theta, \omega)] dG(\omega) &\geq 0, \forall \theta \\ \int_0^1 [\theta \bar{q}_0(\theta, \omega) - \bar{t}_0(\theta, \omega)] dG(\omega) &\geq \int_0^1 [\theta \bar{q}_0(\theta', \omega) - \bar{t}_0(\theta', \omega)] dG(\omega), \forall \theta, \theta' \\ \bar{t}_0(\theta, \omega) - \omega \bar{q}_0(\theta, \omega) &\geq \bar{t}_0(\theta, \omega') - \omega \bar{q}_0(\theta, \omega'), \forall \theta, \omega, \omega' \end{aligned}$$

subject to the constraint that \bar{q}_0 and \bar{t}_0 must be such that there exists a garbling device m and functions q_0 and t_0 such that $\bar{q}_0(\theta, \omega) = q_0(m(\theta), \omega)$ and $\bar{t}_0(\theta, \omega) = t_0(m(\theta), \omega)$ for all θ, ω . Clearly, if we ignore this last constraint, we relax the problem. But this gives us (4). \square

A.3.2 Proof of Proposition 3

Proof. Suppose we have a mechanism in \mathcal{M}^{EAP} that is fully-revealing and solves (4). We can thus focus on choosing $q_0(\theta, \omega)$ and $t_0(\theta, \omega)$. By standard arguments, q_0 must be weakly decreasing in ω . As q_0 is deterministic, for each θ , there exists a marginal cost type $\underline{\omega}(\theta)$ such that $q_0(\theta, \omega) = \mathbf{1}[\omega \leq \underline{\omega}(\theta)]$. Let $\tilde{\theta} = \inf\{\theta : \underline{\omega}(\theta) = 1\}$. By monotonicity, $\underline{\omega}(\theta) = 1$ for all $\theta > \tilde{\theta}$ and $\underline{\omega}(\theta) < 1$ for all $\theta < \tilde{\theta}$. As $\theta = \tilde{\theta}$ is a zero-probability event, we can without loss of generality assume that $\underline{\omega}(\tilde{\theta}) = 1$.

From (ICS) we get that

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

There are two cases. First, for $\theta < \tilde{\theta}$, $\underline{\omega}(\theta) < 1$ and so $q_0(\theta, 1) = 0$. Then by the price mechanism assumption, $t_0(\theta, 1) = 0$. Hence:

$$t_0(\theta, \omega) = \underline{\omega}(\theta) \mathbf{1}[\omega \leq \underline{\omega}(\theta)].$$

This implies that for all $\theta < \tilde{\theta}$, $Q_0^S(\theta) = \int_0^1 q_0(\theta, \omega) dG(\omega) = G(\underline{\omega}(\theta))$ and thus $T_0^S(\theta) = \int_0^1 t_0(\theta, \omega) dG(\omega) = \underline{\omega}(\theta)G(\underline{\omega}(\theta))$. On the other hand, from (IRB) and (ICB), by standard arguments, we get that $Q_0^B(\theta) = G(\underline{\omega}(\theta))$ must be weakly increasing, so that $\underline{\omega}(\theta)$ must be weakly increasing, and we must have

$$T_0^B(\theta) = \theta Q_0^B(\theta) - \int_0^\theta Q_0^B(x) dx = \theta G(\underline{\omega}(\theta)) - \int_0^\theta G(\underline{\omega}(x)) dx.$$

Then $T_0^S(\theta) = T_0^B(\theta)$ gives us a condition

$$[\theta - \underline{\omega}(\theta)]G(\underline{\omega}(\theta)) = \int_0^\theta G(\underline{\omega}(x)) dx.$$

This equality immediately implies that $\underline{\omega}(\theta) \leq \theta$ for all such θ .

We denote $p_0(\theta) = \underline{\omega}(\theta)$ for all $\theta < \tilde{\theta}$. Note that as types $\theta \geq \tilde{\theta}$ choose the guaranteed delivery option in equilibrium, we can set $p_0(\theta)$ for $\theta \geq \tilde{\theta}$ to any value, with a restriction that each type $\theta \geq \tilde{\theta}$ prefers to choose the guaranteed delivery option to any at-will offer $p_0(\theta)$. One convenient way to guarantee this is to choose p_0 so that it satisfies the previous equation for all $\theta \in \Theta$. This is (7).

Next, for all $\theta \geq \tilde{\theta}$, by construction $\underline{\omega}(\theta) = 1$ so that $q_0(\theta, 1) = q_0(\theta, \omega) = 1$ for all ω and thus

$$t_0(\theta, \omega) = \omega + \int_\omega^1 dy + t_0(\theta, 1) - 1 = t_0(\theta, 1).$$

For all such θ , $Q_0^B(\theta) = \int_0^1 q_0(\theta, \omega) dG(\omega) = 1$ and thus using (ICB) we get

$$T_0^B(\theta) = t_0(\theta, 1) = \theta - \int_0^{\tilde{\theta}} Q_0^B(x) dx - \int_{\tilde{\theta}}^\theta 1 dx = \tilde{\theta} - \int_0^{\tilde{\theta}} Q_0^B(x) dx = T_0^B(\tilde{\theta}).$$

We denote $P_0 = T_0^B(\tilde{\theta})$. Note, as there is a discontinuous jump in the allocation rule at $\tilde{\theta}$, we need to determine this value directly from the (ICB) constraints. As the buyer type $\tilde{\theta}$ must be indifferent between guaranteed delivery at price P and at-will offer at price $p_0(\tilde{\theta})$, we get

$$\tilde{\theta} - P_0 = [\tilde{\theta} - p_0(\tilde{\theta})]G(p_0(\tilde{\theta})) \iff P_0 = p_0(\tilde{\theta})G(p_0(\tilde{\theta})) + \tilde{\theta}[1 - G(p_0(\tilde{\theta}))].$$

This is (8).

Finally, we need to show that $\tilde{\theta} < 1$. Suppose not, i.e., $\tilde{\theta} = 1$. Then the profit from

the optimal mechanism is

$$V = \int_0^1 \int_0^{p_0(\theta)} (p_0(\theta) - \omega) dG(\omega) dF(\theta).$$

Let us construct an alternative mechanism, with the same price function $p_0(\theta)$, but such that a guaranteed delivery offer \widehat{P} is chosen by all types $1 - \varepsilon$ for some small $\varepsilon > 0$. By the analysis above,

$$\widehat{P} = p_0(1 - \varepsilon)G(p_0(1 - \varepsilon)) + (1 - \varepsilon)[1 - G(p_0(1 - \varepsilon))].$$

Therefore the new profit can be computed as

$$\widehat{V} = \int_0^{1-\varepsilon} \int_0^{p_0(\theta)} (p_0(\theta) - \omega) dG(\omega) dF(\theta) + \int_{1-\varepsilon}^1 \int_0^1 (\widehat{P} - \omega) dG(\omega) dF(\theta)$$

The difference between the two profits is

$$\widehat{V} - V = [1 - F(1 - \varepsilon)]\widehat{P} - \int_{1-\varepsilon}^1 p_0(\theta)G(p_0(\theta))dF(\theta) - \int_{1-\varepsilon}^1 \int_{p_0(\theta)}^1 \omega dG(\omega) dF(\theta).$$

As p_0 is weakly increasing, $p_0(\theta)G(p_0(\theta)) \leq p_0(1)G(p_0(1))$ for all $\theta \in [1 - \varepsilon, 1]$. Also by monotonicity of p_0 ,

$$\int_{p_0(\theta)}^1 \omega dG(\omega) \leq \int_{p_0(1-\varepsilon)}^1 \omega dG(\omega).$$

Combining these facts, we get

$$\frac{\widehat{V} - V}{\varepsilon} \geq \frac{1 - F(1 - \varepsilon)}{\varepsilon} \left(\widehat{P} - p_0(1)G(p_0(1)) \right) - \frac{1 - F(1 - \varepsilon)}{\varepsilon} \int_{p_0(1-\varepsilon)}^1 \omega dG(\omega).$$

Note that by (7), $p_0(1) < 1$. Taking limits as $\varepsilon \rightarrow 0$, we get that $\frac{1-F(1-\varepsilon)}{\varepsilon} \rightarrow f(1) > 0$, that

$$\widehat{P} - p_0(1)G(p_0(1)) \rightarrow p_0(1)G(p_0(1)) + 1 - G(p_0(1)) - p_0(1)G(p_0(1)) = \int_{p_0(1)}^1 1dG(\omega) > 0,$$

and

$$\int_{p_0(1-\varepsilon)}^1 \omega dG(\omega) \rightarrow \int_{p_0(1)}^1 \omega dG(\omega) > 0.$$

Combining these facts,

$$\lim_{\varepsilon \rightarrow 0} \frac{\widehat{V}(\varepsilon) - V}{\varepsilon} = f(1) \int_{p_0(1)}^1 (1 - \omega) dG(\omega) > 0.$$

Therefore, for a sufficiently small $\varepsilon > 0$, $\widehat{V}(\varepsilon) > V$, which is a contradiction with the assumption that the original mechanism is optimal. This proves that, $\tilde{\theta} < 1$, i.e., the guaranteed delivery offer is chosen with positive probability. \square

A.3.3 Proof of Corollary 3

Proof. As discussed above, the profit from any mechanism satisfying (IRB) and (ICB), can be computed as

$$V = \int_0^1 \int_0^1 [\psi(\theta) - \omega] q(\theta, \omega) dG(\omega) dF(\theta).$$

Specifically, the profit from the benchmark mechanism is

$$V^B = \int_0^1 \int_0^{\psi(\theta)} (\psi(\theta) - \omega) dG(\omega) dF(\theta)$$

and the profit from the optimal ex-ante price mechanism is

$$V^{EAO} = \int_0^{\tilde{\theta}} \int_0^{p_0(\theta)} (\psi(\theta) - \omega) dG(\omega) dF(\theta) + \int_{\tilde{\theta}}^1 \int_0^1 (\psi(\theta) - \omega) dG(\omega) dF(\theta).$$

Note that as for all θ ,

$$\int_0^{\psi(\theta)} (\psi(\theta) - \omega) dG(\omega) \geq \int_0^{p_0(\theta)} (\psi(\theta) - \omega) dG(\omega),$$

the difference

$$V^B - V^{EAO} \geq - \int_{\tilde{\theta}}^1 \int_{\psi(\theta)}^1 (\psi(\theta) - \omega) dG(\omega) dF(\theta) > 0,$$

as $\tilde{\theta} < 1$. \square

A.3.4 Proof of Corollary 4

Proof. We can write the expected profits for the seller with a guaranteed delivery price P_0 as

$$[1 - F(P_0)](P_0 - \mathbb{E}[\omega]) = [1 - F(P_0)](P_0 - 1 + \mathcal{G}(1)).$$

Now, the profits for a mechanism where we still have a single price P_0 but with *at-will* delivery are

$$\int_{P_0}^1 \int_0^{P_0} (P_0 - \omega) dG(\omega) dF\theta = [1 - F(P_0)] \mathcal{G}(P_0).$$

The difference between the latter and the former is therefore:

$$[1 - F(P_0)] [\mathcal{G}(P_0) - P_0 + 1 - \mathcal{G}(1)] = [1 - F(P_0)] \int_{P_0}^1 (1 - G(x)) dx > 0.$$

So, for any P_0 that maximizes the seller's profits with guaranteed delivery, the seller can do better with a mechanism where she uses the same price but with at-will delivery. But this is bettered, as established by Proposition 3, by a menu of offers and the result follows. \square

A.3.5 Proof of Proposition 4

Proof. Clearly, buyers with highest types have the strongest incentive to accept the first offer, so there must be a marginal type $\bar{\theta}$ who is indifferent between accepting and rejecting the first offer.

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

$$V_1^D(\bar{\theta}) = \int_0^{\bar{\theta}} \int_0^{\bar{\theta}} \left[q_1(\theta, \omega) - \int_0^{\theta} q_1(y, \omega) dy - \omega q_1(\theta, \omega) \right] \frac{f(\theta)}{F(\bar{\theta})} d\theta dG(\omega),$$

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

$$V^D(\bar{\theta}, \delta) = (1 - F(\bar{\theta})) \int_0^1 (p_0 - \omega) dG(\omega) + \delta F(\bar{\theta}) V_1^D(\bar{\theta}),$$

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

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

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 standard steps to get $q_1(\theta, \omega) = \mathbf{1} [\omega < \psi(\theta, \bar{\theta})]$, where $\psi(\theta, \bar{\theta}) = \theta - \frac{F(\bar{\theta}) - F(\theta)}{f(\theta)}$, 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,

$$\begin{aligned} V_1^D(\bar{\theta})F(\bar{\theta}) &= \int_0^{\bar{\theta}} \int_0^{\bar{\theta}} \left[q_1(\theta, \omega) - \int_0^{\theta} q_1(y, \omega) dy - \omega q_1(\theta, \omega) \right] f(\theta) d\theta dG(\omega) \\ &= \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} (\theta f(\theta) - F(\bar{\theta}) + F(\theta)) G(\psi(\theta, \bar{\theta})) d\theta \\ &\quad - \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} f(\theta) [\psi(\theta, \bar{\theta}) G(\psi(\theta, \bar{\theta})) - \mathcal{G}(\psi(\theta, \bar{\theta}))] d\theta \\ &= \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} \mathcal{G}(\psi(\theta, \bar{\theta})) f(\theta) d\theta, \end{aligned}$$

Hence,

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

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

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

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

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

where $\xi(\bar{\theta}) = \int_{\theta^{**}(\bar{\theta})}^{\bar{\theta}} G(\psi(\theta, \bar{\theta})) d\theta$, so that $\xi'(1) = 1 - f(1) \int_{\theta^*}^1 \frac{g(\psi(\theta, 1))}{f(\theta)} d\theta \leq 1 < \infty$ and therefore $(1 - F(1))\xi'(1) = 0$. So, for any $\delta < 1$, $V^D(\bar{\theta}, \delta)$ is decreasing at $\bar{\theta} = 1$ and this implies that it is optimized for $\bar{\theta} < 1$. Since $V^D(1, \delta) = V^{EPO}$ then this proves the result. \square

A.4 Proofs for Section 5

A.4.1 Proof of Proposition 5

Proof. Define the set

$$\widehat{\Theta}(c) = \{\theta : \exists \omega \text{ such that } q_0(\theta, \omega) > 0\},$$

then the problem faced by the seller is a simple modification of the benchmark problem where the constraints for the buyer are the same but the objective function is now:

$$\int_0^1 \int_0^1 [t_0(\theta, \omega) - \omega q_0(\theta, \omega)] dF(\theta) dG(\omega) - \int_{\widehat{\Theta}(c)} c dF(\theta).$$

Since the buyer's constraints are unchanged, using the same procedure as in appendix A.2.1 we can rewrite the objective as

$$\int_0^1 \int_0^1 [\psi(\theta) - \omega] q_0(\theta, \omega) dF(\theta) dG(\omega) - \int_{\widehat{\Theta}(c)} c dF(\theta).$$

As the seller observes θ before choosing whether to learn her cost, we can maximize this for any given θ . That is, for any type θ , the seller will have to decide whether to make trade possible for some ω . This will require spending the learning cost c . The other option is not to spend the learning cost and not trading. Hence, the optimal q_0 maximizes:

$$\max \left(\int_0^1 [\psi(\theta) - \omega - c] q_0(\theta, \omega) dG(\omega), 0 \right) \quad \forall \theta \in \Theta.$$

This immediately gives us:

$$\begin{aligned} q_0(\theta, \omega) &= \mathbf{1}[\omega \leq \psi(\theta)] \cdot \mathbf{1} \left[\int_0^{\psi(\theta)} (\psi(\theta) - \omega - c) dG(\omega) \geq 0 \right] \\ &= \mathbf{1}[\omega \leq \psi(\theta)] \cdot \mathbf{1}[\mathcal{G}(\psi(\theta)) \geq c] = \mathbf{1}[\max(\omega, \mathcal{G}^{-1}(c)) \leq \psi(\theta)]. \end{aligned}$$

It is easy to see that this means that $q_0 = 0$ whenever $c > \mathcal{G}(1)$. The rest follows as in the benchmark case noting that $\widehat{\Theta}(c) = \{\theta : \theta \geq \psi^{-1}(\mathcal{G}^{-1}(c))\}$. \square

A.4.2 Proof of Proposition 7

Proof. The proof follows along the lines of the proof of proposition 5. So define $\overline{\Theta}(c)$ as the set of buyer types for which the seller will be willing to invest in learning her type. Contrary to the proof of proposition 5, however, we will not define it ex-ante as we did with $\widehat{\Theta}$. It is immediate, however, $\overline{\Theta}(c) \subseteq \widehat{\Theta}(c)$ for any c , as the seller no longer needs to learn her cost in order to trade. Given this, and given that again we are maximizing for any given θ , we have that whenever $\theta \in \overline{\Theta}(c)$, the seller's expected utility taking into account all the constraints for the buyer is:

$$\int_0^1 [q_0(\theta, \omega)(\psi(\theta) - \omega) - c] dG(\omega) = \int_0^1 q_0(\theta, \omega)(\psi(\theta) - \omega) dG(\omega) - c,$$

which, as we've seen in the proof of proposition 5 is optimized by $q_0(\theta, \omega) = \mathbf{1}[\max(\omega, \mathcal{G}^{-1}(c)) \leq \psi(\theta)]$. If $\theta \notin \overline{\Theta}(c)$ we have:

$$\int_0^1 q_0(\theta)(\psi(\theta) - \omega) dG(\omega),$$

where, with a slight abuse of notation, we have emphasized that in this latter case, q_0 cannot be a function of ω as, by definition, this is not learned. This is clearly maximized by $q_0(\theta) = q_0(\theta, \omega) = \mathbf{1}[\mathbb{E}(\omega) \leq \psi(\theta)]$.

For any θ such that the mechanism gives no probability of trade, it is clearly optimal not to learn ω . Therefore, we now assume that $\theta \geq \psi^{-1}(0)$ as this is the largest possible set of buyer types with whom the seller may wish to trade. For any such θ , define the function:

$$\begin{aligned}\chi(\theta) &= \int_0^{\psi(\theta)} (\psi(\theta) - \omega) dG(\omega) - \int_0^1 (\psi(\theta) - \omega) dG(\omega) = - \int_{\psi(\theta)}^1 (\psi(\theta) - \omega) dG(\omega) \\ &= \mathbb{E}(\omega) + \mathcal{G}(\psi(\theta)) - \psi(\theta).\end{aligned}$$

The $\chi(\theta)$ function tells us, for any θ for which the mechanism gives a positive probability of trade, the gain for the seller from learning her cost. It is easy to check that for any $\theta \in [\psi^{-1}(0), 1)$ this is a positive function because $\mathbb{E}(\omega) = 1 - \mathcal{G}(1)$ so that

$$\chi(\theta) = \int_{\psi(\theta)}^1 [1 - G(x)] dx > 0.$$

Further, $\chi(\psi^{-1}(0)) = \mathbb{E}(\omega)$, $\chi(1) = 0$, and $\chi(\theta)$ is strictly decreasing. The seller will wish to learn her type for any θ such that $\chi(\theta) \geq c$ and $\mathcal{G}(\psi(\theta)) \geq c$. The latter condition tells us that for this θ , the seller prefers to trade and learn her type over not trading, whereas the former guarantees that she also prefers to trade learning her type (and paying the cost of learning) over trading without learning her type. The function $\mathcal{G}(\psi(\theta))$ is strictly positive in the interval $\theta \in (\psi^{-1}(0), 1]$, and is strictly increasing with $\mathcal{G}(\psi(\psi^{-1}(0))) = 0$ and $\mathcal{G}(\psi(1)) = \mathcal{G}(1)$. Thus, the two functions must cross exactly once at some point $\hat{\theta}$. But if $\chi(\hat{\theta}) = \mathcal{G}(\psi(\hat{\theta}))$, then it is immediate to see that it must also be that $\psi(\hat{\theta}) = \mathbb{E}(\omega) \iff \hat{\theta} = \psi^{-1}(\mathbb{E}(\omega))$.

Now, fix c . Given the above, the set $\bar{\Theta}(c)$ is non-empty whenever $c \leq \mathcal{G}(\mathbb{E}(\omega))$. If so,

$$\bar{\Theta}(c) = \{\theta : \chi(\theta) \geq c \text{ and } \mathcal{G}(\psi(\theta)) \geq c\} = [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)],$$

with $\hat{\theta} \in [\psi^{-1}(\mathcal{G}^{-1}(c)), \chi^{-1}(c)]$. Clearly, $\psi^{-1}(\mathcal{G}^{-1}(c))$ is strictly increasing in c with $\psi^{-1}(\mathcal{G}^{-1}(0)) = \psi^{-1}(0)$ and $\psi^{-1}(\mathcal{G}^{-1}(\mathbb{E}(\omega))) = \hat{\theta}$ while $\chi^{-1}(c)$ is strictly decreasing with $\chi^{-1}(0) = 1$ and $\chi^{-1}(\mathcal{G}(\mathbb{E}(\omega))) = \hat{\theta}$.

If $c > \mathcal{G}(\mathbb{E}(\omega))$, then for all $\theta \geq \psi^{-1}(0)$, at least one of the two conditions for the seller wanting to learn her type is not satisfied. The seller then, knows ex-ante that she will never want to learn her type and chooses a mechanism that does not depend on it,

but only on the buyer's type. The optimal mechanism is then to trade whenever the buyer's virtual valuation is no smaller than the seller's *expected* type, as shown above. The corresponding transfers follow immediately as in proposition 1 with $\mathbb{E}(\omega)$ instead of ω .

If $c \leq \mathcal{G}(\mathbb{E}(\omega))$, instead, the interval is well-defined and so the seller will choose the allocation rule $q_0(\theta, \omega) = \mathbf{1}[\max(\omega, \mathcal{G}^{-1}(c)) \leq \psi(\theta)]$ whenever $\theta \in \bar{\Theta}(c)$ which, by construction of $\bar{\Theta}(c)$, becomes $\mathbf{1}[\omega \leq \psi(\theta)]$. If $\theta > \chi^{-1}(c)$, because this implies $\theta > \hat{\theta} = \psi^{-1}(\mathbb{E}(\omega))$, the seller will always trade the good. Noting that the monotonicity conditions on either $q_0(\theta)$ (when there is no learning) or $Q_0^B(\theta)$, $q_0(\theta, \omega)$ (when there is learning) apply and this proves the optimal allocation.

As for transfers, given that our constraints only put restrictions on $T_0^B(\theta)$, it is without loss to have $t_0(\theta, \omega) = T_0^B(\theta)$ whenever a buyer type has probability 0 or 1 of trade and to set

$$t_0(\theta, \omega) = \frac{T_0^B(\theta)}{G(\psi(\theta))} \mathbf{1}[\omega \leq \psi(\theta)],$$

whenever trade also depends on the realization of ω (that is, whenever $\theta \in \bar{\Theta}(c)$). Given the buyer's constraints we then have that $T_0^B(\theta) = \theta Q_0^B(\theta) - \int_0^\theta Q_0^B(y) dy$. It is immediate that if $\theta < \psi^{-1}(\mathcal{G}^{-1}(c))$ then $t_0(\theta, \omega) = T_0^B(\theta) = 0$, whereas if $\theta > \chi^{-1}(c)$ then $t_0(\theta, \omega) = T_0^B(\theta) = \chi^{-1}(c) - \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^{\chi^{-1}(c)} G(\psi(y)) dy$. Finally, if $\theta \in \bar{\Theta}(c)$, then $T_0^B(\theta) = \theta - \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^\theta G(\psi(y)) dy$ which immediately implies the transfers $t_0(\theta, \omega) = \frac{T_0^B(\theta)}{G(\psi(\theta))} \mathbf{1}[\omega \leq \psi(\theta)]$ given in the proposition.

The formulas for V^{BOC} are immediate. As for the comparison with V^{BNC} , note that if $c > \mathcal{G}(1)$ then the former is positive (the seller offers trade with guaranteed delivery at a fixed price $\psi^{-1}(\mathbb{E}(\omega))$) while the latter is zero (the seller doesn't trade). If $\mathcal{G}(1) \geq c > \mathcal{G}(\mathbb{E}(\omega))$ then we have that $\psi^{-1}(\mathcal{G}^{-1}(c))$ is still well-defined and, given the above:

$$\begin{aligned} V^{BOC} &= \int_{\psi^{-1}(\mathbb{E}(\omega))}^1 \int_0^1 (\psi(\theta) - \omega) dG(\omega) dF(\theta) \\ &> \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^1 \int_0^{\psi(\theta)} (\psi(\theta) - \omega - c) dG(\omega) dF(\theta) \\ &> \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^1 -c + \left(\int_0^{\psi(\theta)} (\psi(\theta) - \omega - c) dG(\omega) \right) dF(\theta) = V^{BNC}, \end{aligned}$$

where the first inequality comes from the fact that since $c > \mathcal{G}(\mathbb{E}(\omega))$, then when $\mathcal{G}(\psi(\theta)) \geq$

c , necessarily $\chi(\theta) < c$. Finally, if $c \leq \mathcal{G}(\mathbb{E}(\omega))$ then

$$\begin{aligned} V^{BOC} &> \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^1 \int_0^{\psi(\theta)} (\psi(\theta) - \omega - c) dG(\omega) dF(\theta) \\ &> \int_{\psi^{-1}(\mathcal{G}^{-1}(c))}^1 \left(-c + \int_0^{\psi(\theta)} (\psi(\theta) - \omega) dG(\omega) \right) dF(\theta) = V^{BNC}. \end{aligned}$$

Now, the first inequality comes from simple inspection of V^{BOC} . □

A.4.3 Proof of Proposition 9

Proof. We only need to check that the ex-post optimal allocation and the transfers $t_0(\theta, \omega) = \theta \mathbf{1}[\omega \leq \theta] - \mathcal{G}(\theta)$ satisfy all the constraints. The usual monotonicity conditions on $Q_0^B(\theta)$ and $q_0(\theta, \omega)$ are clearly satisfied. We then have

$$\begin{aligned} T_0^S(\theta) &= \int_0^1 t_0(\theta, \omega) dG(\omega) = v(\theta, 1) + \int_0^1 \left(\int_{\omega}^1 \mathbf{1}[x \leq \theta] dx + \omega \mathbf{1}[\omega \leq \theta] \right) dG(\omega) \\ &= \lambda^E(\theta) + \int_0^1 \mathbf{1}[\omega \leq \theta] G(\omega) d\omega + \int_0^1 \omega \mathbf{1}[\omega \leq \theta] dG(\omega) \\ &= \lambda^E(\theta) + \mathcal{G}(\theta) + \theta G(\theta) - \mathcal{G}(\theta) = \lambda^E(\theta) + \theta G(\theta), \end{aligned}$$

whereas

$$\begin{aligned} T_0^B(\theta) &= \theta Q_0^B(\theta) - \int_0^{\theta} Q_0^B(y) dy = \theta \int_0^1 \mathbf{1}[\omega \leq \theta] dG(\omega) - \int_0^{\theta} \int_0^1 \mathbf{1}[\omega \leq y] dG(\omega) dy \\ &= \theta G(\theta) - \int_0^{\theta} G(y) dy = \theta G(\theta) - \mathcal{G}(\theta), \end{aligned}$$

and, of course, equating them gives us $\lambda^E(\theta) = -\mathcal{G}(\theta)$. To check that this is individually rational for the seller we can show that

$$\begin{aligned} \int_0^1 v(\theta, \omega) dG(\omega) &= \int_0^1 [t_0(\theta, \omega) - \omega q_0(\theta, \omega)] dG(\omega) = \int_0^{\theta} (\theta - \omega) dG(\omega) - \mathcal{G}(\theta) \\ &= \theta G(\theta) - [\theta G(\theta) - \mathcal{G}(\theta)] - \mathcal{G}(\theta) = 0, \end{aligned}$$

which means that the “reverse” individual rationality condition for the seller is satisfied with equality for all buyer types. This, of course, also implies the seller’s ex-ante individual rationality. □