

Information Design in the Hold-up Problem*

Daniele Condorelli[†] and Balázs Szentes[‡]

November 15, 2018

Abstract

We analyze a bilateral trade model where the buyer can choose the probability distribution of her valuation for the good. The seller, after observing the buyer's choice of the distribution but not the realized valuation, makes a take-it-or-leave-it offer. If the buyer's choice of the distribution is costless, the price and the payoffs of both the buyer and the seller are shown to be $1/e$ in the unique equilibrium outcome. The equilibrium distribution of the buyer's valuation generates a unit-elastic demand and trade is ex-post efficient. These two properties are shown to be preserved even when different distributions are differentially costly as long as the cost is monotone in the dispersion of the distribution.

*We have benefited from discussions with Dirk Bergemann, Ben Brooks, Jeremy Bulow, Sylvain Chassang, Eddie Dekel, Andrew Ellis, Jeff Ely, Faruk Gul, Yingni Guo, Sergiu Hart, Johannes Horner, Emir Kamenica, Francesco Nava, Santiago Oliveros, Ludovic Renou, Phil Reny and Andreas Uthemann. The comments of several seminar audiences are also gratefully acknowledged.

[†]Department of Economics, University of Essex, Colchester, UK. E-mail: d.condorelli@gmail.com.

[‡]Department of Economics, London School of Economics, London, UK. E-mail: b.szentes@lse.ac.uk.

1 Introduction

A central result in Information Economics is that the distribution of economic agents' private information is a key determinant of their welfare. In auctions, for example, asymmetric information between the seller and the bidders regarding bidders' valuations affects all agents' payoffs. In the context of labour markets, asymmetric information between the employers and workers regarding workers' productivities affects workers' wages. When analysing environments with asymmetric information, most microeconomic models assume that the distribution of types is exogenous. However, if an agent's payoff depends on her information rent then she is likely to take actions in order to generate this rent optimally. This paper contributes towards this problem by reconsidering the hold-up problem from the perspective of optimal information design.

The literature on hold-up typically focuses on environments where the buyer's investment decision has a deterministic impact on her valuation for the seller's good. If the seller has full bargaining power, he appropriates all the surplus generated by investments. Hence, if investments are costly, the buyer does not invest at all and ends up with a payoff of zero. In contrast, if investments have privately observed stochastic returns, the buyer's investment decision induces asymmetric information at the contracting stage. Therefore, the buyer may capture some of the surplus through receiving information rent which, in turn, could partially restore her incentives to invest on the first place. Our objective is to analyze the buyer's problem of optimally designing the investment's stochastic returns.

For example, consider bargaining between a patent holder (seller) and an innovative firm (buyer). The firm intends to develop a product that requires patent-protected technology. Prior to contracting, the firm implements an investment strategy which determines its future profits from bringing the product to market and hence, its valuation for the license of the patent holder.¹ The patent holder may detect this strategy perhaps through observing a prototype, the size of investment and capital goods purchased by the firm. However, it is conceivable that the firm has private information regarding the realized profitability generated by its investments. Then the contract offered by the patent holder may depend on the investment strategy but not on the actual profitability. Of course, if the choice of investment perfectly reveals the profitability, the patent holder's optimal contract might hold the firm down to its reservation payoff — a phenomenon often referred to as “patent hold-up”, see, for example, Shapiro (2010).² Therefore, the firm may benefit from implementing a less effective investment strategy which, even when detected, leaves the patent holder uncertain about the realized profitability.

¹Before the investment stage, the patent holder is unaware of the firm's intentions regarding the use of the license and, in the absence of fully contingent contracts, negotiations may not be practical. Indeed, a substantial share of patent licensing is conducted after investments have already been made, see Siebert (2015).

²This is different from the so called “patent ambush” whereby the firm is unaware that the technology is patent protected and the holder withholds this information until the technology becomes widely adopted.

In order to focus on the buyer's incentive to generate information rent, we first abstract from the cost of investment and consider a benchmark model where distributions are costless and the only constraint faced by the buyer is that the maximum valuation is bounded. Specifically, the buyer can choose any CDF with support in $[0, 1]$ which then determines her valuation for the object. After observing the buyer's CDF, the seller sets a price at which either the buyer trades or the game ends. The central trade-off arising in this model can be explained as follows. If the price of the good was given, the buyer would be better off the higher is her valuation. However, the seller who knows that the buyer's valuation is likely to be high will charge a higher price. As a result, when choosing her value-distribution, the buyer faces a trade-off between the higher payoff she would receive from the good and the higher price she would have to pay.

We show that this game has a unique equilibrium outcome. In each equilibrium, the price is $1/e$ and the buyer's distribution is supported on $[1/e, 1]$, so that trade always occurs and the seller's payoff is $1/e$. The buyer's value-distribution is a combination of a continuous distribution on $[1/e, 1)$ defined by the CDF $F^*(v) = 1 - 1/(ev)$ and an atom of size $1/e$ at $v = 1$. The total equilibrium surplus is found to be shared equally between the seller and the buyer, so that the buyer also receives an equilibrium payoff of $1/e$. The efficiency loss that results from the buyer's desire to generate information rent is found to be around one quarter of the first-best social surplus.

Our baseline model provides a rationale for the endogenous emergence of unit-elastic demands. Indeed, the equilibrium CDF of the buyer, F^* , generates a unit-elastic demand function on $(1/e, 1)$, such that the probability that the buyer is willing to buy the good at price p is $1 - F^*(p) = 1/(ep)$. When faced with this demand function, the seller is indifferent between charging any price in $[1/e, 1]$, since any price in this range will result in an expected profit of $p[1/(ep)] = 1/e$. We demonstrate that this result arises because the buyer can always increase her payoff by choosing a different CDF unless the seller is indifferent between any price on the support of the buyer's CDF. If the latter does not hold, the buyer can move probability mass from suboptimal prices to higher valuations without changing the seller's incentive to set a certain price.

Our second main insight is that the equilibrium outcome is ex-post efficient, i.e., trade occurs with probability one. As we explained above, given the buyer's equilibrium CDF F^* , the seller finds it optimal to set the price at the lowest possible valuation of the buyer. In contrast, if the buyer's valuation was given exogenously, the seller would often find it optimal to set a price which excludes low-value buyers from trade. However, if a distribution induces a price above the smallest valuation on the support, the buyer can increase her payoff by choosing the same distribution conditional on the value being higher than the price. The reason is that, conditional on trade, the buyer's payoff is the same but she trades more often if she chooses the conditional distribution.

We extend our benchmark model by introducing a cost associated to the buyer's choice of a CDF. It turns out that the crucial property of such a cost for our analysis to carry forward is

monotonicity in risk. More formally, the buyer's cost is said to be *increasing in risk* if a mean-preserving spread of a CDF is more expensive than the CDF itself. To give an example, suppose that each particular valuation has a deterministic cost and that the cost of a CDF F is the expected cost generated by F . As will be explained, this cost function defined over the CDFs is increasing in risk whenever the deterministic cost function defined over the valuations is convex. Our main result is that if the buyer's cost is increasing in risk, there exist equilibria in which the buyer's demand is unit elastic and trade occurs with probability one. Furthermore, if the cost is strictly increasing in risk, all equilibria have these features.

We also consider the case when the buyer's cost is *decreasing in risk*, that is, adding a mean-preserving spread to a CDF reduces the buyer's cost. We demonstrate the existence of equilibria where the buyer's CDF is a convex combination of an atom at zero and an equal-revenue distribution. This means that conditional on the valuation being strictly positive, the buyer's demand is again unit-elastic. The equilibrium price is the smallest strictly positive element of the support. Therefore, trade does not always occur but it is ex-post efficient because the buyer rejects the equilibrium price only if her valuation is zero. In addition, if the cost is strictly decreasing in risk, there is no other kind of equilibrium.

Finally, we discuss how the equilibria of our model are modified if there is a positive production cost, the buyer is not risk-neutral, and the seller does not have full bargaining power.

There are papers on the hold-up problem which also consider the buyer's ability to generate asymmetric information. In the model by Gul (2001) and Lau (2008), the buyer can generate asymmetric information by randomizing over imperfectly observable but deterministic investment decisions. Gul (2001) shows that if investment is unobservable, the hold-up problem is as severe as in the fully observable case.³ In the unique mixed-strategy equilibrium, the efficiency gain from the buyer's randomization over positive investments is fully offset by the ex-post efficiency loss due to bargaining disagreement. Lau (2008) assumes that the seller observes the buyer's investment perfectly with a given probability, and receives an uninformative signal otherwise. Her main result is that efficiency is maximized when this probability is strictly between zero and one. Nevertheless, the buyer does not earn information rent and ends up with a payoff of zero. Hermalin and Katz (2009) consider the hold-up problem incorporating asymmetric information after the buyer's investment decision. In particular, a larger investment results in a first-order stochastic shift in the buyer's value distribution. The authors consider several scenarios in which the degree of observability of the investment is varied and find that the buyer is made worse off if her investment is unobservable.⁴ Finally, Skrzypacz (2004) also emphasizes, as we do, that the desire to generate information rent

³Dilmé (2017) considers a related model, but assumes that the costly investment also increases the outside option of the investor.

⁴The consequences of the observability assumption in hold-up problems were also explored in Tirole (1986) and in Gul (2001) under various bargaining protocols.

can shape the buyer’s incentives to invest. However, the author’s analysis focuses on dynamic bargaining, where Coasian dynamics imply that the buyer obtains full bargaining power as the discount rate vanishes.

Unit-elastic demands also appear in a literature on non-Bayesian monopoly pricing. Bergemann and Schlag (2008) consider a monopolist with the min-max regret criterion. The seller’s ex-post regret is defined as the buyer’s payoff if trade occurs and the buyer’s valuation otherwise. The authors argue that the optimal pricing policy coincides with the seller’s equilibrium strategy in a zero-sum game played against nature in which nature chooses the buyer’s valuation to maximize the seller’s regret. The authors show that nature’s equilibrium strategy in this case is the same as our equilibrium CDF and the seller fully randomizes on its support.⁵ Neeman (2003) considers second-price auctions with private values and addresses the problem of finding the value-distribution which minimizes the ratio between the seller’s profit and the expected value. The author shows that the solution generates a unit-elastic demand function. Unit-elastic demands also play a role in models where randomization by the seller is required for the existence of an equilibrium, since this form of demand function is required to ensure that the seller is indifferent on the support of valuations. Renou and Schlag (2010) consider models with min-max regret and imperfect competition and Hart and Nisan (2012) approximate the seller’s maximum revenue in a multiple-item auction.⁶ Related results also appear in Brooks (2013) and in Kremer and Snyder (2015).

Finally, our paper also relates to a recent literature on information design, see for example Kamenica and Gentzkow (2011). Bergemann and Pesendorfer (2007) consider the seller’s problem of designing the information structure to determine how buyers learn about their valuations prior to participating in an auction. Bergemann, Brooks and Morris (2015) analyze a model where the buyer’s value distribution is given and the seller receives a signal about this value. The authors characterize the entire set of payoff outcomes that can arise from *some* signal structure. Unit-elastic demand plays a crucial role in their construction. Roesler and Szentes (2017) analyze a setup where the buyer has a given value distribution and designs a signal structure at no cost to learn about her valuation. After observing the signal structure, the seller sets a price and the buyer trades if the expected valuation conditional on her signal exceeds the price. The buyer in Roesler and Szentes (2017) faces a trade-off between signal precision improving the efficiency of her purchase decision for a given price, but potentially increasing the price since the buyer’s demand is determined by her signal. As will be shown, the buyer’s problem in Roesler and Szentes (2017) is equivalent to

⁵The seller’s regret is the buyer’s payoff if there is trade, but differs from it in case of no trade. Therefore, while it’s not a coincidence that nature chooses the CDF which maximizes the buyer’s expected surplus, their result does not directly imply ours, and vice-versa.

⁶In a different context, Ortner and Chassang (2014) show that, in order to eliminate collusion between an agent and the monitor, the principal benefits from introducing asymmetric information between the agent and the monitor by making the monitor’s wage random. The optimal wage scheme is determined by a distribution similar to our equilibrium CDF.

the buyer’s problem in our paper if the buyer cost is specified such that a CDF is free if the prior distribution is a mean-preserving spread of it and it is prohibitively costly otherwise. Consistent with our result on cost functions which are increasing in risk, Roesler and Szentes (2017) show that there is an equilibrium in which the buyer always trades and the seller is indifferent between any price in the signal’s support. We show that this result remains valid even if signals are costly and this cost is increasing in the informativeness of the signals.

2 Model

There is a seller who has an object to sell to a single buyer. Prior to interacting with the seller, the buyer can choose the distribution of her valuation subject to the constraint that the valuation is below one. Formally, the buyer can choose any $F \in \mathcal{F}$ where \mathcal{F} is the set of CDFs with a support in $[0, 1]$. The cost of a CDF is given by the function $C : \mathcal{F} \rightarrow \mathbb{R}$. The seller observes the choice of the buyer and gives a take-it-or-leave-it price offer to the buyer, p . Finally, the buyer’s valuation, v , is realized and she decides whether to trade at price p . If trade takes place, the payoff of the seller is p and that of the buyer is $v - p - C(F)$. Otherwise, the buyer’s payoff is $-C(F)$ and the seller’s payoff is zero. For simplicity, we assume that the buyer trades whenever $v \geq p$.⁷ The seller and the buyer are expected payoff maximizers. We restrict attention to Subgame Perfect Nash Equilibria of this game.

In the remainder of this section, we introduce some properties of the function C and discuss them in the context of examples.

The Cost Function.— The key property of C in our analysis is going to be monotonicity in risk. Formally, we say that C is *increasing in risk* if $C(F) \geq C(G)$ whenever F is a mean-preserving spread of G . If this inequality is strict whenever $F \neq G$, we call C *strictly increasing in risk*. Similarly, if $C(F) \leq (<) C(G)$ whenever F is a mean-preserving spread of G and $F \neq G$, the function C is said to be (strictly) decreasing in risk. Finally, we call C *mean-based* if $C(F) = C(G)$ whenever F is a mean-preserving spread of G . In most economic applications it would also be reasonable to expect C to be increasing in the sense of first-order stochastic dominance, that is, $C(F) \geq C(G)$ if $F \leq G$. We emphasize that our results do not depend on such monotonicity. In fact, we do not require C to be positive either.

Furthermore, we require C to be *lower semi-continuous* with respect to the weak*-topology on \mathcal{F} . That is, if $F \in \mathcal{F}$, $\{F_n\}_{n \in \mathbb{N}} \subset \mathcal{F}$, and the sequence $\{F_n\}_{n \in \mathbb{N}}$ converges to F in distribution then $\liminf_{n \rightarrow \infty} C(F_n) \geq C(F)$.⁸

⁷The assumption that the buyer trades when indifferent has no effect on our results but makes the analysis simpler.

⁸Recall that F_n converges to F in distribution if $\lim_{n \rightarrow \infty} F_n(v) = F(v)$ for all $v \in [0, 1]$ at which F is continuous.

Examples.— Our first example illustrates that requiring C to be increasing in risk is in the same spirit as assuming convexity of a traditional, deterministic investment cost function. Here, we assume that each particular valuation has a deterministic cost and the cost of a CDF F is the expected cost generated by F .

Example 1. Let $c : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous function and let

$$C(F) = \int_0^1 c(v) dF(v).$$

By Portmanteu's Theorem, the function C is continuous. In addition, By Proposition 6.D.2 of Mas-Colell et al (1995)), the cost function C is (strictly) increasing in risk if and only if the function c is (strictly) convex. The function c can be interpreted as a moment function and a cost of F as a general moment of F . Similarly, one could assume that the cost of a CDF is its variance. Since the variance is increasing in a mean-preserving spread, such a cost function is also increasing in risk.

In the following example, the cost of a CDF depends only on its expectation.

Example 2. Let $\tilde{c} : \mathbb{R} \rightarrow \mathbb{R}$ be a function and let

$$C(F) = \tilde{c} \left(\int_0^1 v dF(v) \right).$$

Note that any mean-based cost function can be described in this form. The reason is that any CDF is a mean-preserving spread of the degenerate distribution specifying an atom of size one at the mean of the CDF. Therefore, if C is mean-based then $C(F)$ depends only on the mean of the CDF F . If the function \tilde{c} is lower semi-continuous then this cost function is also lower semi-continuous.

One can combine the previous two examples and assume that the cost of a CDF is a function of many general moments of the distribution, that is,

$$C(F) = \tilde{c} \left(\int_0^1 c_1(v) dF(v), \dots, \int_0^1 c_n(v) dF(v) \right).$$

If c_i is convex for each $i = 1, \dots, n$ and \tilde{c} is increasing, this cost function is increasing in risk.

In the next example, the buyer first incurs a cost to determine her largest possible valuation. Then she can freely and stochastically destroy this value. So the cost of a distribution depends only on the largest value on its support.

Example 3. Let $\hat{c} : \mathbb{R} \rightarrow \mathbb{R}$ be a continuous function, let $\bar{s}(F)$ denote the upper bound of F 's support and let

$$C(F) = \hat{c}(\bar{s}(F)).$$

Since a mean-preserving spread enlarges the distribution's support and \bar{s} is lower semi-continuous, this cost is increasing in risk and lower semi-continuous as long as \hat{c} is increasing.

2.1 Alternative Interpretations of the Model

In this section, we first argue that our model can also be interpreted as a hold-up problem where the buyer has all the bargaining power and the seller can generate uncertainty about his production cost. Then we explain that by assuming a particular cost function, the buyer's choice of a CDF can be interpreted as a learning strategy about a given prior value distribution instead of investment in valuation.

Cost Uncertainty.— Consider the following alternative model. There is a seller who can produce and sell an object to a single buyer. The buyer's valuation for the object is normalized to be one. Prior to interacting with the buyer, the seller can choose any CDF $\bar{F} \in \mathcal{F}$ which then determines his production cost. The cost of a CDF is given by the function $\Gamma : \mathcal{F} \rightarrow \mathbb{R}$. The buyer observes the choice of the seller and gives a take-it-or-leave-it price offer to the buyer, τ . Finally, the seller's cost, γ , is realized and he decides whether to trade at price τ . If trade takes place, the payoff of the seller is $\tau - \gamma - \Gamma(\bar{F})$ and that of the buyer is $1 - \tau$. Otherwise, the seller's payoff is $-\Gamma(\bar{F})$ and the buyer's payoff is zero.

We explain that this game is strategically equivalent to our original model except that the roles of the buyer and the seller are switched. To this end, let p denote $1 - \tau$ and let v denote $1 - \gamma$. Furthermore, if γ is distributed according to \bar{F} then let F denote the distribution of $1 - \gamma (= v)$. Note that $\bar{F} \in \mathcal{F}$ if and only if $F \in \mathcal{F}$. Finally, define the function $C : \mathcal{F} \rightarrow \mathbb{R}$ such that $C(F) = \Gamma(\bar{F})$ for all $\bar{F} \in \mathcal{F}$. Further observe that, if trade takes place, the payoff of the seller can be rewritten as $v - p - C(F)$ and that of the buyer can be expressed as p . Conditional on trade, these are exactly the payoffs of the buyer and the seller, respectively, in our original model. Therefore, there is an equilibrium in this alternative model where the seller chooses $\bar{F} \in \mathcal{F}$ and the buyer sets price γ if and only if there is an equilibrium in our original model where the buyer chooses F and the seller sets price $p (= 1 - \gamma)$.

Learning About Valuation.— Let us recall the model of Roesler and Szentes (2017). There is a seller who has an object to sell to a single buyer. The buyer's valuation, v , for the object is given by the CDF $H \in \mathcal{F}$. The buyer can choose any signal s about v at no cost. After the seller observes the joint distribution of v and s the same bargaining game proceeds as in our model.

Roesler and Szentes (2017) argue that there exists a signal such that the CDF of the buyer's posterior value estimate is F if and only if H is a mean-preserving spread of F . In other words, assuming that the buyer can choose any signal structure is the same as assuming that, instead of learning about H , the buyer can choose any CDF F which then determines her valuation as long as H is a mean-preserving spread of F .

This model is equivalent to ours if the cost of a CDF F is defined to be zero if H is a mean-preserving spread of F and to be a constant, larger than one, say κ , otherwise. The reason

is that the buyer's payoff from trade cannot exceed one, so she would contemplate choosing the CDF F only if it is free, that is, if H is a mean-preserving spread of F . We argue that this cost function is increasing in risk, that is, $C(F) \geq C(G)$ whenever F is a mean-preserving spread of G . To see this, note that if H is a mean-preserving spread of F and hence, $C(F) = 0$, then H is also a mean-preserving spread of G , so $C(G) = 0$. If H is not a mean-preserving spread of F then $C(F) = \kappa$ and consequently, $C(F) \geq C(G)$ is satisfied because $C(G)$ is either zero or κ . Finally, we note that C is also lower semi-continuous because if H is a mean-preserving spread of F_n for all $n \in \mathbb{N}$ and $\{F_n\}_n$ converges to F in distribution then H is also a mean-preserving spread of F .

3 Main results

This section accomplishes three goals. We first analyze the special case where CDFs are costless, that is, $C \equiv 0$. This allows us to abstract from trade-offs arising from cost-benefit considerations and focus on the buyer's incentive to generate information rent. It also enables us to isolate the inefficiency due to this incentive and separate it from that emerging in standard hold-up problems stemming from the costly nature of the buyer's choice. We fully characterize the unique equilibrium in this case and demonstrate that the buyer's equilibrium demand is unit-elastic and trade always occurs. Second, we show that these two attributes are preserved in equilibria even when CDFs are costly as long as the cost function C is increasing in risk. Finally, we examine the case of cost function which are decreasing in risk. We show that the equilibrium CDF of the buyer typically specifies an atom at zero. However, conditional on the buyer's valuation being strictly positive, this CDF again induces a unit-elastic demand. Since the buyer rejects the seller's price only if her valuation is zero, trade is ex-post efficient also in this case.

Before stating our results, we introduce some notations and formally define the class of equal-revenue distributions which plays an important role in equilibrium constructions.

Notations.— For each $F \in \mathcal{F}$ and $p \in \mathbb{R}$, let $D(F, p)$ denote the *demand* at price p generated by F , that is, the probability of trade. Observe that $D(F, p) = 1 - F(p) + \Delta(F, p)$, where $\Delta(F, v)$ denotes the probability of v according to F .⁹ If the buyer chooses a CDF $F \in \mathcal{F}$, then the seller's profit is $\Pi(F) = \max_p pD(F, p)$.¹⁰ The seller might be indifferent between charging different prices which result in different payoffs to the buyer. For each $F \in \mathcal{F}$, let $P(F)$ denote the set of profit maximizing prices, that is, $P(F) = \arg \max_p pD(F, p)$. Furthermore, let $U(F, p)$ denote the buyer's *expected payoff from trade* if the seller sets price p , that is, $U(F, p) = \int_p^1 (v - p) dF(v)$. Finally, we refer to the pair (F, p) as an *equilibrium outcome* if there is an equilibrium in which the buyer chooses F and the seller responds by setting price p .

⁹That is, $\Delta(F, v) = F(v) - \sup_{z < v} F(z)$. If F does not specify an atom at v then $\Delta(F, v) = 0$.

¹⁰This maximum is well-defined because the buyer trades when she is indifferent.

Equal-revenue distributions.— For each $\pi \in (0, 1]$, let $F_\pi \in \mathcal{F}$ be defined as follows:

$$F_\pi^B(v) = \begin{cases} 0 & \text{if } v \in [0, \pi), \\ 1 - \frac{\pi}{v} & \text{if } v \in [\pi, B), \\ 1 & \text{if } v \in [B, 1]. \end{cases} \quad (1)$$

Since $F_\pi^B(\pi) = 0$, the valuation of the buyer is never below π . The function F_π^B is continuous and strictly increasing on $[\pi, B)$. On this interval, F_π^B is defined by the decreasing density $f_\pi(v) = \pi/v^2$. Finally, there is an atom of size π/B at $v = B$, that is, $\Delta(F_\pi^B, B) = \pi/B$. For completeness, let F_0^B be defined to be the degenerate distribution specifying an atom of size one at $v = 0$ for all $B \in [0, 1]$. Observe that this is a continuous extension in the sense that $F_{\pi_n}^B$ converges F_0^B in distribution if π_n goes to zero.

Note that the seller is indifferent between setting any price on the support of F_π^B , that is, $P(F_\pi^B) = [\pi, B]$. To see this, suppose first that the seller sets price $p \in [\pi, B)$. Then the seller's payoff is $pD(F_\pi^B, p) = p[1 - F_\pi^B(p)] = p(\pi/p) = \pi$. If the seller sets a price of B then his payoff is $\Delta(F_\pi^B, B) = B[\pi/B] = \pi$. Therefore, the seller's payoff is π as long as he sets a price in $[\pi, B]$, that is, $\Pi(F_\pi^B) = \pi$. As a consequence, these distributions exhibit unit-elastic demand.

3.1 Costless Distributions

Throughout this section we assume that $C \equiv 0$. So, when choosing a CDF, the only constraint faced by the buyer is that the upper bound of the valuation-support must be smaller than one. The main theorem of this section is stated as follows.

Theorem 1 *If $C \equiv 0$ then the unique equilibrium outcome is $(F_{1/e}^1, 1/e)$. In addition, $U(F_{1/e}^1, 1/e) = \Pi(F_{1/e}^1) = 1/e$.*

Recall that $P(F_{1/e}^1) = [1/e, 1]$, so the seller is indifferent between setting any price on the support of the buyer's equilibrium CDF. Therefore, this theorem implies that the buyer's equilibrium demand is unit-elastic. Furthermore, trade occurs with probability one since the seller's price, $1/e$, is the lowest valuation of the support of the buyer's equilibrium CDF, $F_{1/e}^1$. The seller's equilibrium strategy is not determined uniquely off the equilibrium path. He might charge different prices if he is indifferent after observing out-of-equilibrium CDFs.

The proof consists of two steps. First, for a given profit of the seller π , we compute the CDF $G_\pi \in \mathcal{F}$ which maximizes the buyer's payoff subject to the constraint that the seller's profit is π . This step essentially reduces our search for an equilibrium to a one-dimensional problem, since the buyer's equilibrium CDF must be in the set $\{G_\pi\}_\pi$. In the second step, we characterize the profit π and the corresponding CDF G_π which maximizes the buyer's payoff. To this end, we next show that $G_\pi = F_\pi^1$.

Lemma 1 Suppose that $G \in \mathcal{F}$, $p \in P(G)$ and let π denote $\Pi(G)$.

(i) Then F_π^1 first-order stochastically dominates G , that is, for all $v \in [0, 1]$

$$F_\pi^1(v) \leq G(v). \quad (2)$$

(ii) Furthermore, $U(F_\pi^1, \pi) \geq U(G, p)$ and the inequality is strict if $F_\pi^1 \neq G$.

Part (i) of the lemma says that the expected value induced by F_π^1 is larger than that induced by G . Recall that $\pi \in P(F_\pi^1)$, that is, the seller finds it optimal to set price π in the subgame generated by F_π^1 . In addition, $\Pi(G) = \pi = \Pi(F_\pi^1)$. Therefore, part (ii) implies that the maximum payoff the buyer can achieve subject to the constraint that the seller's profit is π is $U(F_\pi^1, \pi)$.

Proof. To prove part (i), note that for all $v \in [0, 1]$,

$$vD(G, v) \leq \pi,$$

or equivalently,

$$1 - \frac{\pi}{v} + \Delta(G, v) \leq G(v). \quad (3)$$

Since $\Delta(G, p) \in [0, 1]$, the left-hand side is weakly larger than $1 - (\pi/v)$. Therefore, the previous inequality implies (2).


To see part (ii), note that

$$U(F_\pi^1, \pi) = \int_\pi^1 (v - \pi) dF_\pi^1(v) \geq \int_\pi^1 (v - \pi) dG(v) \geq \int_p^1 (v - p) dG(v) = U(G, p),$$

where the first inequality follows from F_π^1 first-order stochastically dominating G and the second one from $\pi = pD(G, p) \leq p$. In addition, the first inequality is strict unless $F_\pi^1 = G$ (see Proposition 6.D.1 of Mas-Colell et al (1995)). ■

Figure 1 illustrates the statement and the proof of Lemma 1 for the case when G is the uniform distribution, that is, $G(v) = v$ on $[0, 1]$. Following Marshall's convention, the dependent variable, the price, is plotted on the vertical axis. In this case, $D(G, v) = 1 - v$, the seller's optimal price, p , is $1/2$ and the seller's profit, π , is $1/4$. The consumer surplus, $U(G, 1/2)$, is the integral of the demand curve on $[p, 1]$, represented by the darker shaded triangle. If the buyer's valuation is determined by the CDF F_π^1 , the seller's profit is π irrespective of which price he charges on the interval $[\pi, 1]$. This implies that $D(F_\pi^1, p) = D(G, p)$ and $D(F_\pi^1, v) > D(G, v)$ at $v \neq p$. In other words, the demand curve generated by F_π^1 is above that of G and the two curves are tangent to each other's at p . This explains part (i) of the lemma. To see part (ii), note that switching from G to F_π^1 increases the buyer's payoff for two reasons. First, since $D(F_\pi^1, v) > D(G, v)$, the integral of the new demand curve on $[p, 1]$ goes up. Second, if the seller charges π instead of p , the consumer surplus is the integral of the demand on a larger range than before. The increase in the buyer's payoff from moving (G, p) to (F_π^1, π) is the lighter shaded area on the figure.

Figure 1: Improving on the Uniform Distribution



We are ready to prove our main result.

Proof of Theorem 1. First, we show that there is an equilibrium in which the buyer chooses $F_{1/e}$ and the seller responds by setting price $1/e$. Recall that $P(F_\pi^1) = [\pi, 1]$ for all π , so the seller optimally charges the price $1/e$ in the subgame generated by $F_{1/e}^1$. Next, we argue that the buyer has no incentive to deviate either. If the buyer does not deviate, her payoff is $U(F_{1/e}^1, 1/e)$. By part (ii) of Lemma 1, any deviation payoff of the buyer is weakly smaller than $U(F_\pi^1, \pi)$ for some $\pi \in [0, 1]$. So, it is sufficient to show that $U(F_\pi^1, \pi)$ is maximized at $\pi = 1/e$. Note that

$$U(F_\pi^1, \pi) = \int_\pi^1 (v - \pi) dF_\pi^1(v) = \int_\pi^1 v f_\pi(v) dv + \Delta(F_\pi^1, 1) - \pi = \int_\pi^1 \frac{\pi}{v} dv = -\pi \log(\pi), \quad (4)$$

where the second equality follows from $f_\pi(v) = \pi/v^2$ and $\Delta(F_\pi^1, 1) = \pi$. The function $-\pi \log(\pi)$ is indeed maximized at $1/e$. From the equality chain (4), the buyer's payoff is $U(F_{1/e}^1, 1/e) = -1/e \log(1/e) = 1/e$ in this equilibrium. Finally, since the price is $1/e$ and the buyer always trades, $\Pi(F_{1/e}^1) = 1/e$ follows.

It remains to prove the uniqueness of the equilibrium outcome. Part (ii) of Lemma 1 implies that the equilibrium payoff of the buyer, $U(F^*, p^*)$, is weakly smaller than $U(F_{\Pi(F^*)}^1, \Pi(F^*))$. We showed that $U(F_{\Pi(F^*)}^1, \Pi(F^*)) \leq U(F_{1/e}^1, 1/e)$ in the previous paragraph, so the buyer's payoff

cannot exceed $1/e$ in any equilibrium. Since $1/e$ is the unique maximizer of $U(F_\pi^1, \pi)$, part (ii) of Lemma 1 also implies that in order for the buyer to achieve this payoff, she must choose $F_{1/e}^1$ and the seller must charge $1/e$. Therefore, to establish uniqueness, it is sufficient to show that the buyer's equilibrium payoff is at least $1/e$. To this end, for each $\varepsilon \in (0, 1/e)$, let G^ε be $\varepsilon + F_{1/e}^1$ on $[p, 1)$ and $F_{1/e}^1$ otherwise. The CDF G^ε is constructed from $F_{1/e}^1$ by moving a mass of size ε from the atom at $v = 1$ to $v = 1/e$. If the buyer chooses G^ε then the seller strictly prefers to set price $1/e$.¹¹ In addition, the buyer's payoff is ε -close to $U(F_{1/e}^1, 1/e) = 1/e$. Therefore, the buyer can achieve a payoff arbitrarily close to $1/e$ irrespective of the seller's strategy to resolve ties and hence the buyer's equilibrium payoff cannot be smaller than $1/e$. ■

From the proof of Theorem 1 it follows that the buyer's expected valuation generated by $F_{1/e}^1$ is $2/e$. Indeed, from (4),

$$\int_{1/e}^1 v dF_{1/e}(v) = \int_{1/e}^1 v f_{1/e}(v) dv + \Delta(F_{1/e}^1, 1) = \int_{1/e}^1 \frac{1}{ev} dv + \frac{1}{e} = \frac{2}{e}.$$

Note that efficiency requires the buyer to choose a distribution which would specify $v = 1$ with probability one, and the seller to quote a price which is weakly less than one. The total surplus would be one. In contrast, the total surplus in equilibrium is only $U(F_{1/e}^1, 1/e) + \Pi(F_{1/e}^1) = 2/e$. Hence, the efficiency loss due to the buyer's desire to generate information rent is $1 - 2/e \approx 0.26$.

Of course, requiring the upper bound of the buyer's CDF's support to equal one is just a normalization. If this upper bound is k then the equilibrium price, the payoff of the buyer and the payoff of the seller are k/e . As k goes to infinity, the payoffs also converge to infinity.

3.2 Increasing Cost

In this section, we consider the case where the buyer's cost is increasing in risk and show that the search for an equilibrium CDF can again be restricted to the set of equal-revenue distributions. To this end, let S^* denote the set of parameters corresponding to the buyer's optimal CDFs among the equal-revenue distributions, that is,

$$S^* = \arg \max_{\substack{(\pi, B) \\ \pi \in [0,1], B \in [\pi,1]}} \{U(F_\pi^B, \pi) - C(F_\pi^B) : F_\pi^B \in \mathcal{F}\}. \quad (5)$$

Since the domain of this maximization problem is compact, U is continuous and C is lower semi-continuous in (π, B) , the maximum is attained, so $S^* \neq \{\emptyset\}$. The main result of this section is the following

¹¹By setting price $1/e$, the seller achieves the same payoff as if the buyer chose $F_{1/e}^1$. For any price in $(1/e, 1]$, the seller's payoff is strictly smaller than after the choice of $F_{1/e}^1$ because the probability of trade is smaller by ε .

Theorem 2 *Suppose that the function C is increasing in risk. Then*

- (i) $(F_{\pi^*}^{B^*}, \pi^*)$ is an equilibrium outcome if $(\pi^*, B^*) \in S^*$ and
- (ii) there is no other equilibrium outcome if C is strictly increasing in risk or mean-based.

Part (i) implies that there always exist equilibria in which the buyer's demand is unit-elastic and trade occurs with probability one as long as the cost is increasing in risk. Part (ii) says that if the cost is strictly increasing in risk or depends only on the distribution's mean then all equilibria have these features. We observe that this theorem does not imply, but leaves open the possibility of, the existence of equilibria in which the buyer's CDF is not an equal-revenue distribution if the cost only weakly increases in risk but is not mean-based. We construct such equilibria in Section 4 in the context of Example 1.

We need two lemmas to prove this theorem. The first one states that for each CDF G there is an equal-revenue distribution which second-order stochastically dominates G and increases the buyer's expected payoff from trade relative to G while generating the same profit to the seller as G .

Lemma 2 *Suppose that $G \in \mathcal{F}$, $p \in P(G)$ and let π denote $\Pi(G)$. Then there exists a unique $B \in [\pi, 1]$ such that*

- (i) G is a mean-preserving spread of F_{π}^B ,
- (ii) $U(G, p) \leq U(F_{\pi}^B, \pi)$ and
- (iii) $B < 1$ unless $G = F_{\pi}^1$.

Proof. See Roesler and Szentes (2017). ■

A key implication of this lemma is that if the function C is increasing in risk, the buyer can weakly improve her payoff from any CDF G by choosing an equal-revenue distribution. Indeed, by part (ii), the buyer's payoff from trade is higher at F_{π}^B than at G . In addition, the CDF F_{π}^B is cheaper than G because G is a mean-preserving spread of F_{π}^B by part (i) and C is increasing in risk.

To outline the argument behind part (i), recall that Lemma 1 states that F_{π}^1 first-order stochastically dominates G and hence, the expectation of F_{π}^1 exceeds that of G . To decrease its mean, F_{π}^1 can be truncated by moving all the probability mass from above B to build an atom at B . The threshold of the truncation, B , can be chosen so that the expectation of the new CDF, F_{π}^B , is exactly the same as that of G . Furthermore, since F_{π}^1 first-order stochastically dominates G and $F_{\pi}^B = F_{\pi}^1$ on $[0, B]$, it follows that $G \geq F_{\pi}^B$ on $[0, B]$. In addition, $G \leq F_{\pi}^B$ on $[B, 1]$ because $F_{\pi}^B = 1$ on $[B, 1]$. It is not hard to see that these two observations imply that G is indeed a mean-preserving spread of F_{π}^B .

To explain part (ii), note that the sum of the payoffs from trade of the buyer and the seller cannot be larger than the expectation of the buyer's CDF. That is, the buyer's payoff from trade,

$U(G, p)$, cannot exceed the difference between the mean of G and the seller's profit, π . But this difference is exactly the buyer's payoff from trade if she chooses F_π^B and the seller sets price π . The reason is that part (i) implies that G and F_π^B have the same expectations and trade always occurs because π is the lowest value in the support of F_π^B .

Lemma 3 *Let $\pi \in (0, 1)$ and $B \in [\pi, 1)$. Then for each $\bar{B} \in (B, 1]$, there exists a unique $\bar{\pi} \in (0, \pi)$ such that $F_{\bar{\pi}}^{\bar{B}}$ is a mean-preserving spread of F_π^B . In addition, $U(F_\pi^B, \pi) < U(F_{\bar{\pi}}^{\bar{B}}, \bar{\pi})$.*

Proof. See the Appendix. ■

This lemma states that if the upper bound of an equal-revenue distribution is less than one then the buyer's expected payoff from trade can be increased by expanding the support of this CDF in both directions without affecting its mean. A consequence of this lemma is that if the cost function is mean-based, the upper bound of the support of any optimal equal-revenue distribution is one. This observation plays an important role in the proof part (ii) of Theorem 2 for mean-based cost functions.

Finally, we are ready to prove the main result of this section.

Proof of Theorem 2. To prove part (i), fix $(\pi^*, B^*) \in S^*$ and note that if the buyer chooses $F_{\pi^*}^{B^*}$ then the seller optimally responds by setting price π^* since $P(F_{\pi^*}^{B^*}) = [\pi^*, B^*]$. We have to show that the buyer has no incentive to deviate. To this end, let $G \in \mathcal{F}$, $p \in P(G)$ and let π denote $\Pi(G)$. By part (i) of Lemma 2, there exists a B such that G is a mean-preserving spread of F_π^B . Then

$$U(G, p) - C(G) \leq U(F_\pi^B, \pi) - C(G) \leq U(F_\pi^B, \pi) - C(F_\pi^B) \leq U(F_{\pi^*}^{B^*}, \pi^*) - C(F_{\pi^*}^{B^*}), \quad (6)$$

where the first inequality follows from part (ii) of Lemma 2. The second inequality follows from the fact that $C(G) \geq C(F_\pi^B)$ since G is a mean-preserving spread of F_π^B and C is increasing in risk. The third inequality follows from $(\pi^*, B^*) \in S^*$. The previous inequality chain implies that choosing G instead of $F_{\pi^*}^{B^*}$ is not a profitable deviation.

To show part (ii), we only argue that for all $G \in \mathcal{F}$, $p \in P(G)$, $U(G, p) - C(G) < U(F_{\pi^*}^{B^*}, \pi^*) - C(F_{\pi^*}^{B^*})$ for all $(\pi^*, B^*) \in S^*$ unless $G = F_{\hat{\pi}}^{\hat{B}}$ for some $(\hat{\pi}, \hat{B}) \in S^*$. In other words, we show that one of the inequalities is strict in (6). The rest of the proof is essentially identical to the uniqueness proof of Theorem 1. Note that if $G = F_\pi^B$ then the last inequality in (6) is strict unless $(\pi, B) \in S^*$. We now consider the case when $G \neq F_\pi^B$. Suppose first that C is strictly increasing in risk. Since G is a mean preserving spread of F_π^B , $C(G) > C(F_\pi^B)$ and the second inequality in (6) is strict. Suppose now that C is mean-based. Since $G \neq F_\pi^B$, part (iii) of Lemma 2 implies that $B < 1$. Then, by Lemma 3, there exists $\bar{\pi} \in (0, \pi)$ and $\bar{B} \in (B, 1]$ such that $F_{\bar{\pi}}^{\bar{B}}$ is a mean-preserving spread of F_π^B . Hence,

$$U(F_\pi^B, \pi) - C(F_\pi^B) = U(F_\pi^B, \pi) - C(F_{\bar{\pi}}^{\bar{B}}) < U(F_{\bar{\pi}}^{\bar{B}}, \bar{\pi}) - C(F_{\bar{\pi}}^{\bar{B}}),$$

where the equality follows from C being mean-based and $F_{\bar{\pi}}^B$ being a mean preserving spread of F_{π}^B . The inequality follows from Lemma 3. This inequality chain implies that $(\pi, B) \notin S^*$ and hence, the last inequality in (6) is strict. ■

Note that even if the hypothesis of part (ii) of Theorem 2 is satisfied and each equilibrium CDF is an equal-revenue distribution, the parameters defining these distribution are not characterized. However, in the proof of this theorem, we have shown that if the cost function is mean-based then Lemma 3 implies that the upper bound of the support is one. Formally, we state the following

Remark 1 *If C is mean-based and $(F_{\pi^*}^B, \pi^*)$ is an equilibrium outcome then $B^* = 1$.*

We conclude this section with a discussion of the relationship between Theorem 2 and the main result of Roesler and Szentes (2017). As mentioned in Section 2.1, the model of Roesler and Szentes (2017) is equivalent to ours if the cost of a CDF is defined to be zero if the buyer's prior value distribution is a mean-preserving spread of it and its cost is $\kappa (> 1)$ otherwise. We also explained that this cost function is increasing in risk. Therefore, Theorem 2 is applicable and part (i) of this theorem implies the existence of a buyer-optimal signal structure which generates a unit-elastic demand and induces trade with probability one, which is the main result of Roesler and Szentes (2017).

In fact, Theorem 2 can be used to generalize the main result of Roesler and Szentes (2017) to environments where learning is costly for the buyer as long as more informative signals cost more. The key observation to this is that the distribution of the buyer's posterior value estimate generated by a more informative signal is a mean-preserving spread of that generated by the less informative signal. Therefore, if more informative signals are more expensive, the buyer's cost of inducing a distribution of her posterior value estimate is increasing in risk. Provided that this cost is also lower semi-continuous, part (i) of Theorem 2 implies that the result of Roesler and Szentes (2017) remains valid. In addition, if more informative signals are strictly more expensive then, by part (ii) of Theorem 2, the buyer-optimal signal structure is unique.

3.3 Decreasing Cost


In this section, we assume that the buyer's cost function, C , is decreasing in risk. Next, we introduce a class of distributions and then show that there always exists an equilibrium where the buyer chooses a CDF from this class. For each $\pi \in (0, 1]$ and $l \in [\pi, 1]$ define

$$F_{\pi, l}(v) = \begin{cases} 1 - \frac{\pi}{l} & \text{if } v < l, \\ 1 - \frac{\pi}{v} & \text{if } v \in [l, 1), \\ 1 & \text{if } v \geq 1. \end{cases}$$

The CDF $F_{\pi,l}$ is a convex combination of an atom at zero and the equal-revenue distribution F_l^1 . Indeed, if the buyer's valuation is zero with probability $1 - (\pi/l)$ and it is determined by F_l^1 with probability π/l then the resulting CDF is $F_{\pi,l}$. Of course, when the seller decides what price to set, he conditions on the event that the buyer's value is determined by F_l^1 . Hence, $P(F_{\pi,l}) = [l, 1]$ and $\Pi(F_{\pi,l}) = (\pi/l)l = \pi$. Observe that $F_{\pi,\pi} = F_\pi^1$ and the CDF $F_{\pi,1}$ specifies an atom of size π at one and places the rest of the probability mass at zero. For completeness, let $F_{0,l}$ be the degenerate distribution specifying an atom of size one at $v = 0$ for all $l \in [0, 1]$. This is a continuous extension in the sense that $F_{\pi_n,l}$ converges $F_{0,l}$ in distribution if π_n goes to zero.

Figure 2 illustrates the comparison of $F_{\pi,l}$ and F_π^B for the case of $l < B < 1$. Note that F_π^1 coincides with $F_{\pi,l}$ on $[l, 1]$ and with F_π^B on $[\pi, B]$. Since $l < B$, the lower envelope of the CDFs $F_{\pi,l}$ and F_π^B on this figure is F_π^1 .

Figure 2: An illustration of F_π^B and $F_{\pi,l}$



Let S^{**} denote the set of parameters corresponding to the buyer's optimal CDFs in the set $\{F_{\pi,l}\}$, that is,

$$S^{**} = \arg \max_{\substack{(\pi,l) \\ \pi \in [0,1], l \in [\pi,1]}} \{U(F_{\pi,l}, l) - C(F_{\pi,l}) : F_{\pi,l} \in \mathcal{F}\}. \quad (7)$$

Again, since the domain of this maximization problem is compact, U is continuous and C is lower semi-continuous in (π, B) , the maximum is attained, so $S^{**} \neq \{\emptyset\}$. The main result of this section is the following

Theorem 3 *Suppose that the function C is decreasing in risk. Then*

- (i) (F_{π^*, l^*}, l^*) is an equilibrium if $(\pi^*, l^*) \in S^{**}$ and
- (ii) there is no other equilibrium outcome if C is strictly decreasing in risk.

As mentioned above, $P(F_{\pi^*, l^*}) = [l^*, 1]$ and $F_{\pi^*, l^*}(l^*) = F_{\pi^*, l^*}(0)$, that is, the buyer's valuation is never in the interval $(0, l^*]$. Therefore, since $F_{\pi^*, l^*} = F_{\pi^*}^1$ on $[l^*, 1]$, part (i) implies that there exists an equilibrium where, conditional on the valuation being strictly positive, the buyer's demand is unit-elastic. Furthermore, since the seller's equilibrium price is l^* , trade occurs unless the buyer's valuation is zero. Therefore, trade is ex-post efficient. Part (ii) says that if the cost is strictly decreasing in risk then all equilibria have these features.

The next lemma shows that for each CDF G there is an element in $\{F_{\pi,l}\}$ which is a mean-preserving spread of G , induces the same profit for the seller and a weakly higher payoff from trade to the buyer.

Lemma 4 *Suppose that $G \in \mathcal{F}$, $p \in P(G)$ and let π denote $\Pi(G)$. Then there exists a unique $l \in [\pi, 1]$ such that*

- (i) $F_{\pi,l}$ is a mean-preserving spread of G ,
- (ii) $U(G, p) \leq U(F_{\pi,l}, l)$ and the inequality is strict if $D(G, p) < 1 - G(0)$.

Proof. See the Appendix. ■

A key implication of this lemma is that if the function C is decreasing in risk, the buyer can weakly improve her payoff from any CDF G by choosing $F_{\pi,l}$ instead. Indeed, by part (ii), the buyer's payoff from trade is higher at $F_{\pi,l}$ than at G . In addition, the CDF $F_{\pi,l}$ is cheaper than G because $F_{\pi,l}$ is a mean-preserving spread of G by part (i) and C is decreasing in risk.

We observe that Lemma 4 is the counterpart of Lemma 2. Recall that for each $G \in \mathcal{F}$, Lemma 2 identified a CDF, F_{π}^B , such that G is a mean-preserving spread of F_{π}^B and the buyer's payoff from trade generated by F_{π}^B exceeds that generated by G . In contrast, Lemma 4 identifies

a CDF $F_{\pi,l}$ which is a mean-preserving spread of G and improves the buyer's payoff from trade relative to G .

We are ready to prove the main theorem of this section.

Proof of Theorem 3. To prove part (i), fix $(\pi^*, l^*) \in S^{**}$ and note that if the buyer chooses F_{π^*, l^*} then the seller optimally responds by setting price l^* since $P(F_{\pi^*, l^*}) = [l^*, 1]$. We have to show that the buyer has no incentive to deviate. To this end, let $G \in \mathcal{F}$, $p \in P(G)$ and let π denote $\Pi(G)$. By part (i) of Lemma 4, there exists an l such that $F_{\pi,l}$ is a mean-preserving spread of G . Then

$$U(G, p) - C(G) \leq U(F_{\pi,l}, l) - C(G) \leq U(F_{\pi,l}, l) - C(F_{\pi,l}) \leq U(F_{\pi^*, l^*}, l^*) - C(F_{\pi^*, l^*}), \quad (8)$$

where the first inequality follows from part (ii) of Lemma 4. The second inequality follows from the fact that $C(G) \geq C(F_{\pi,l})$ since $F_{\pi,l}$ is a mean-preserving spread of G and C is decreasing in risk. The third inequality follows from $(\pi^*, l^*) \in S^{**}$. The previous inequality chain implies that choosing G instead of F_{π^*, l^*} is not a profitable deviation.

To show part (ii), we only argue that for all $G \in \mathcal{F}$, $p \in P(G)$ and $\pi = \Pi(G)$, $U(G, p) - C(G) < U(F_{\pi^*, l^*}, l^*) - C(F_{\pi^*}^{B*})$ for all $(\pi^*, l^*) \in S^*$ unless $G = F_{\hat{\pi}, \hat{l}}$ for some $(\hat{\pi}, \hat{l}) \in S^{**}$. In other words, we show that one of the inequalities is strict in (8). The rest of the proof is essentially identical to the uniqueness proof of Theorem 1. Note that if $G = F_{\pi,l}$ then the last inequality is strict in (6) unless $(\pi, l) \in S^{**}$. We now consider the case when $G \neq F_{\pi,l}$. Since F_{π}^B is a mean preserving spread of G and C is strictly decreasing in risk, $C(G) > C(F_{\pi}^B)$ and the second inequality in (6) is strict. ■

We conclude this section with an application of Theorem 3 to a problem of optimally designing riskiness. To this end, consider a scenario where the buyer's prior value distribution is given, say by $H \in \mathcal{F}$, and she can reshape it by adding risk. That is, before observing the realization of H , the buyer can choose any CDF $F \in \mathcal{F}$ which is a mean-preserving spread of H at no cost.¹² This new CDF will then determine the buyer's valuation for the object. After the seller observes the buyer's choice, the same bargaining game ensues as in our model. In this case, the buyer's problem is to maximize $U(F, \min P(F))$, over the set of CDFs which are mean-preserving spreads of H . Note that this is also the problem of finding an equilibrium CDF in our model if the cost of a CDF F is defined to be zero if it is a mean-preserving spread of H and its cost is one otherwise. The reason is that the buyer's payoff from trade cannot exceed one, so she would only contemplate choosing the CDF F if it is free, that is, F is a mean-preserving spread of H . We argue that this cost function is decreasing in risk, that is, $C(F) \geq C(G)$ whenever G is a mean-preserving spread of F . To see this, note that if F is a mean-preserving spread of H and hence, $C(F) = 0$, then G is also a mean-preserving spread of H , so $C(G) = 0$. If F is not a mean-preserving spread of

¹²We refer to Johnson and Myatt (2006) for an extensive discussion of environments, such as product design and advertisement, where the change in demand are modeled as mean-preserving spreads.

H then $C(F) = 1$ and consequently, $C(F) \geq C(G)$ is satisfied because $C(G)$ is either zero or one. Since C is also lower semi-continuous¹³, part (i) of Theorem 3 implies the existence of an equilibrium CDF which generates unit-elastic demand conditional on the buyer's valuation being strictly positive and ex-post efficient trade.

4 Examples Revisited

Theorems 2 and 3 provide only a partial characterization of equilibria in the case of costly CDFs. In particular, they allow for multiplicity and, unless C is strictly monotonic in risk or mean-based, there might be equilibria in which the buyer's CDF is not an equal-revenue distribution. To further discuss these issues, we revisit the examples of Section 2 and accomplish two goals. First, we provide specifications of the cost function under which the equilibrium outcome is unique and we characterize the buyer's equilibrium equal-revenue distribution. We also discuss alternative specifications under which the equilibrium is not unique and construct equilibrium CDFs which are not equal-revenue distributions.

Example 1. (Continued.) Recall that $C(F) = \int_0^1 c(v) dF(v)$ and let us assume that $c : \mathbb{R} \rightarrow \mathbb{R}$ is an increasing, differentiable and convex function such that $c(0) = 0$ and $c'(1) \geq 1$.

Recall that, since c is convex, the function C is increasing in risk. The next claim characterizes the unique equilibrium outcome if c is strictly convex and describes equilibrium CDFs which are not equal-revenue distributions if c becomes linear above a certain threshold.

Claim 1 (i) *If c is strictly convex then the unique equilibrium CDF is $F_{\pi^*}^{B^*}$ where $B^* = c'^{-1}(1)$ and $\pi^* \leq B^*/e$.*

(ii) *If c is strictly convex on $[0, \underline{B}]$ where $\underline{B} \in (0, 1)$ and $c'(v) = 1$ on $[\underline{B}, 1]$ then there exists a π^* such that any CDF G is an equilibrium CDF if $G(v) = F_{\pi^*}^1(v)$ on $[\pi^*, \underline{B}]$ and $G(v) \geq F_{\pi^*}^1(v)$ on $[\underline{B}, 1]$.*

Proof. See the Appendix. ■

Part (i) of this claim illustrates that the equilibrium characterization of Theorem 2 can be further sharpened for certain cost functions. In particular, the upper bound of the support of the equilibrium CDF is pinned down by the function c in this example. Part (ii) demonstrates that equilibrium multiplicity can indeed be a consequence of certain cost functions. Note that, in this example, adding a mean-preserving spread to a distribution is costless if and only if it affects only valuations on $[\underline{B}, 1]$. Therefore, the cost function C generated by c in part (ii) is neither strictly increasing in risk nor it is mean-based. Hence, the hypothesis of part (ii) of Theorem 2 is violated.

¹³The reason is that if F_n is a mean-preserving spread of H for all $n \in \mathbb{N}$ and $\{F_n\}_n$ converges to F in distribution then F is also a mean-preserving spread of H .

Next, we show that the buyer's equilibrium CDF is always an equal-revenue distribution in the context of Examples 2 and 3.

Example 2. (Continued.) Recall that for each $F \in \mathcal{F}$, $C(F) = \tilde{c}(\mu_F)$ where μ_F is the expectation of F . Let us assume that the function \tilde{c} is continuous and increasing.

In what follows, we reduce the problem of finding an equilibrium CDF to a one-dimensional maximization problem.

Claim 2 *The CDF F^* is an equilibrium CDF if and only if $F^* = F_{\pi^*}^1$ where*

$$\pi^* \in \arg \max_{\pi \in [0, 1/e]} [-\pi \log(\pi) - \tilde{c}(\pi - \pi \log(\pi))].$$

Proof. See the Appendix. ■

Example 3 (Continued.) Recall that if the buyer chooses a distribution and the upper bound of its support is k then her cost is $\hat{c}(k)$ and let us assume that \hat{c} is continuous.

We use the results of Section 3.1 to characterize equilibria for this example.

Claim 3 *The CDF F^* is an equilibrium CDF if and only if $F^* = F_{k^*/e}^{k^*}$ where solves*

$$k^* \in \arg \max_{k \in [0, 1]} \left[\frac{k}{e} - \hat{c}(k) \right].$$

Proof. See the Appendix. ■

5 Discussion

To conclude, we discuss various assumptions of our model and describe equilibria under alternative assumptions.

Risk Preferences.— We first examine the extent to which the buyer's risk-neutrality affects our results. To this end, suppose that the buyer's ex-post payoff from trade is $u(v - p)$ if her valuation is v and trades at price p and $u(0)$ otherwise and that u is a strictly increasing function. We first argue that in the benchmark model, where distributions are costless, the buyer's equilibrium CDF still generates a unit-elastic demand. The statement of Lemma 1 is still valid, that is, the buyer-optimal CDF which provides the seller with a given profit π is still F_{π}^1 . The reason is that a monotonic transformation of the buyer's payoff has no effect on the buyer's demand and hence, F_{π}^1 first-order stochastically dominates any other CDF generating profit π . Since the buyer's payoff is increasing in $v - p$, this implies that the buyer's expected payoff is maximized at F_{π}^1 . Therefore, the search for an equilibrium CDF can again be restricted to the class of equal-revenue distribution,

$\{F_\pi^1\}_{\pi \in (0,1]}$. Of course, the equilibrium value of π depends on the Bernoulli utility, u , and it is not necessarily $1/e$.

Suppose now that CDF's are costly and the function C is increasing in risk.¹⁴ Then, in order for a version of Theorem 2 to be valid, the buyer needs to be risk-averse. To see this, we observe that the statement of Lemma 2 holds only if the buyer is risk-averse. The reason is that since G is a mean-preserving spread of F_π^B , the Bernoulli utility, u , must be concave to guarantee that the buyer prefers F_π^B over G . A similar argument yields that, if C is decreasing in risk, the utility function needs to be convex for a version of Theorem 3 to hold.

Production Cost.— We have implicitly assumed that the seller's production cost is zero. It is not hard to generalize our results to the case where the seller has to pay a cost $c \in (0, 1)$ if trade occurs. For example, if CDFs are costless, one can follow the same two-step procedure to solve the problem as in Section 3.1. Given that the seller's profit must be π , the distribution which maximizes the buyer's payoff is defined by the continuous CDF $1 - \pi/(v - c)$ on $[\pi + c, 1)$ and an atom of size $\pi/(1 - c)$ at $v = 1$. This distribution makes the seller indifferent between setting any price on $[\pi + c, 1]$. The profit which maximizes the buyer's payoff is $(1 - c)/e + 2c$.

Bargaining Power.— As it will be argued, the assumption that the seller has all the bargaining power can be relaxed. To this end, let us assume that, after the buyer has chosen the CDF, a bargaining game with random proposer ensues. With probability $\alpha \in (0, 1)$, the buyer that makes a take-it-or-leave-it offer to the seller; while with probability $1 - \alpha$ the seller makes the offer. Of course, if the buyer makes the offer the price is zero. Therefore, in choosing the CDF, the buyer maximizes $\alpha\mu_F + (1 - \alpha)U(F, p) - C(F)$ where μ_F denotes the expectation of F , $p = \min P(F)$ and $C(F)$ is the cost of F . This problem is equivalent to maximizing $U(F, p) - (C(\mu_F) - \alpha\mu_F)/(1 - \alpha)$, which is precisely the buyer's problem if the seller has all the bargaining power and the cost of F is $(C(F) - \alpha\mu_F)/(1 - \alpha)$. Observe that this cost function is increasing (decreasing) in risk if and only if C is increasing (decreasing) in risk. Therefore, Theorems 2 and 3 are applicable even when the seller does not have all the bargaining power.

Restricted Choice Set.— In many applications of our model, the buyer may not have access to every distribution. The information design approach adopted in this paper then provides an upper bound on the buyer's payoff. Furthermore, the buyer's equilibrium CDF still approximates an equal-revenue distribution in a sense explained below. Consider a model where CDFs are costless and the buyer's action space, say \mathcal{A} , is restricted, that is, $\mathcal{A} \subsetneq \mathcal{F}$. Further assume that (F, p) is an equilibrium outcome. Then the CDF F is closest to F_p^1 on average on the interval $[p, 1]$ among all CDFs for which the seller finds it optimal to set price p , that is,

¹⁴If the buyer chooses F , her ex-post payoff from trade can either be defined as $u(v - p) - C(F)$ or $u(v - p - C(F))$.

$$F \in \arg \min_{\substack{G \in \mathcal{A} \\ p \in P(G)}} \int_p^1 [G(v) - F_p^1(v)] dv. \quad (9)$$

To see this, note that

$$\begin{aligned} U(F_p^1, p) - U(G, p) &= \int_p^1 (v - p) dF_p^1(v) - \int_p^1 (v - p) dG(v) \\ &= \int_p^1 [1 - F_p^1(v)] dv - \int_p^1 [1 - G(v)] dv = \int_p^1 [G(v) - F_p^1(v)] dv, \end{aligned}$$

where the second equality follows from integration by parts. Since the equilibrium CDF of the buyer solves

$$\min_{\substack{G \in \mathcal{A} \\ p \in P(G)}} [U(F_p^1, p) - U(G, q)],$$

the statement in (9) follows.¹⁵ Of course, if CDFs are not free the approximation is going to take the cost into account.

References

- [1] Bergemann, Dirk, Benjamin Brooks, and Stephen Morris, “The Limits of Price Discrimination,” *American Economic Review*, March 2015, 105 (3), 921–57.
- [2] Bergemann, Dirk and Martin Pesendorfer, “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 137 (1), 580-609.
- [3] Bergemann, Dirk and Karl H. Schlag, “Pricing without Priors,” *Journal of the European Economic Association*, 04-05 2008, 6 (2-3), 560–569.
- [4] Brooks, Benjamin, “Surveying and selling: Belief and surplus extraction in auctions,” *mimeo*, University of Chicago, 2013.
- [5] Dilmé, Francisc, “Pre-Trade Risky Investments,” *mimeo*, University of Bonn, 2017.
- [6] Gul, Faruk, “Unobservable Investment and the Hold-Up Problem,” *Econometrica*, March 2001, 69 (2), 343–76.

¹⁵It can also be shown that F approximates F_p^1 in the entire set \mathcal{A} in the sense that

$$F \in \arg \min_{\substack{G \in \mathcal{A} \\ q \in P(G)}} \int_p^1 [G(v) - F_p^1(v)] dv - \int_q^p [1 - G(v)] dv.$$

- [7] Hart, Sergiu and Noam Nisan, “Approximate Revenue Maximization with Multiple Items,” in “Proceedings of the 13th ACM Conference on Electronic Commerce” EC ’12 ACM New York, NY, USA 2012, pp. 656–656.
- [8] Hermalin, Benjamin E. and Michael L. Katz, “Information and the hold-up problem,” *RAND Journal of Economics*, 2009, *40* (3), 405–423.
- [9] Johnson, Justin and David Myatt, “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 2006, *96* (3), 756–784.
- [10] Kamenica, Emir and Gentzkow, Matthew, “Bayesian Persuasion ” *American Economic Review*, 2011, *101* (6), 2590–2615
- [11] Kremer, Michael and Christopher Snyder, “Preventives versus Treatments, *Quarterly Journal of Economics*, 2015, 1167–1239.
- [12] Lau, Stephanie, “Information and bargaining in the hold-up problem,” *RAND Journal of Economics*, 2008, *39* (1), 266–282.
- [13] Mas-Colell A., M. Whinston and J. Green, *Microeconomic Theory*, Oxford University Press, 1995.
- [14] Neeman, Zvika, “The effectiveness of English auctions,” *Games and Economic Behavior*, May 2003, *43* (2), 214–238.
- [15] Ortner, Juan and Sylvain Chassang, “Making collusion hard: asymmetric information as a counter-corruption measure”, *Mimeo*, Princeton University.
- [16] Renou, Ludovic and Karl H. Schlag, “Minimax regret and strategic uncertainty,” *Journal of Economic Theory*, January 2010, *145* (1), 264–286.
- [17] Roesler Ann-Katrin and Balazs Szentes, “Buyer-optimal Learning and Monopoly Pricing,” *American Economic Review*, 2017, *107*(7): 2072–80.
- [18] Shapiro, Carl “Injunctions, Hold-Up, and Patent Royalties,” *American Law and Economics Review*, 2010, *12*(2): 280–319.
- [19] Siebert, Ralph “What Determines Firms’ Choices Between Ex Ante and Ex Post Licensing Agreements?” *Journal of Competition Law & Economics*, *11*(1), 165–199
- [20] Skrzypacz, Andrzej, “Bargaining under Asymmetric Information and the Hold-up Problem”, *mimeo*, Stanford University, 2004.
- [21] Tirole, Jean, “Procurement and Renegotiation,” *Journal of Political Economy*, 1986, *94* (2), 235–59.

6 Appendix

Proof of Lemma 3. Fix a $\bar{B} \in (B, 1]$ and let μ denote the expectation of F_{π}^B . We first show that there exists a unique $\bar{\pi} \in (0, \pi)$ such that the expected value generated by $F_{\bar{\pi}}^{\bar{B}}$ is μ . Note that

$$\int_0^1 v dF_0^{\bar{B}}(v) = 0 < \int_0^1 v dF_{\pi}^B(v) = \mu < \int_0^1 v dF_{\bar{\pi}}^{\bar{B}}(v),$$

where the second inequality follows from the fact that $F_{\bar{\pi}}^{\bar{B}}$ first-order stochastically dominates F_{π}^B , that is, $F_{\pi}^B \geq F_{\bar{\pi}}^{\bar{B}}$ because $\bar{B} > B$. This inequality chain implies that the expectation of $F_{\bar{\pi}}^{\bar{B}}$ is smaller than μ at $\tilde{\pi} = 0$ and larger than μ at $\tilde{\pi} = \pi$. Furthermore, this expectation is strictly increasing because $F_{\tilde{\pi}_1}^{\bar{B}}$ strictly first-order stochastically dominates $F_{\tilde{\pi}_2}^{\bar{B}}$ whenever $\tilde{\pi}_1 > \tilde{\pi}_2$. Since this expectation is also continuous in $\tilde{\pi}$, the Intermediate Value Theorem implies the existence of a unique $\bar{\pi}$, such that the expectation of $F_{\bar{\pi}}^{\bar{B}}$ is μ . It remains to show that for all $x \in [0, 1]$,

$$\int_0^x F_{\pi}^B(v) dv \leq \int_0^x F_{\bar{\pi}}^{\bar{B}}(v) dv.$$

If $x \leq B$ this follows from the observation that $F_{\pi}^B(v) \leq F_{\bar{\pi}}^{\bar{B}}(v)$ for all $v \leq B$ since $\bar{\pi} < \pi$. If $x > B$ then

$$\int_0^x F_{\pi}^B(v) dv = 1 - \mu - \int_x^1 F_{\pi}^B(v) dv = 1 - \mu - [1 - x] \leq 1 - \mu - \int_x^1 F_{\bar{\pi}}^{\bar{B}}(v) dv = \int_0^x F_{\bar{\pi}}^{\bar{B}}(v) dv,$$

where the second equality follows from $F_{\pi}^B(v) = 1$ if $v \geq B$ and the inequality follows from $F_{\bar{\pi}}^{\bar{B}}(v) \leq 1$.

Finally, note that

$$U(F_{\pi}^B, \pi) = \int_0^1 (v - \pi) dF_{\pi}^B(v) = \mu - \pi < \mu - \bar{\pi} = \int_0^1 (v - \bar{\pi}) dF_{\bar{\pi}}^{\bar{B}}(v) = U(F_{\bar{\pi}}^{\bar{B}}, \bar{\pi}),$$

where the second and third equalities follow from the fact that π and $\bar{\pi}$ are the lowest points of the supports of F_{π}^B and $F_{\bar{\pi}}^{\bar{B}}$, respectively. The inequality follows from $\pi > \bar{\pi}$. ■

Proof of Lemma 4. Let μ denote the expectation of G . First, we show that there exists an $l \in [\pi, 1]$ such that the expectation of $F_{\pi, l}$ is μ . Note that

$$\int v dF_{\pi, 1}(v) = \pi \leq \mu = \int v dG(v) \leq \int v dF_{\pi}(v) = \int v dF_{\pi, \pi}(v),$$

where the second inequality follows from part (i) of Lemma 1, and the last equality from $F_{\pi, \pi} = F_{\pi}$. Since $\int v dF_{\pi, l}(v)$ is continuous and strictly decreasing in l on $[\pi, 1]$, the Intermediate Value Theorem implies that there exists an l such that

$$\int v dF_{\pi, l}(v) = \mu. \tag{10}$$

Next, we prove that there exists a $z \in (0, l)$ such that $F_{\pi,l}(v) \geq G(v)$ if $v \leq z$ and $F_{\pi,l}(v) \leq G(v)$ if $v \geq z$. To this end, let

$$z = \inf \{y \in [0, 1] : F_{\pi}(v) \leq G(v) \text{ for all } v \geq y\}.$$

We have to show that $F_{\pi,l}(v) \leq G(v)$ if $v \geq z$. Note that if $v \geq l$ then

$$F_{\pi,l}(v) = F_{\pi}^1(v) \leq G(v),$$

where the inequality follows from part (i) of Lemma 1. Therefore, $z \leq l$. Suppose that $y < z$ and $F_{\pi,l}(y) < G(y)$. Since $F_{\pi,l}$ is constant on $[0, l]$ and G is increasing, it follows that $G(v) > F_{\pi,l}(v)$ for all $v \in [y, z] \subset [y, l]$ which contradicts to the definition of z . We conclude that $F_{\pi,l}(v) \geq G(v)$ on $[0, z]$.

We are ready to prove that $F_{\pi,l}$ is a mean-preserving spread of G . If $x \leq z$ then $F_{\pi,l}(v) \geq G(v)$ for all $v \in [0, x]$ and hence,

$$\int_0^x F_{\pi,l}(v)dv \geq \int_0^x G(v)dv.$$

If $x \geq z$ then

$$\int_0^x F_{\pi,l}(v)dv = 1 - \mu - \int_x^1 F_{\pi,l}(v)dv \geq 1 - \mu - \int_x^1 G(v)dv = \int_0^x G(v)dv,$$

where the inequality follows from $F_{\pi,l}(v) \leq G(v)$ on $[z, 1]$. The previous two displayed inequalities imply part (i).

To prove part (ii), first note that

$$U(F_{\pi,l}, l) = \int_l^1 (v - l) dF_{\pi,l}(v) = \int_l^1 v dF_{\pi,l}(v) - lD(F_{\pi,l}, l) = \int_0^1 v dF_{\pi,l}(v) - lD(F_{\pi,l}, l) = \mu - \pi,$$

where the third equality follows the fact that $F_{\pi,l}(v) = F_{\pi,l}(0)$ for all $v \in [0, l]$ and the last equality follows from (10) and $\pi = \Pi(F_{\pi,l})$. Finally, observe that

$$\mu - \pi \geq \int_p^1 (v - p) dG(v) = U(G, p),$$

where the inequality follows from the facts that the expectation of G is μ and that the buyer's payoff cannot exceed the first-best total surplus, μ , minus the seller's profit, π . The previous two displayed equations imply that $U(G, p) \leq U(F_{\pi,l}, l)$. The last inequality is strict whenever the total surplus in the outcome (G, p) is strictly less than μ , that is, trading is ex-post inefficient. Trading is ex-post inefficient if and only if the probability that $v \in (0, p)$ is positive, that is, $D(G, p) < 1 - G(0)$. ■

Proof of Claim 1. To prove part (i), note that since c is strictly convex, C is strictly increasing in risk and hence, by part (ii) of Theorem 2 we only need to show that

$$U(F_{\pi}^B, \pi) - C(F_{\pi}^B) \leq U(F_{\pi^*}^{B^*}, \pi^*) - C(F_{\pi^*}^{B^*})$$

for all $\pi \in [0, 1]$ and $B \in [\pi, 1]$ and the inequality is strict if $F_\pi^B \neq F_{\pi^*}^{B^*}$. Observe that

$$U(F_\pi^B, \pi) - C(F_\pi^B) = \pi \log\left(\frac{B}{\pi}\right) - \int_\pi^B c(v) dF_\pi^B(v) = \pi \log\left(\frac{B}{\pi}\right) - \left[\int_\pi^B c(v) \frac{\pi}{v^2} dv + \frac{\pi}{B} c(B) \right], \quad (11)$$

where the first equality essentially follows from (4) and the second one from the facts that F_π^B is defined by the density π/v^2 on $[\pi, B)$ and an atom of size π/B at B . The derivative of this expression according to B is $[\pi/B][1 - c'(B)]$, and hence, the buyer's payoff is strictly decreasing in B if $B \geq B^*$ and strictly increasing otherwise. This implies that for each π , the optimal B is B^* if $\pi < B^*$ and it is π if $\pi \geq B^*$. Since the buyer's payoff is negative in the latter case, we conclude that the optimal B is B^* . The second derivative of the right-hand side of (11) is $[c(\pi) - 1]/\pi$. Since $c(\pi) < 1 (= c(B^*))$ for all $\pi < B^*$, the second equality implies that the buyer's payoff is strictly concave in π , and hence, it has a unique maximizer for B^* . It remains to show that $\pi^* \leq B^*/e$. To this end, note that for all $\pi \in [B^*/e, B^*]$,

$$\begin{aligned} & U(F_\pi^{B^*}, \pi) - C(F_\pi^{B^*}) < U(F_{B^*/e}^{B^*}, B^*/e) - \int_0^1 c(v) dF_\pi^{B^*}(v) \\ & < U(F_{B^*/e}^{B^*}, B^*/e) - \int_0^1 c(v) dF_{B^*/e}^{B^*}(v) = U(F_{B^*/e}^{B^*}, B^*/e) - \int_0^1 c(v) dF_{B^*/e}^{B^*}(v), \end{aligned}$$

where the first inequality follows from $U(F_\pi^{B^*}, \pi)$ being strictly decreasing in π on $[B^*/e, B^*]$ (see the proof of Theorem 1) and the definition of C , the second inequality follows from the facts that c is increasing and $F_\pi^{B^*}$ first-order stochastically dominates $F_{B^*/e}^{B^*}$ (see Proposition 6.D.1 of Mas-Colell et al (1995)), and the equality again follows from the definition of C . This inequality chain implies that the buyer's payoff is strictly decreasing in π on $[B^*/e, B^*]$ and hence, $\pi^* \leq B^*/e$.

To show part (ii), first note that, by the arguments of the previous paragraph, the buyer's payoff from choosing F_π^B is constant on $B \in [B, 1]$ and is strictly increasing on $B \in (\pi, B]$ if $\pi < B$. This implies that for each π any $B \in [\max\{\pi, B\}, 1]$ is optimal. Again, the arguments of the previous paragraph implies that the buyer's payoff is strictly concave in π , and hence, it has a unique maximizer for B , π^* . Therefore, each CDF F_π^B , $B \in [B, 1]$, corresponds to an equilibrium outcome. In fact, any CDF G for which $G(v) = F_{\pi^*}^1(v)$ on $[\pi^*, B]$ and $G(v) \geq F_{\pi^*}^1(v)$ on $[B, 1]$ is an equilibrium CDF. ■

Proof of Claim 2. By Theorem 2 and Remark 1, each equilibrium CDF is in the set $\{F_\pi^1\}_{\pi \in [0, 1]}$. Furthermore, in the proof of Theorem 1, we have shown that the expected value of F_π^1 is $-\pi \log(\pi) + \pi$. Hence, the buyer's problem when choosing a CDF can be written as

$$\max_{\pi \in [0, 1]} U(F_\pi^1, \pi) - \tilde{c}(\mu_{F_\pi^1}) = \max_{\pi \in [0, 1]} [-\pi \log(\pi) - \tilde{c}(\pi - \pi \log(\pi))].$$

where the equality follows from $U(F_\pi^1, \pi) = \mu_{F_\pi^1} - \pi$. It remains to show that the solution to this problem is weakly smaller than $1/e$. To this end, note that for all $\pi > 1/e$

$$U(F_\pi^1, \pi) - \tilde{c}(\mu_{F_\pi^1}) < U(F_{1/e}^1, 1/e) - \tilde{c}(\mu_{F_\pi^1}) \leq U(F_{1/e}^1, \mu_{F_{1/e}^1}) - \tilde{c}(\mu_{F_{1/e}^1}),$$

where the first inequality follows from Theorem 1. The second inequality follows from the facts that $\mu_{F_\pi^1} > \mu_{F_{1/e}^1}$ since $\pi > 1/e$ and that \tilde{c} is increasing. This inequality chain implies that the buyer prefers the CDF $F_{1/e}^1$ to F_π^1 . ■

Proof of Claim 3. As mentioned after the proof of Theorem 1, if the buyer can choose any distribution supported on $[0, k]$ at no cost, then $(F_{k/e}^k, k/e)$ is the unique equilibrium outcome and the buyer's payoff is k/e . So, when determining the upper bound of the support, the buyer solves $\max_k [k/e - \hat{c}(k)]$. ■