Buyer-Optimal Demand and Monopoly Pricing[*]

Daniele Condorelli[†]and Balázs Szentes[‡]

February 5, 2016

Abstract

This paper analyzes a bilateral trade model where the buyer can choose any cumulative distribution function (CDF) supported on $[0, 1]$, which then determines her valuation. The seller, after observing the buyer's choice of the CDF but not its realization, gives a take-it-or-leave-it offer to the buyer. We characterize the unique equilibrium outcome of this game and show that in this outcome, the price and the payoffs of both the buyer and the seller are equal to $1/e$. The equilibrium CDF of the buyer generates a unit-elastic demand on $[1/e, 1]$.

1 Introduction

For a given price of a good, a buyer is better off the higher is her valuation for the good. However, a monopolist who knows that the buyer's valuation is likely to be high will charge a higher price. As a result, a buyer who could increase her willingness-to-pay stochastically would face a trade-off between the higher payoff she would receive from the good and the higher price she would have to pay. This paper analyses this trade-off without imposing any restriction on the set of value-distributions available to the buyer, except that the maximum valuation is less than one.[1]

In our model, the seller has full bargaining power. Specifically, after observing the distribution chosen by the buyer, the seller sets a price at which either the buyer trades or the game ends. We show that this game has a unique equilibrium outcome. In each equilibrium, the price is $1/e$ and the distribution of the buyer's 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

[*]We have benefited from discussions with Dirk Bergemann, Ben Brooks, Sylvain Chassang, Eddie Dekel, Jeff Ely, Faruk Gul, Yingni Guo, Sergiu Hart, Francesco Nava, Santiago Oliveros, Ludovic Renou and Phil Reny.

[†]Department of Economics, University of Essex, Colchester, UK. E-mail: dcond@essex.ac.uk.

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

[1]Requiring the upper bound to be one is simply a normalization. Such an upper bound can be interpreted, for example, as a technological frontier determining the object's maximum value.

$1/e$ 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 more than one quarter of the first-best social surplus.

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 $p[1/(ep)] = 1/e$.

The problem analysed in our paper is a hold-up one. Indeed, if the buyer's choice of the CDF was contractible, the equilibrium outcome would be efficient. Our model can therefore be motivated by the same applications as the literature on hold-up problems. Most of this literature considers the case where the buyer's investment is costly, observable and shifts her valuation deterministically. However, it seems likely that the returns to many investments are stochastic and the investor has superior information about these returns. Consider, for example, human capital investment. Prior to going on the job market, an individual decides the type and length of her education. It is conceivable that education has a stochastic impact on productivity and hence the type of education can be thought of as a distribution on productivities. In this case, an employer can offer a contract based on the observed education but not on the actual productivity. The goal of our paper is to reconsider the hold-up problem from this perspective and to analyze the buyer's incentive to generate information rent and the resulting inefficiencies.

We consider an extension of our baseline model to incorporate an additional constraint on the buyer that her expected valuation cannot exceed a given threshold. We show that, as in the baseline model, trade occurs with probability one and the buyer's demand is unit-elastic. We demonstrate that the buyer's equilibrium CDF minimizes the seller's profit among all CDFs generating the threshold expectation. Furthermore, we show that the seller's payoff converges to zero while the buyer's payoff converges to the threshold expectation as the upper bound of the CDF's support tends to infinity. As such, if the buyer is not constrained by the support of the CDFs, she can choose a value distribution which induces the seller to set an arbitrarily small price and generates the threshold expectation.

There are papers on the hold-up problem which also consider the buyer's ability to generate asymmetric information. 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. This result arises because the buyer can generate information rent by randomizing on various investment levels, since randomization generates asymmetric information conditional on the uninformative signal. In turn, the buyer's information rent makes it more profitable for her to invest. The equilibrium randomization of the buyer ensures that the seller is indifferent between setting prices on the support of valuations, as in our model. In contrast to our equilibrium, however, trade often does not occur since the seller also randomizes on the full range of valuations in order to ensure that the buyer is indifferent. 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.[2] Similarly, we demonstrate that in our model, the buyer's equilibrium payoff is zero if the seller is unable to observe the buyer's CDF.

Roesler and Szentes (2015) analyze the same bilateral trade protocol as considered here but with a different information structure. In their setup, the buyer has a given value distribution and designs a signal structure 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 (2015) 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. Roesler and Szentes (2015) show that the latter determines the buyer's equilibrium signal structure so that the equilibrium price is the lowest across all signal structures and the buyer always trades. As in our equilibrium, the seller is indifferent between any prices in the signal's support.

Our buyer-optimal distribution also appears 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 a 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.[3] 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.[4]

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

[3] Since the seller's regret is the buyer's payoff if there is trade, it is not a coincidence that nature chooses the CDF which maximizes the buyer's expected surplus.

[4] In 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

Our paper also relates to a recent literature on the importance of information structures in selling problems. For example, 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.

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 supported on $[0,1]$. The seller observes the choice of the buyer, F, and gives a take-it-or-leave-it price offer to the buyer, p. Finally, the buyer's valuation, v, is realized and she trades with the seller if and only if $v \geq p$.[5] If trade takes place, the payoff of the seller is p and that of the buyer is $v - p$; both receive a payoff of zero otherwise. The seller and the buyer are expected payoff maximizers. We restrict attention to Subgame Perfect Nash Equilibria of this game.

We introduce a few pieces of notation. 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.[6] If the buyer chooses a CDF $F \in \mathcal{F}$, then the seller's profit is $\Pi(F) = \max_p pD(F,p)$.[7] 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, $P(F) = \arg\max_p pD(F,p)$. Finally, let $U(F,p)$ denote the buyer's payoff if the seller sets price p, that is, $U(F,p) = \int_p^1 v - pdF(v)$.

3 Results

We solve our problem in two steps. First, for a given profit of the seller π, we compute the CDF $F_\pi \in \mathcal{F}$ 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 $\{F_\pi\}_{\pi\in[0,1]}$. In the second step, we characterize the profit π and the corresponding CDF F_π which maximizes the buyer's payoff.

similar to our equilibrium CDF.

[5] Assuming that the buyer is non-strategic in the final stage of the game and buys the object whenever her valuation weakly exceeds the price has no effect on our results but makes the analysis simpler.

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

[7] This maximum is well-defined because the buyer trades when she is indifferent.

To this end, we next define the CDF F_π and show that it maximizes the buyer's payoff subject to the constraint that the seller's profit is π.

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

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

Since $F_\pi(\pi) = 0$, the valuation of the buyer is never below π. The function F_π is continuous and strictly increasing on $[\pi,1)$. On this interval, F_π is defined by the decreasing density $f_\pi(v) = \pi/v^2$. Finally, there is an atom of size π at $v = 1$, that is, $\Delta(F_\pi,1) = \pi$.

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

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

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

$$F_\pi(v) \leq G(v). \tag{1}$$

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

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

Proof. To prove part (i), note that since $p \in P(G)$, it must be that for all $v \in [0,1]$,

$$vD(G,v) \leq pD(G,p).$$

or equivalently,

$$1 - \frac{pD(G,p)}{v} + \Delta(G,v) \leq G(v). \tag{2}$$

Since $\Delta(G,p) \in [0,1]$ and $\pi = pD(G,p)$, the previous inequality implies (1).

To see part (ii), note that

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

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

Let F^* and p^* denote the equilibrium CDF of the buyer and the equilibrium price of the seller, respectively. We are now ready to state our main result.

Theorem 1 *In the unique equilibrium outcome,* $F^* = F_{1/e}$ *and* $p^* = U(F^*, p^*) = \Pi(F^*) = 1/e$.

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.

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

$$U(F_\pi, \pi) = \int_\pi^1 v - \pi dF_\pi(v) = \int_\pi^1 v f_\pi(v)\,dv + \Delta(F_\pi, 1) - \pi = \int_\pi^1 \frac{\pi}{v} dv = -\pi \log \pi, \quad (3)$$

where the second equality follows from $f_\pi(v) = \pi/v^2$ and $\Delta(F_\pi, 1) = \pi$. The function $-\pi \log \pi$ is indeed maximized at $1/e$. From the equality chain (3), the buyer's payoff is $U\left(F_{1/e}, 1/e\right) = -1/e \log 1/e = 1/e$ in this equilibrium. Finally, note that $\Pi\left(F_{1/e}\right) = 1/e$.

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 $U\left(F_{\Pi(F^*)}, \Pi(F^*)\right)$. We showed that $U\left(F_{\Pi(F^*)}, \Pi(F^*)\right) \leq U\left(F_{1/e}, 1/e\right)$ in the previous paragraph, so the buyer's payoff cannot exceed $1/e$ in any equilibrium. Since $1/e$ is the unique maximizer of $U(F_\pi, \pi)$, part (ii) of Lemma 1 also implies that in order for the buyer to achieve this payoff, she must choose $F_{1/e}$ 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 $F^\varepsilon_{1/e}$ be $\varepsilon + F_{1/e}$ on $[p, 1)$ and $F_{1/e}$ otherwise. The CDF $F^\varepsilon_{1/e}$ is constructed from $F_{1/\varepsilon}$ by moving a mass of size ε from the atom at $v = 1$ to $v = 1/e$. If the buyer chooses $F^\varepsilon_{1/e}$ then the seller strictly prefers to set price $1/e$.[8] In addition, the buyer's payoff is ε-close to $U\left(F_{1/e}, 1/e\right) = 1/e$. 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

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

$F_{1/e}$ is $2/e$. Indeed, from (3),

$$\int_{1/e}^{1} vdF_{1/e}(v) = \int_{1/e}^{1} vf_{1/e}(v)\,dv + \Delta\left(F_{1/e}, 1\right) = \int_{1/e}^{1} \frac{1}{ev}dv + \frac{1}{e} = \frac{2}{e}.$$

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 $U\left(F_{1/e}, 1/e\right) + \Pi\left(F_{1/e}\right) = 2/e$. Hence, the efficiency loss due to the buyer's desire to generate information rent is $1 - 2/e \approx 0.26$.

As mentioned earlier, 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.

An Extension: Bounded Expectations

In our baseline model, the only constraint faced by the buyer is that the upper bound of the valuation-support must be smaller than one. Suppose now that the buyer faces another constraint when choosing the distribution F, namely that the expected value cannot exceed μ, that is

$$\int_{0}^{1} vdF(v) \leq \mu.$$

Let $\mathcal{F}_\mu \subset \mathcal{F}$ denote the set of CDF's satisfying this constraint. We assume that $\mu \leq 2/e$, since otherwise the new constraint does not bind and the unique equilibrium outcome is identical to the one identified by Theorem 1.

We now characterize the equilibrium CDF of the buyer, F_μ^*, and the equilibrium price, p_μ^*.

Proposition 1 *If the buyer's action space is $\mathcal{F}_\mu$ then, in the unique equilibrium outcome, $F_\mu^* = F_{\pi_\mu^*}$ and π_μ^* is defined by $-\pi_\mu^* \log \pi_\mu^* + \pi_\mu^* = \mu$.*

The expected valuation of the buyer generated by F_π is increasing in π. This expectation is $-\pi \log \pi + \pi$ by the proof of Theorem 1. The previous proposition states that the buyer chooses $F_{\pi_\mu^*}$ so that her expected valuation is exactly μ, so $F_{\pi_\mu^*} \in \mathcal{F}_\mu$. Since the seller charges π_μ^*, trade always occurs and $U\left(F_{\pi_\mu^*}, \pi_\mu^*\right) = \mu - \pi_\mu^*$.

Proof. We only show that there exists an equilibrium outcome with the properties described in the statement of the theorem. The uniqueness proof is essentially identical to that in the proof of Theorem 1.

If the buyer chooses $F_{\pi_\mu^*}$, then the seller optimally responds by setting price π_μ^* since $P\left(F_{\pi_\mu^*}\right) = \left[\pi_\mu^*, 1\right]$. We have to show that the buyer has no incentive to deviate. Let $G \in \mathcal{F}_\mu$, $p \in P(G)$ and let π denote $\Pi(G)$. We show that $U(G,p) \leq U\left(F_{\pi_\mu^*}, \pi_\mu^*\right)$. Note that

$$U(G,p) = \int_{p}^{1} v - pdG(v) = \int_{p}^{1} vdG(v) - pD(G,p) \leq \mu - \pi,$$

where the inequality follows from $G \in \mathcal{F}_\mu$ and $\pi = pD(G,p)$. If $\pi \geq \pi^*_\mu$ then this inequality chain implies $U(G,p) \leq U\left(F_{\pi^*_\mu}, \pi^*_\mu\right)$ because $U\left(F_{\pi^*_\mu}, \pi^*_\mu\right) = \mu - \pi^*_\mu$. If $\pi \leq \pi^*_\mu$ then

$$U(G,p) \leq U(F_\pi, \pi) \leq U\left(F_{\pi^*_\mu}, \pi^*_\mu\right),$$

where the first inequality follows from part (ii) of Lemma 1 and the second from $U(F_\pi, \pi)$ being strictly increasing in π on $(0, 1/e]$ (see (3)). ■

Next, we show that the CDF $F_{\pi^*_\mu}$ is not only optimal for the buyer but is also the CDF which minimizes the seller's profit among those CDFs which generate an expected valuation of exactly μ.

Remark 1 *If* $\int_0^1 vdG(v) = \mu$ *then* $\Pi(G) \geq \Pi\left(F_{\pi^*_\mu}\right)$.

This result is also reported in Neeman (2003) and Kremer and Snyder (2016).

Proof. Let π denote $\Pi(G)$. Note that

$$\int_0^1 vdF_\pi(v) \geq \int_0^1 vdG(v) = \mu = \int_0^1 vdF_{\pi^*_\mu}(v),$$

where the inequality follows from part (i) of Lemma 1, the first equality is the hypothesis of this remark, and the second equality follows from the statement of Proposition 1. Since $\int_0^1 vdF_p(v)$ is strictly increasing in p, the previous inequality chain implies that $\pi \geq \pi^*_\mu$. ■

Again, if the upper bound of the CDF's support is required to be k instead of one, the equilibrium price and the profit would be $k\pi^*_{\mu/k}$ and the equilibrium payoff of the buyer would be $kU\left(F_{\pi^*_{\mu/k}}, \pi^*_{\mu/k}\right)$. What happens if k goes to infinity for a fixed μ?

Remark 2 *For all* $\mu \in \mathbb{R}_{++}$, $\lim_{k\to\infty} kU\left(F_{\pi^*_{\mu/k}}, \pi^*_{\mu/k}\right) = \mu$ *and* $\lim_{k\to\infty} k\pi^*_{\mu/k} = 0$.

This remark implies that if the upper bound k can be arbitrarily large, the buyer extracts the full surplus, μ, by choosing a CDF which induces a price arbitrarily close to zero.

Proof. By Proposition 1, $-\pi^*_{\mu/k} \log \pi^*_{\mu/k} + \pi^*_{\mu/k} = \mu/k$ or equivalently,

$$-\log \pi^*_{\mu/k} + 1 = \frac{\mu}{\pi^*_{\mu/k} k}.$$

Since $\pi^*_{\mu/k}$ goes to zero as k goes to infinity, the left-hand side converges to infinity. Therefore, the right-hand side must also go to infinity. Hence, the price $k\pi^*_{\mu/k}$ must converge to zero, which implies that the seller's payoff goes to zero and the buyer's payoff, $kU\left(F_{\pi^*_{\mu/k}}, \pi^*_{\mu/k}\right)$, goes to μ. ■

4 Discussion

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

Unobservable Distributions.— If the seller is unable to observe the buyer's choice, the price cannot depend on the chosen distribution. Hence, the buyer always prefers stochastically higher valuations. As a consequence, the seller sets a price of one in each equilibrium and the buyer chooses a distribution which specifies a large enough mass at one.

Investment Cost.— It is straightforward to introduce some costs associated with the buyer's choice of distribution. Suppose, for example, that if the buyer chooses a distribution and the upper bound of its support is k then she has to pay an additively separable cost $c(k)$. Recall that we have shown that the buyer's payoff is k/e if she can choose any distribution supported on $[0, k]$ at no cost. So, when determining the upper bound of the support, the buyer solves $\max_k k/e - c(k)$. If, for example, $c(k) = k^2/2$, the optimal choice is $k = 1/e$.

Production Cost.— We have implicitly assumed that the seller's production cost is zero. It is straightforward to generalize Theorem 1 to the case where the seller has to pay a cost $c \in (0, 1)$ if trade occurs. One can follow the same two-step procedure to solve the problem as in Section 3. 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 $(1 - c)/e + 2c$.

Multiple Buyers.— If there is more than one buyer and the seller is restricted to set a single price then there are multiple equilibria. In each of these equilibria, the price is one and at least two buyers choose distributions which specify sufficiently large atoms at one. There is an equilibrium in which the value of at least one buyer is always one, so full efficiency is achieved. In this sense, competition eliminates the inefficiency due to the hold-up problem.

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 Karl H. Schlag, "Pricing without Priors," *Journal of the European Economic Association*, 04-05 2008, *6* (2-3), 560–569.

[3] Bergemann, Dirk and Martin Pesendorfer, "Information structures in optimal auctions," *Journal of Economic Theory*, November 2007, *137* (1), 580–609.

[4] Gul, Faruk, "Unobservable Investment and the Hold-Up Problem," *Econometrica*, March 2001, *69* (2), 343–76.

[5] 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.

[6] Hermalin, Benjamin E. and Michael L. Katz, "Information and the hold-up problem," *RAND Journal of Economics*, 2009, *40* (3), 405–423.

[7] Klein, B., Crawford, R. and Alchian, A, 1978. "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process." Journal of Law and Economics 21: 297-326.

[8] Kremer, Michael and Christopher Snyder, "Worst-Case Bounds on R&D and Pricing Distortions: Theory and Disturbing Conclusions if Consumer Values Follow the World Income Distribution," 2016, *Mimeo*, NBER.

[9] Lau, Stephanie, "Information and bargaining in the hold-up problem," *RAND Journal of Economics*, 2008, *39* (1), 266–282.

[10] Mas-Colell A., M. Whinston and J. Green, *Microeconomic Theory*, Oxford University Press, 1995.

[11] Neeman, Zvika, "The effectiveness of English auctions," *Games and Economic Behavior*, May 2003, *43* (2), 214–238.

[12] Ortner, Juan and Sylvain Chassang, "Making collusion hard: asymmetric information as a counter-corruption measure", *Mimeo*, Princeton University.

[13] Roesler A-K. and B. Szentes, "Buyer-optimal Learning and Monopoly Pricing," *Mimeo*, London School of Economics.

[14] Renou, Ludovic and Karl H. Schlag, "Minimax regret and strategic uncertainty," *Journal of Economic Theory*, January 2010, *145* (1), 264–286.

[15] Tirole, Jean, "Procurement and Renegotiation," *Journal of Political Economy*, 1986, *94* (2), 235–59.