Investing in Outside Options in Bargaining

Enrico De Magistris*

JOB MARKET PAPER

THIS VERSION: November 15, 2021

Click here for the most recent version.

Abstract

I study bargaining between a buyer and a seller when the buyer can invest in generating outside options at a cost, and the seller cannot observe his investment choice. I model the negotiation phase as an incentive compatible and ex-post individually rational direct mechanism, which maximizes a weighted average of the buyer's and seller's expected payoff. When the weight on the buyer is larger, trade happens with certainty, the price equals the seller's cost, and the buyer does not invest in generating outside options. When the weight on the seller is larger, the optimal mechanism is a posted price mechanism at the worst outside option that the buyer could have, provided that the cost function of the buyer is decreasing in first order stochastic dominance. The probability of trade is strictly below 1, even if it would be socially efficient to trade.

*Department of Economics, Boston University. I am extremely grateful to my advisors Bart Lipman, Chiara Margaria and Juan Ortner for the invaluable help, support, and guidance they have provided for this project, and throughout the years. I also thank Arjada Bardhi, Ludovica Ciasullo, Roberto Corrao, Stefania Garetto, Giacomo Lanzani, Duoxi Li, Leslie Marx, Dilip Mookherjee, Andy Newman, Jawwad Noor, Fernando Payró, Julia Salmi, Martina Uccioli, Frank Verboven, Leeat Yariv, and the participants to the BU Micro Theory workshop for useful comments and suggestions.

1 Introduction

The availability of an outside option in a bargaining situation can improve on the terms of trade for the party with such alternative opportunity. For example, a seller may be willing to offer a discount to a buyer with an alternative price offer. The seller could convince the buyer to trade by slightly undercutting the other offer, if she knew its size. But if the seller only knew that an outside option exists, the seller may offer a discount larger than that needed to trade.

Traditionally, economic theory has modelled asymmetric information about outside options using an exogenus prior distribution over the possible alternatives that the buyer could have. This approach is well suited to model a situation in which how the buyer obtained his outside option does not impact the outcome of the negotiation (Fudenberg et al., 1987). In some applications, it is more realistic to take a different approach, one in which outside options arise because of some action taken by the buyer. In other words, the buyer may have to exert effort to find a good alternative price, rather than being born with one at his disposal. In addition, the seller may not know either the outside opportunity that the buyer found, or how hard he looked for it.

As an example, consider an employer (the *buyer*, he) negotiating with a job applicant (the *seller*, she). The worker typically does not know whethe the employer has interviewed other applicants before negotiating with her, and if he found a good match willing to take the job at a low wage. When deciding how much effort to put into the search, the employer has to take into account two things. First, effort is costly: finding a good match is difficult and finding one willing to take the job for a low wage may be even harder. Moreover, the employer may be unable to foresee how good an alternative he could find before negotiating with the worker: interviewing many candidates does not guarantee finding a good one willing to accept the job at a low wage. Second, the buyer has to anticipate what effect will this outside option have when bargaining with the buyer. Specifically, having an alternative offer at hand can offset its cost by benefitting the buyer in two ways during the negotiations. On the one hand, the outside option provides him with a worst-case scenario payoff: if the worker adamantly demands a high wage, the employer knows he can resort to the alternative applicant. On the other hand, it creates the opportunity of extracting *information rents*, since the seller does not observe either his effort or the resulting alternative price: the worker may end up with a lower wage than what the employer was ready to agree on.

This paper studies the effect that the endogeneity of outside options has on the outcome of a negotiation and the distribution of bargaining power in the relation. To this end, I analyse bargaining between a buyer and a seller assuming that, prior to negotiating, the buyer can invest in generating outside options. I assume that the buyer's willingness to pay is larger than the seller's opportunity cost, so that trade is socially efficient. As a result, the buyer's investment in outside options is a form of ex-post opportunistic behavior (Klein et al., 1978; Morita and Servátka, 2018) to extract quasi-rents from the seller. I model investment as choosing a cumulative distribution function (CDF) on $[0, 1]$, where 0 is the opportunity cost of the seller, and 1 is the willingness to pay of the buyer.

We typically model bargaining between the buyer and the seller as an extensive form game with incomplete information. Such game specifies the whole *bargaining protocol*: which party makes an offer at a given point, any potential exchange of information (via cheap talk), and so on. Sometimes, there are good reasons to study a specific bargaining protocol. For example, a monopolist facing many potential customers has all the bargaining power in a negotiation, which can be modelled by assuming she makes all the offers (as in Gul et al., 1986). If we believe the parties have equal bargaining power, perhaps, it is natural to assume that they alternate in making an offer (as in Rubinstein, 1982). The choice of a specific bargaining protocol to model the negotiations inevitably adds an assumption about the distribution of bargaining power in the relationship between buyer and seller. It is therefore risky to study the impact of endogenous outside options on bargaining power within a specific protocol, as the conclusion of the analysis might depend on the protocol chosen. Therefore, rather than arbitrarily choosing a game form describing the negotiation process, I use Mechanism Design to represent its equilibrium outcome in reduced form. More precisely, by the *Revelation Principle*, for any equilibrium of any bargaining protocol, there exists a *direct revelation mechanism*, implementing the same allocation as a Bayes-Nash equilibrium in truthful reporting.

Thus, I model the negotiation stage as a direct revelation mechanism, which associates with each outside option that the buyer could report an *outcome*, specifying the probability and the price at which the buyer and the seller trade. In the main model, I assume that the bargaining mechanism is *ex-ante Pareto efficient*, given the distribution over outside options chosen by the buyer. That is, conditional on the chosen distribution neither party can obtain a larger expected payoff in an alternative mechanism without the other party getting a lower payoff.

An equilibrium of the game is then represented by a pair, consisting of the distribution over outside options and a direct revelation mechanism, with the following property. First, the distribution is optimal for the buyer, given the mechanism governing trade at the bargaining stage. Second, the mechanism is ex-ante Pareto optimal given the distribution chosen by the buyer. To find the equilibria of the model, it is useful to consider a (fictitious) game between the buyer and a (fictitious) *mechanism designer*. In this game, the buyer invests in a CDF over outside options, and the designer selects a mechanism to maximize a weighted average of the buyer's and seller's expected payoffs. Thus, when analysing the model I will use expressions like "the designer *favors* the buyer", if the buyer's weight is larger than the seller's, or that a Pareto efficient mechanism for distribution F is a *best-reply* of the designer to F. I emphasize here that this is just a metaphor, albeit a useful one: the true game is between the buyer and the seller, and the mechanism is just a reduced form representation of its equilibrium.

Perhaps surprisingly, in equilibrium the buyer is unable to extract any information rent from the seller. If the buyer's Pareto weight is larger, in equilibrium he obtains the whole surplus, without having to invest in generating outside options (Theorem 2). If that is the case, there is no efficiency loss, and the social welfare reaches its maximum level. Since this result is robust to all variations of the main model I consider (see Section 6), in the rest of the introduction I will only discuss on the case in which the seller's Pareto weight is larger. If the seller is favored, then the buyer invests in a random outside option with a two-point support, taking either value 0, or 1. The buyer and the seller only trade when the buyer's alternative price is 1, so that the seller matches the high outside option. When the two do not trade, the buyer simply exercises his outside option. Consequently, the buyer is unable to obtain any information rents from the seller: either the seller extracts his full surplus, or he has to resort to a (costly) outside option (Theorem 3).

The crucial assumption behind this result is that the investment cost in decreasing in First Order Stochastic Dominand (FSD). Using the employer-job applicant example, this assumption corresponds to the intuitive idea that it is easy to find someone willing to work for a high wage, whereas finding someone who would accept a low wage may require interviewing many candidates. While this assumption is not enough to fully characterize the equilibrium, it implies (Lemma 1) that the seller extracts all the buyer's surplus whenever there is trade. To characterize the equilibrium, I also assume that the cost function is decreasing in Mean Preserving Spread (MPS), which implies that the buyer invests in a distribution supported on the extreme points of the interval, 0 and 1. Monotonicity in MPS means that a distribution with some variability around its mean should cost more than one which is more concentrated around the same mean.

One may wonder if the absence of information rents is determined by the unobservability of investment. In a similar setting, (Condorelli and Szentes, 2020, CS hereafter) show how the buyer generates information rents if he invests in valuations for the good and the seller observes his investment choice before making a take it or leave it offer to the buyer. When the investment generates outside options, instead, the buyer in equilibrium does not obtain any information rents. Rather, he invests in the same two-point distribution he would use under unobservable investment.

The non-contractibility of the investment decision of the buyer creates a *hold-up problem* inefficiency: the buyer's investment level is above the socially optimal one. This inefficiency arises whether or not the seller observes the buyer's action. The size of the inefficiency is determined by how costly is it for the buyer to find an alternative outside option. But how large is the welfare loss due to the rent-seeking behavior of the buyer in this model? To study this issue, in Section 5.2 I analyse a variation of the model in which the investment choice is contractible and observable. Specifically, the buyer signs a binding contract with the designer prior to his investment choice, which the designer observes. The contract that the parties sign specifies the rules governing trade that would follow each CDF that the buyer could choose. In the optimal contract, the buyer is rewarded with a low price if he does not invest in outside options, and punished with a price of 1 otherwise. Thus, the buyer has to decide whether to invest in generating outside options, at the cost of not trading with the seller, and not investing and receiving some surplus within the trade relationship. In equilibrium, the buyer is offered a price which makes him indifferent between accepting that price without looking for an outside option, versus investing at the cost of this relationship.

The paper is organized as follows. Section 2 reviews the literature. Section 3 introduces the main model, with unobservable investment, and the notation. Section 4 studies its equilibria. Section 5 analyses two variations of the model in which the investment choice is observed by the seller. In particular, 5.1 studies what happens if investment is still not contractible, while 5.2 analyses the best possible contract assuming that the buyer and the seller can agree on a price before the buyer invests in generating outside options. Section 6 studies the robustness of the conclusion of the main model to weakening of the assumptions on the cost function. Section 7 concludes.

2 Related Literature

This paper contributes to, and combines approaches from, three main strands of the literature: the *hold-up problem*, *bargaining with outside options*, and *mechanism design with endogenous private information*.

Hold-up problem

The hold-up problem has long been recognized as a source of efficiency loss in bargaining and the organization of the firm (Klein et al., 1978; Grossman and Hart, 1986). In its basic version, a party invests less than the efficient level because another party is the residual claimant of the returns from investment. Many hold-up models posit that a buyer invests in his valuation of the good prior to negotiating with the seller. Different papers make different assumptions about how investment translates into valuations. Some assume a deterministic link (Grossman and Hart, 1986), some that the valuation gets determined stochastically (Hermalin and Katz, 2009), while in Condorelli and Szentes (2020)'s model the buyer chooses the (stochastic) technology that generates valuations. In all these examples, the buyer invests less than the socially optimal level. Indeed, he anticipates the fact that the seller will appropriate (part of) the returns from investment, since its cost is sunk at the bargaining stage. Several solutions for the hold-up problem have been proposed over the years: allocation of property rights (Hart and Moore, 1990), contractual solutions, depending on informational asymmetries (Rogerson, 1992), repeated investment through a long relationship (Che and Sákovics, 2004), unobservability of investment together with frequently repeated trade (Gul, 2001), just to name a few. Some papers also consider what happens if the buyer's investment impacts his outside option. In De Meza and Lockwood (1998)'s model, investment impacts both the trade surplus that can be achieved (which determines the buyer's *inside* option) as well as what he could obtain by not trading with the seller (his *outside* option). Indeed, as my paper shows, investment in outside options can be seen as a form of rent-seeking behavior, and the equilibrium investment level is above the socially optimum level.

Bargaining with outside options

While outside options have long been recognized as a source of bargaining power, incorporating them in formal models presents difficulties and ambiguities.

For instance, a seller may have an outside option either because she could trade with an alternative partner, or because she could consume the good herself (Fudenberg et al., 1987; Shaked and Sutton, 1984). Despite these ambiguities, outside opportunities have been used in a variety of settings. For example, they neutralize the effects of crazy types *à la* Abreu and Gul (2000) in Compte and Jehiel (2002), and imply the failure of the Coase conjecture (Coase, 1972) in Board and Pycia (2014). While all these papers assume that the outside option is constant over time, this needs not be the case. For example, traders in Wolinsky (1987) can interrupt the negotiation to look for an alternative partner; the buyer in Hwang and Li (2017) receives the outside option at a random time during the bargaining process; and the outside option itself changes stochastically over time in McClellan (2021). I contribute to this literature by showing how the investment in outside options impacts the bargaining outcome when the bargaining process is ex-ante efficient.

Mechanism design with endogenous private information

The Harsanyi type space approach models incomplete information as an *exogenous* component of the model, and has been used in most of the classical papers in mechanism design (e.g. Myerson, 1981; Myerson and Satterthwaite, 1983). When treating information as an endogenous component, at least two approaches emerge.

Under the first approach, the state of the world is an exogenous random variable, and agents can choose a signal containing information about it. In the early literature this choice was constrained in several possible ways. For instance, in Persico (2004), there is a unique signal that the agent can acquire, or not acquire. In Shi (2012), the signals that can be chosen are ordered by the monotone likelihood ratio property.[1] The recent literature has mostly focused on *fully flexible* information acquisition, assuming that the agent can choose *any* signal correlated with the state of the world. For example, in Roesler and Szentes (2017) and Ravid et al. (forthcoming) the state of the world is the buyer's valuation of the good, and the buyer can choose how precisely to learn about it before trading with the seller. Mensch (2020) analyses how a designer can influence the learning strategy of the agent by announcing a mechanism prior to his choice of a signal.

Under the second approach, the state of the world is enogenously determined by the agents' actions. Private information arises either because the principal does not observe the agent's action, as in the auction model of Gershkov et al. (2021), or because the mapping from actions to states is stochastic, and only the agent observes the realized state (as in Condorelli and Szentes, 2020). This is also the approach of the current paper.

I contribute to this literature by studying how *endogenous* private information about outside options differs from endogenous private information about valuations in a bargaining model.

3 The Model

3.1 Mathematical preliminaries and notation

A *cumulative distribution function* (CDF) on $\mathbb{R}$ is a weakly increasing and right continuous function such that $\lim_{x\to-\infty} F(x) = 0$, $\lim_{x\to\infty} F(x) = 1$. The *support* of F is the set (see Chung, 2001)

$$\operatorname{supp} F = \{x \in \mathbb{R} : \forall \epsilon > 0, F(x + \epsilon) - F(x - \epsilon) > 0\}.$$

[1] Bergemann and Välimäki (2006) surveys the early literature on information acquisition in mechanism design.

The set of all CDFs with support contained in $[0,1]$ is denoted $\mathcal{D}[0,1]$, and is endowed with the L_1 distance:

$$d_1(F,G) = \int_0^1 |F(x) - G(x)|\mathrm{d}x.$$

Classical results by Machina (1982) show that d_1 metrizes the weak convergence topology of probability measures, so that $(\mathcal{D}[0,1], d_1)$ is a compact and convex subset of $L_1[0,1]$, the vector space of all real valued functions on $[0,1]$ whose absolute value is Lebesgue integrable.

Given two CDFS $F, G \in \mathcal{D}[0,1]$, say that F *first order stochastically dominates* G, written $F \gtrsim_{FSD} G$ if for all $x \in [0,1]$, $F(x) \leqslant G(x)$. Equivalently (Quirk and Saposnik, 1962), $F \gtrsim_{FSD} G$ if and only if for all (weakly) increasing functions $u : [0,1] \to \mathbb{R}$,

$$\int_0^1 u(x)\mathrm{d}F(x) \geqslant \int_0^1 u(x)\mathrm{d}G(x).$$

Say that F is a mean preserving spread of G, written $F \gtrsim_{MPS} G$ if for all $x \in [0,1]$,

$$\int_0^x F(y) - G(y)\mathrm{d}y \geqslant 0,$$

with equality at $x = 1$. Equivalently (Rothschild and Stiglitz, 1970), $F \gtrsim_{MPS} G$ if and only if $\mathbb{E}_F b = \mathbb{E}_G b$ and, for all (weakly) increasing and (weakly) concave $u : [0,1] \to \mathbb{R}$,

$$\int_0^1 u(x)\mathrm{d}F(x) \leqslant \int_0^1 u(x)\mathrm{d}G(x).$$

Finally, I use $H_b \in \mathcal{D}[0,1]$ to denote the CDF concentrated at $b \in [0,1]$:

$$H_b(x) = \begin{cases} 0 & x < b, \\ 1 & x \geqslant b. \end{cases}$$

It is easy to see that H_b is a mean preserving *contraction* of any distribution with average b. On the opposite extreme of the MPS order there is the distribution supported on $\{0,1\}$, $(1-b)H_0 + bH_1$, which is a mean preserving spread of any other CDF with mean b. These two CDFs play an important role in the analysis of Section 4.

3.2 Model

A buyer (he) and a seller (she) negotiate for the ownership of an indivisible good. The buyer's valuation, equal to 1 and the seller's marginal cost, equal to 0, are common knowledge. Prior to the negotiation, the buyer can exert effort in order to obtain an *outside option*, which can be interpreted as a price offer from an alternative vendor.[2] The seller does not observe the effort exerted by the buyer or the value of the outside option, when the negotiation starts.

The game is divided in two parts: an investment phase and a bargaining phase. In the *investment phase*, the buyer chooses a costly cumulative distribution function (CDF) over outside options. The buyer can choose any distribution function F from some collection of CDFs, $\mathcal{F} \subseteq \mathcal{D}[0,1]$. A cost function $C : \mathcal{D}[0,1] \to \mathbb{R}_+$ associates with each distribution F the cost of investing in it, $C(F)$. Assuming that C be defined on the whole space $\mathcal{D}[0,1]$ is without loss of generality, as one can simply assign an arbitrarily large value to $C(G)$ whenever $G \notin \mathcal{F}$. A *technology* is a pair composed of the cost function C, and the set of distributions available to the buyer, $\mathcal{F}$.

For an illustration, consider again the employer-job applicant example of the Introduction. Assume that there is a large mass of candidates that the employer could interview, and a corresponding distribution of wages, $\Phi \in \mathcal{D}[0,1]$, that they would accept. For simplicity, the employer decides the number n of candidates to interview, so that his outside option is the lowest wage that another candidate would accept. Thus, the buyer could generate the distribution of the minimum of n random draws from Φ, which I denote by Φ_n. The cost of distribution Φ_n is the opportunity cost of the employer of making n interviews.

Throughout the paper, I will maintain three minimal assumptions on the technology $(C, \mathcal{F})$. First, I assume that $\mathcal{F}$ is closed and C, restricted to $\mathcal{F}$, is continuous. Continuity of $C_{|\mathcal{F}}$ corresponds to the idea that similar distributions should have similar cost. Notice that, since $\mathcal{D}[0,1]$ is compact, so is $\mathcal{F}$. Second, I

[2]It is easy to specify the model alternatively, by assuming that b is the outside option payffo that the buyer receives, rather than an alternative price.

assume that $\mathcal{F}$ is a convex set, and $C_{|\mathcal{F}}$ a convex function. This is just a *reduction of compounded lotteries* assumption: if F and G are both elements of $\mathcal{F}$, then the buyer can generate the distribution $\gamma F + (1-\gamma)G$ by playing a mixed action, choosing F with probability γ and G with probability $(1-\gamma)$, so that assuming that $\gamma F + (1-\gamma)G \in \mathcal{F}$ is without loss of generality. It is also without loss to assume that $C_{|\mathcal{F}}$ is convex, since otherwise the buyer could generate the distribution F by playing the mixed action with the least expected cost that generates F. Finally, I assume that the buyer can always generate an outside option equal to his valuation, at 0 cost. This is modelled by assuming that the distribution concentrated at 1, denoted H_1, is an element of $\mathcal{F}$ and $C(H_1) = 0$.

Once distribution F has been chosen, the buyer observes the realization of a random variable distributed according to F, which represents his outside option. The seller observes neither the investment choice of the buyer, nor the realized outside option. At this point, the game moves to the next stage: the *negotiation phase*. As mentioned in the introduction, in principle, the negotiation phase could be modelled as an extensive form game with incomplete information, which would specify the whole protocol. Instead, I use a reduced form approach, and model bargaining as a *direct revelation mechanism*, which satisfies three conditions. First, the mechanism is ex-ante Pareto optimal. The second requirement is *Incentive Compatibility*: the buyer should prefer reporting his outside option b truthfully rather than lying and reporting another value b'. The third and final requirement is *Ex-post Individual Rationality*: the seller should prefer trading to not trading, the buyer should prefer trading to exercising his outside option, whenever trade has positive probability. This condition corresponds to the idea that trade is always voluntary, and the buyer could interrupt the negotiations to exercise his outside option, if that benefits him. To find the equilibria of the game between the buyer and the seller, I study an auxiliary game between the buyer, who chooses a distribution over outside options, and a (fictitious) mechanism designer (she), who chooses the mechanism governing trade. The payoff function of the designer is a weighted average between the payoffs of the buyer and the seller. The equilibria of this auxiliary game correspond to the (Pareto efficient) equilibria of the investment-then-bargaining game between the buyer and the seller.

The set of actions of the designer in the auxiliary game is the set of direct revelation mechanisms that are incentive compatible and ex-post individually rational. A *revelation mechanism* is composed of two elements: a closed set X of possible reports by the buyer, and a pair of functions $(q, p) : X \to [0, 1]^2$. The value $q(x)$ specifies the probability of trade if the buyer reports x, whereas $p(x)$ is the price that the buyer pays to the seller if trade happens.[3] In a *direct* revelation mechanisms the set of reports X is the subset of $[0, 1]$ of all types deemed possible by the designer's belief.[4] By reporting $x \in X$, a buyer with outside option b obtains an expected payoff of

$$q(x)(1 - p(x)) + (1 - q(x))(1 - b) = 1 - b + q(x)(b - p(x)),$$

whereas the seller's payoff is

$$q(x)p(x).$$

A direct revelation mechanism defined on $X \subseteq [0, 1]$ is *incentive compatible* (IC) if truthful reporting is optimal for the buyer: for each $b, b' \in X$,

$$1 - b + q(b)(b - p(b)) \geqslant 1 - b + q(b')(b - p(b')).$$

It is *ex-post individually rational* (epIR) if players prefer trading to not trading whenever the probability of trade is positive: for each $b \in X$,

$$q(b) > 0 \implies 0 \leqslant p(b) \leqslant b.$$

The designer selects an incentive compatible and ex-post individually rational mechanism to maximize a weighted average of the traders' payoff, given her belief about the distribution of the buyer's outside options. Let $\tilde{F}$ be the CDF

[3] Remember that the mechanism has to be ex-post individually rational, which is why $p(x) \in [0, 1]$.

[4] More precisely: X is a closed subset of $[0, 1]$ which coincides with the support of the belief of the designer.

that the designer *thinks* that the seller has selected, and let $X = \operatorname{supp} \tilde{F}$. The designer's expected payoff from selecting an incentive compatible and ex-post individually rational mechanism $(q,p) : X \to [0,1]^2$ is

$$W_\alpha(q,p) = \alpha \underbrace{\int_0^1 q(b)p(b)\mathrm{d}\tilde{F}(b)}_{\text{Seller's payoff}} + (1-\alpha) \underbrace{\int_0^1 \Big(1 - b + q(b)\big(p(b) - b\big)\Big)\mathrm{d}\tilde{F}(b)}_{\text{Buyer's payoff}},$$

where $\alpha \in [0,1]$ is a parameter capturing how much the designer "cares" about the seller. The *best reply* of the designer to the buyer choosing distribution $\tilde{F}$ is the solution of the following program:

$$\begin{aligned}
\max_{(q,p):X\to[0,1]^2} \quad & \alpha\int_0^1 q(b)p(b)\mathrm{d}\tilde{F}(b) + (1-\alpha)\int_0^1 \Big(1 - b + q(b)\big(p(b) - b\big)\Big)\mathrm{d}\tilde{F}(b), \\
\text{s.t.} \quad & \forall b, b' \in X, \;\; 1 - b + q(b)(b - p(b)) \geqslant 1 - b + q(b')(b - p(b')), \\
& \forall b \in X, \;\; q(b) > 0 \implies 0 \leqslant p(b) \leqslant b.
\end{aligned}$$

The buyer has to take a decision at three points in the game. First, he has to select a distribution $F \in \mathcal{F}$. Second, once he has observed his outside option and he has been offered a mechanism, he has to decide whether or not to participate. Finally, provided he takes part in the mechanism, the buyer chooses which report to communicate to the designer.

Consider the buyer's problem, starting from this last decision. Suppose the buyer's outside option is b, and that he's participating in the mechanism $(q,p) : X \to [0,1]^2$, where X is a closed subset of $[0,1]$. The buyer chooses report $x \in X$ to solve

$$\max_{x\in X} 1 - b + q(x)(b - p(x)).$$

If the buyer's outside option, b is deemed possible by the mechanism (that is, $b \in X$), and if the mechanism is incentive compatible, then b solves this maximization problem. However, since the designer does not observe the distribution chosen by the buyer, it is possible that $b \notin X$; that is, b may not be one of the types that the designer thought possible when choosing the mechanism. In that case, the solution to the maximization problem will be some element $b' \in X$. It is also possible that participating in the mechanism is suboptimal for the buyer as no report gives an expected payoff larger than $1 - b$. That is, if

$$1 - b > \max_{x \in X}\{1 - b + q(x)(b - p(x))\},$$

then the buyer is better off not taking part in the mechanism.

At the time of the first decision, the buyer does not know what mechanism will be offered by the designer. So suppose he believes that mechanism $(q, p) : X \to [0, 1]^2$ will be the designer's choice. Then the largest expected payoff the buyer can obtain after playing distribution $F \in \mathcal{F}$ is

$$\int_0^1 \max\Big\{\max_{x \in X}\Big\{1 - b + q(x)\big(b - p(x)\big)\Big\}, 1 - b\Big\}\mathrm{d}F(b) - C(F).$$

The *best reply* of the buyer to mechanism $(q, p) : X \to [0, 1]^2$ is the solution to the following program:

$$\max_{F \in \mathcal{F}} \int_0^1 \max\Big\{\max_{x \in X}\Big\{1 - b + q(x)\big(b - p(x)\big)\Big\}, 1 - b\Big\}\mathrm{d}F(b) - C(F).$$

Definition 1. An *Equilibrium* is a pair $(F^*, (q^*, p^*))$ such that

- $(q^*, p^*) : \operatorname{supp} F^* \to [0, 1]^2$ is incentive compatible and ex-post individually rational,
- (q^*, p^*) is a best reply to F^* for the designer, and
- F^* is a best reply to (q^*, p^*) for the buyer.

The first formal result of the paper shows that, without loss of optimality, one can restrict attention to the class of *posted price mechanisms*. More precisely, the designer's best reply to CDF F always contains at least one posted price mechanism. This is an extension of classical results in the literature to allow for arbitrary cumulative distribution functions.

Definition 2. Let X be a closed subset of $[0, 1]$. A posted price mechanism (on X) is a mechanism $(q, p) : X \to [0, 1]^2$ such that there exists $\tilde{p} \in [0, 1]$, and $\tilde{x} \in X$ with

$$q(x) = \begin{cases} 1 & x \geqslant \tilde{x} \\ 0 & \text{else,} \end{cases} \qquad \text{and} \qquad p(x) = \begin{cases} \tilde{p} & x \geqslant \tilde{x} \\ 0 & \text{else.} \end{cases}$$

If either $\tilde{p} = \tilde{x}$, or $\tilde{p} = 0$ and $\tilde{x} = \min\{x \in X\}$, then the mechanism is called a *posted price at* $\tilde{p}$.

Theorem 1. *Fix* $F \in \mathcal{D}[0,1]$*. Then there exists* $p \in \operatorname{supp} F$ *such that the posted price mechanism at* p *is a best reply to* F.

Proof. See the Appendix. □

So by Theorem 1, it is without loss to assume that the Designer posts a price, and the buyer decides whether or not to trade at that price. Indeed, suppose that the designer plays a posted price mechanism at $p \in X$, and the buyer's outside option is b. Then the buyer is willing to trade at price p if and only $1 - b \leqslant 1 - p$, or, equivalently, $p \leqslant b$. This is true whether b is an element of X, the set of outside options that the Designer deems possible, or not. Indeed, if $p \leqslant b$, then the buyer reports either b, if $b \in X$, or any type in X which is allowed to trade, if $b \notin X$. If $p > b$, then the buyer can either report b (if $b \in X$), or report a type in X who doesn't trade, or even refuse to participate in the mechanism altogether. Therefore, from now I will simplify the terminology, and say that the Designer selects a price $p \in [0,1]$.

When facing a posted price of p, the buyer's expected payoff from selecting distribtion F is

$$\int_0^1 \max\{1-p, 1-b\} \mathrm{d}F(b) - C(F) = \int_0^p (1-b)\mathrm{d}F(b) + \int_p^1 (1-p)\mathrm{d}F(b) - C(F),$$

which equals

$$1 - p + \int_0^p F(b)\mathrm{d}b - C(F),$$

by integration by parts. On the other hand, the Designer's expected payoff from posting price p when the buyer chooses distribution F is

$$\alpha p(1 - F(p)) + (1-\alpha)\Big(1 - p + \int_0^p F(b)\mathrm{d}b\Big).$$

Thus, a pair (F^*, p^*) is an *equilibrium* if

1. Distribution F^* is optimal for the buyer, given that the Designer posts price p^*. That is, F^* is a solution to

$$\max_{G \in \mathcal{F}} 1 - p^* + \int_0^{p^*} G(b)\mathrm{d}b - C(G).$$

2. The posted price mechanism at p^* is optimal for the Designer, if the distribution of outside options of the buyer is F^*. That is, p^* solves

$$\max_{p\in[0,1]} \alpha p(1 - F^*(p)) + (1 - \alpha)\Big(1 - p + \int_0^p F^*(b)\mathrm{d}b\Big).$$

Before moving to the analysis of the general model in the next Section, consider again the employer-job applicant example. The employer decides the number of candidates to interview before negotiating with the job applicant. Suppose for simplicity that interviewing n candidates has an associated cost kn, where $k \in \mathbb{R}$ is a constant. If the employer interviews n candidates, and the i-th one asks for wage b_i, then the employer's outside option is $b = \min_{i=1,\dots,n} b_i$. Therefore, this process generates the distribution

$$\Phi_n(b) = \Pr_\Phi(\min_{i=1,n} b_i \leqslant b) = 1 - (1 - \Phi(b))^n,$$

where Φ is the underlying distribution of wages asked by the alternative candidates. If the buyer invests in a sample of size 0, his outside option is 1 with probability 1. So after accounting for the possibility of mixed strategies of the buyer, the set of distributions that the buyer can choose from is

$$\mathcal{F} = \overline{\mathrm{conv}}\left(\{1 - (1 - \Pr(\cdot))^n\}_{n\in\mathbb{N}} \cup \{H_1\}\right),$$

where $\overline{\mathrm{conv}}(S)$ is the closed convex hull of the set S. A (pure) equilibrium is then constituted by a pair (Φ_{n*}, p^*) where n^* solves the buyer's problem when the Designer posts price p^*:

$$n^* \in \mathrm{argmax}_{n\in\mathbb{N}}\, 1 - \int_0^{p^*} (1 - \Phi(b))^n \mathrm{d}b - kn,$$

and p^* solves the Designer's problem when the buyer's sample contains n^* prices:

$$p^* \in \mathrm{argmax}_p\, \alpha p(\Phi^{n^*}(p)) + (1 - \alpha)\Big(1 - \int_0^p (1 - \Phi(b))^{n^*}\mathrm{d}b\Big).$$

As an illustration, consider the case where the underlying distribution is a binomial distribution on $\{0, 1\}$, with $\varphi = \Pr_\Phi(1)$. That is,

$$\Phi(b) = \begin{cases} 0 & b < 0, \\ 1 - \varphi & 0 \leqslant b < 1, \\ 1 & b \geqslant 1. \end{cases}$$

Suppose that the Designer cares more about the seller than the buyer ($\alpha < \frac{1}{2}$). Since, regardless of the sample dimension of the buyer, his outside option will be either 0 or 1, in any equilibrium the Designer chooses price $p^* = 1$. This means that the equilibrium sample dimension, n^*, must solve

$$\max_n 1 - \int_0^1 \big(1 - (1-\varphi)^n\big)\mathrm{d}b - kn = \max_n 1 - \varphi^n - kn.$$

The function $f(x) = 1 - \varphi^x - kx$ is concave on $\mathbb{R}$, and is maximized at

$$x^* = \log\left(\frac{k}{-\log\varphi}\right)\frac{1}{\log\varphi}.$$

Therefore, the solution to the buyer's problem will either be $n = \lfloor x^* \rfloor$ or $n = \lceil x^* \rceil$, where $\lfloor x \rfloor$ is the largest integer below x, and $\lceil x \rceil$ is the smallest integer above x.

4 Analysis and Equilibrium

In order to fully characterize the equilibrium of the model, I impose three assumptions. First, I assume that the buyer can choose *any* distribution supported on $[0,1]$.

Assumption 1. $\mathcal{F} = \mathcal{D}[0,1]$.

Second, I assume that the cost function C is decreasing in first order stochastic dominance. This corresponds to the idea that it is easy to find someone willing to sell an object for a high price, and hard to find a low price.

Assumption 2. C is strictly decreasing in First Order Stochastic Dominance: if $F \gtrsim_{FSD} G$ then $C(F) \leqslant C(G)$, and the inequality is strict whenever $F \neq G$.

Third, the cost function is decreasing in Mean Preserving Spread. This corresponds to the idea that it should be easier to obtain a distribution with some dispersion around the mean than one that attains the mean with high probability.

Assumption 3. C is strictly decreasing in Mean Preserving Spread: if F is a mean preserving spread of G, then $C(F) \leqslant C(G)$, and the inequality is strict whenever $F \neq G$.

A consequence of Assumption 2 is that H_1, the distribution concentrated on 1, is the only distribution which costs 0. The following three simple examples show that the requirements of Assumptions 1-3 can be satisfied by simple functional forms.

Example 1. Let $\tilde{\kappa} : [0,1] \to \mathbb{R}_+$ be strictly decreasing and strictly concave, with $\tilde{\kappa}(1) = 0$. Then the function $\tilde{C} : \mathcal{D}[0,1] \to \mathbb{R}_+$ defined by

$$\tilde{C}(F) = \int_0^1 \tilde{\kappa}(b)\mathrm{d}F(b)$$

is continuous (by Portmanteau theorem), and satisfies Assumptions 2 and 3. By linearity of the Riemann-Stieltjes with respect to the integrator, the cost function C is affine (hence, convex).

Example 2. Let $h : \mathbb{R}_+ \to \mathbb{R}_+$ be strictly increasing and strictly convex, with $h(0) = 0$, and define

$$\hat{C}(F) = h\left(\int_0^1 \hat{\kappa}(b)\mathrm{d}F(b)\right),$$

where $\hat{\kappa}$ is strictly decreasing, strictly concave and satisfies $\hat{\kappa}(1) = 0$ as in Example 1. Then $\hat{C}$ satisfies Assumptions 2 and 3, and is strictly convex:

$$\begin{aligned}
\hat{C}(\lambda F + (1-\lambda)G) &= h\left(\int_0^1 \hat{\kappa}\mathrm{d}\big(\lambda F + (1-\lambda)G\big)\right), \\
&= h\left(\lambda \int_0^1 \hat{\kappa}\mathrm{d}F + (1-\lambda)\int_0^1 \hat{\kappa}\mathrm{d}G\right), \\
&< \lambda\hat{C}(F) + (1-\lambda)\hat{C}(G).
\end{aligned}$$

Example 3. Let $\ell : [0,1] \to \mathbb{R}_+$ be continuous, convex, strictly decreasing and such that $\ell(1) = 0$. Then the function

$$\mathcal{L}(F) = \ell\Big(\int_0^1 b\mathrm{d}F(b)\Big)$$

satisfies Assumption 2, and is constant in Mean Preserving Spread: if $F \gtrsim_{MPS} G$, then $\mathcal{L}(F) = \mathcal{L}(G)$. The cost function $\mathcal{L}$ is said to be *mean based*.

In the rest of this Section, I characterize the equilibrium of the game as a function of the parameter α, capturing how much the Designer cares about the seller, relative to the buyer.

4.1 Case I: the Designer Cares More About the Buyer ($\alpha < \frac{1}{2}$)

Since the buyer's valuation is larger than the seller's cost, trade is socially optimal in this model. So, if the designer posts a price that ensures that trade happens with probability 1, then the social optimum would be achieved. If the buyer's outside option is distributed according to F, then probability 1 trade can be achieved only by posting a price smaller than $\min\{\operatorname{supp} F\}$. Moreover, the buyer may has an incentive to look for an outside option that gives him a *larger* payoff than that granted by the posted price, especially if he expects it to be high. As an extreme example, if the buyer expects the designer to post price 1, then he would have a strong incentive to look for an alternative price that does not extract his entire willingness to pay.

One possibility for the designer to achieve the social efficient allocation is to post price 0. Indeed, the buyer would effortlessly attain his maximum payoff possible, subject to the seller being individually rational, and so he would have no incentive to look for an alternative vendor.

Theorem 2. *If the Designer cares more about the buyer ($\alpha < \frac{1}{2}$), then in the unique equilibrium, the posted price is 0, the buyer chooses the distribution concentrated on 1, H_1, and trade happens with probability 1.*

Proof. First, I will show that this is an equilibrium. If the Designer posts a price of 0, then the only best reply of the buyer is to choose the cheapest distribution available: no alternative price could possibly be below 0, so the buyer should exert no effort, and choose distribution H_1. Because of Assumption 2, any other distribution that the buyer could choose costs more than $0 = C(H_1)$, and therefore H_1 is the unique best reply of the buyer. If the buyer chooses the point mass at 1, then any price p of the Designer would have trade happening with probability 1. The Designer's payoff from posting price p would be

$$\alpha p + (1-\alpha)(1-p) = (1-\alpha) + p(2\alpha - 1) < \alpha \cdot 0 + (1-\alpha) \cdot 1,$$

and so the Designer would have no incentive to deviate.

As for uniqueness: suppose by contradiction that $(\hat{F}, \hat{p})$ is another equilibrium, with $\hat{p} > 0$, so that the probability of trade is $\hat{q} = 1 - \hat{F}(\hat{p})$. I claim that $p = 0$ is a profitable deviation. Indeed, the Designer's payoff in the candidate equilibrium $(\hat{F}, \hat{p})$ is

$$(2\alpha - 1)\hat{q}\hat{p} + (1 - \alpha)\hat{q} + (1 - \alpha)\int_0^{\hat{p}} b\mathrm{d}\hat{F}(b).$$

The first term in the summation is negative, so putting $p = 0$ would (strictly) improve on it. The sum of the last two terms is dominated by

$$(1 - \alpha)\big(1 - \hat{F}(\hat{p})(1 - \hat{p})\big) \leqslant (1 - \alpha)\big(1 - \hat{F}(0)(1 - 0)\big) = 1 - \alpha,$$

so that $p = 0$ is a profitable deviation, a contradiction. □

As the second part of the proof of Theorem 2 shows, if $\alpha < \frac{1}{2}$ then in any equilibrium the designer posts price 0. The equilibrium uniqueness depends on Assumption 2, which guarantees that H_1 is the only distribution with 0 cost, and so is the only best reply to a price of 0. If Assumption 2 is violated, then it is possible that another distribution – say, G – has zero cost, in which case $(G, 0)$ would be another equilibrium. In such equilibrium too, trade would happen with certainty.[5]

4.2 Case II: the Designer Cares More About the Seller ($\alpha > \frac{1}{2}$)

The intuition behind Theorem 2 is simple: the Designer cares more about the privately informed party, so she gives him as much surplus as possible. Matters are not so simple when the Designer cares more about the uninformed party, however. In this case, the classical static monopoly trade-off between setting a large price and having a high probability of trade emerges. When optimally balancing this trade-off, the Designer must try to correctly guess which was the distribution over outside options that the buyer chose.

[5] I'm implicitly assuming here that a buyer with an outside option of 0 trades with the seller if offered price 0.

The problem is greatly simplified by the following crucial observation: any two types of the buyer who are willing to trade have the same ex post payoff. To be more precise, suppose that the posted price is $p \in [0, 1)$, and consider two outside options $b > b' > p$. Then both these type accept the price they are offered, trade, and obtain a payoff of $1 - p$. But since b and b' have the same payoff, the buyer would rather concentrate all the probability on b, the largest of these two values. This would decrease the cost of the distribution chosen by the buyer, since the resulting CDF would FSD-dominate the starting one.

Lemma 1. *Under Assumptions 1 and 2, if F is a best reply of the buyer to $p \in (0, 1)$, then F assigns zero measure to the interval $[p, 1)$, or, equivalently, $F(p) = \lim_{b\uparrow 1} F(b)$.*

Proof. Let F be a best reply to p, and suppose by contradiction that F puts some probability mass between p and 1. The buyer's expected payoff is

$$\int_0^p (1 - b)\mathrm{d}F(b) + (1 - p)(1 - F(p)) - C(F).$$

Consider the distribution $\hat{F}$ defined by

$$\hat{F}(b) = \begin{cases} F(b) & b \leqslant p, \\ F(p) & p < b < 1, \\ 1 & b \geqslant 1. \end{cases}$$

In plain English, $\hat{F}$ coincides with F below p, is constant and equal to $F(p)$ on $[p, 1)$, and has an atom of size $1 - F(p)$ at 1. Clearly, $F \gtrsim_{FSD} \hat{F}$, so that, by Assumption 2, $C(F) > C(\hat{F})$. But then $\hat{F}$ has a larger payoff than F for the buyer, a contradiction. □

The characterization of the equilibrium is basically a corollary to Lemma 1. Indeed, if the buyer has type b and accepts a price offer of p, then $b = 1$, so that the mass of traders is concentrated on an atom at 1. Consequently, the designer, who favors the seller, has an incentive to post price 1. Thus, if an equilibrium exists, then the equilibrium price equals the buyer's valuation of the good, 1. So, when $\alpha > \frac{1}{2}$, the buyer's expected payoff from trading with the seller is 0. His best reply is the distribution F that maximizes

$$\int_0^1 F(b)\mathrm{d}b - C(F) = 1 - \mathbb{E}_F b - C(F).$$

Notice that the buyer's expected payoff in this case only depends on F through its mean, and its cost. Since the buyer is risk neutral, he chooses the cheapest distribution among those that have the same mean. So define a function $c : [0, 1] \to \mathbb{R}_+$ by

$$c(\mu) = \min\{C(F) : \mathbb{E}_F b = \mu\}.$$

The value $c(\mu)$ is the least cost needed for the buyer to generate an average outside option μ.[6] It is easy to see that the function c is convex (by convexity of C) and strictly decreasing (since C is decreasing in First Order Stochastic Dominance). This suggests that the best reply to $p = 1$ can be found by solving two problems. The first problem is to compute the best mean outside option μ^*, i.e.,

$$\mu^* \in \text{argmax}_{\mu\in[0,1]}\, 1 - \mu - c(\mu).$$

The second problem is to find a distribution G^* that has mean μ^* and achieves the smallest cost among CDFs with that mean, i.e. such that $C(G^*) = c(\mu^*)$. The solution to this second problem is the distribution supported on $\{0, 1\}$ with mean μ^*, $(1 - \mu^*)H_0 + \mu^* H_1$. Indeed, this distribution is a mean preserving spread of any CDF F with average μ^*, and so, by Assumption 3, has a cost lower than $C(F)$. Notice that the best price that the Designer could offer against this two-point distribution is indeed 1. This discussion is summarized in the second main result of the paper.

Theorem 3. *Suppose that the cost function C satisfies Assumptions 1, 2, and 3. If the Designer cares more about the seller $\left(\alpha > \frac{1}{2}\right)$, then the equilibrium posted price is 1. The buyer chooses a distribution supported on $\{0, 1\}$, $(1 - \mu^*)H_0 + \mu^* H_1$,*

[6] The minimization problem in the formula for $c(\mu)$ is well defined, since C is continuous and the set of all CDFs that have average μ is a compact subset of $\mathcal{D}[0, 1]$.

where the mean μ^* solves

$$\max_{\mu \in [0,1]} 1 - \mu - c(\mu).$$

The equilibrium of Theorem 3 is not necessarily unique, however. While in any equilibrium the price must be 1, there may be several different means $\hat{\mu}$ that solve

$$\max_{\mu \in [0,1]} 1 - \mu - c(\mu),$$

since the function c is convex, but possibly not strictly convex. For each $\hat{\mu}$, the distribution supported on $\{0,1\}$ with mean $\hat{\mu}$ is a best reply to price 1, and so $\big((1-\hat{\mu})H_0 + \hat{\mu}H_1,\, 1\big)$ is an equilibrium. Consequently, the set

$$\text{argmax}_{\mu \in [0,1]} 1 - \mu - c(\mu)$$

is a compact (possibly degenerate) interval $[\mu^*, \mu^{**}] \subseteq [0,1]$.[7] Thus, the set of equilibria either contains a single element, or contains a continuum of elements indexed by the mean $\hat{\mu} \in [\mu^*, \mu^{**}]$ that the buyer chooses. Notice that the buyer obtains the same payoff,

$$1 - \hat{\mu} - c(\hat{\mu}) = 1 - \mu^* - c(\mu^*)$$

in each of these equilibria. The designer's payoff in equilibrium, instead, is

$$\alpha\hat{\mu} + (1-\alpha)(1-\hat{\mu}) = (2\alpha - 1)\mu + (1-\alpha),$$

which is strictly increasing in $\hat{\mu}$. Thus, the equilibria of the game are Pareto-ranked, and the Pareto-best one is the equilibrium in which the buyer chooses $(1-\mu^{**})H_0 + \mu^{**}H_1$.

Before concluding the Section, consider again the cost function $\tilde{C}$ of Example 1. Then the optimal mean outside option that the buyer targets, $\tilde{\mu}^*$, is either 0 if $\tilde{\kappa}(0) \leqslant 1$ or 1 otherwise. With the cost function $\hat{C}$ of Example 2, $\hat{\mu}^*$ must satisfy the first order condition

$$0 \in \partial\Big(1 - \hat{\mu}^* - h\big((1-\hat{\mu}^*)\hat{\kappa}(0)\big)\Big),$$

where $\partial(g(x))$ is the superdifferential of g at x.

[7]Compactness follows from Berge's Maximum Theorem. It is an interval since the set of maximizers of a concave function is convex.

5 Observable Investment

The model of Section 3 assumes that the investment choice of the buyer is neither observable by the Designer, nor contractible. In equilibrium, unless the designer favors the buyer, a hold-up inefficiency arises: the buyer invests too much in generating outside options, and the probability of trade is below 1. In this section, I analyse two variations of the model. In the first one, I assume that the seller can observe the distribution chosen by the buyer, but not the realized outside option. I show that, even in this case, in equilibrium the buyer chooses the same distribution supported on $\{0, 1\}$ as in Theorem 3. Thus, the hold-up inefficiency arises even if the seller observes the investment decision of the buyer. Second, I study how large is the welfare loss due to hold-up. To this end, I analyse a version of the model in which investment is contractible, so the buyer and the seller sign a contract prior to the buyer's choice. In order to avoid that moral hazard issues affect the analysis, I also assume that the seller is able to observe the buyer's choice, so the contract specifies an outcome for each distribution that the buyer could choose.

5.1 Non-contractible Investment

(Throughout this section, assume that Assumption 1 holds.) A striking implication of Theorem 2 is that the buyer does not obtain any *information rent* in equilibrium, despite being privately informed about his choice and the value of his outside option. Indeed, either there is trade and the seller extracts the whole consumer surplus, or there is no trade and the buyer exercises his outside option. Is the unobservability of his investment choice behind this lack of information rents?

In other words, if the Designer could observe the investment choice, would the buyer's equilibrium payoff be larger? The intuition would be the same as in a Stackelberg leader-follower game: by moving first, the buyer is able to dictate the best-reply of the Designer that he prefers. Notice that the buyer cannot obtain a lower equilibrium payoff with observability than without it. After all, he could always choose the two point distribution of Theorem 3. Could he obtain a *strictly* larger payoff than that?

Condorelli and Szentes (2020) study a similar model, in which the buyer invests in a distribution over *valuations* for the good. The seller observes the buyer's distribution, and makes a take-it-or-leave-it offer p to the buyer. In equilibrium, the buyer accepts the offer if and only if it is (weakly) below his realized valuation. The authors argue that if the cost function of the buyer is strictly decreasing in mean preserving spread, the buyer's equilibrium distribution belongs to a certain family of distributions. Specifically, for $\pi \in [0,1]$, and $\rho \in [\pi, 1]$ define $F_{\pi,\rho} \in \mathcal{D}[0,1]$ by

$$F_{\pi,\rho}(v) = \begin{cases} 0 & v < 0, \\ 1 - \frac{\pi}{\rho} & v \in [0, \rho), \\ 1 - \frac{\pi}{v} & v \in [\rho, 1), \\ 1 & v \geqslant 1. \end{cases}$$

The seller responds to this choice by posting price ρ, as any price in the interval $[\rho, 1]$ leave her with an expected profit of π. I will now use show that the Condorelli and Szentes (2020) result still holds when the buyer invests in outside options instead. To facilitate the comparison, I will assume that the Designer only cares about the seller ($\alpha = 1$) – or, equivalently, that the Designer *is* the seller. First, a preliminary lemma.

Lemma 2. *Suppose $G \in \mathcal{D}[0,1]$, let π denote the monopoly profit if the buyer chooses G, and p be the monopoly price. Then there exists a unique $\rho \in [\pi, 1]$ such that*

1. *$F_{\pi,\rho}$ is a mean preserving spread of G, and*
2. *The buyer's expected payoff is larger if he chooses $F_{\pi,\rho}$ and the seller posts price ρ than if he chooses G and the seller posts price p:*

$$1 - \rho + \int_0^{\rho} F_{\pi,\rho}(b)\mathrm{d}b - C(F_{\pi,\rho}) \geqslant 1 - p + \int_0^{p} G(b)\mathrm{d}b - C(G),$$

and the inequality is strict unless $G = F_{\pi_G,\rho_G}$ *for some* $\pi_G \in [0,1]$*,* $\rho_G \in [\pi_G, 1]$.

Proof. The geometric property 1 is taken *verbatim* from Condorelli and Szentes (2020), Lemma 4. As for 2, suppose G does not belong to the family, and notice that the value of the left hand side of the inequality is

$$1 - \rho + \int_0^\rho 1 - \frac{\pi}{\rho} \mathrm{d}b - C(F_{\pi,\rho}) = 1 - \pi - C(F_{\pi,\rho}).$$

Since $F_{\pi,\rho}$ is a mean preserving spread of G, then, $C(F_{\pi,\rho}) \leqslant C(G)$. Therefore it is enough to argue that

$$\pi \leqslant p - \int_0^p G(b)\mathrm{d}b.$$

Since π is the monopoly profit, $\pi = p(1 - \lim_{b\uparrow p} G(b'))$. Since G is increasing,

$$p \lim_{b'\uparrow p} G(b') \geqslant \int_0^p \lim_{b'\uparrow p} G(b')\mathrm{d}b \geqslant \int_0^p G(b)\mathrm{d}b,$$

which concludes the proof. □

Corollary 1. *Under Assumption 3, in equilibrium the buyer chooses a distribution supported on* $\{0,1\}$ *with average* μ^**, where*

$$\mu^* \in \operatorname{argmax}_\mu 1 - \mu - c(\mu).$$

Proof. Because of Lemma 2, it is without loss of generality to constrain the buyer's choice to the set

$$\mathcal{H} = \{F_{\pi,\rho} \in \mathcal{D}[0,1] : \pi \in [0,1], \rho \in [\pi,1]\}.$$

Pick any $\pi \in [0,1]$, and any $\rho \in [\pi,1]$. Then

$$1 - \pi - C(F_{\pi,\rho}) \leqslant 1 - \pi - c(\pi),$$

and the inequality is strict unless $\rho = 1$. In other words, the buyer's optimal distribution involves having $\rho = 1$, i.e. that he chooses a distribution supported on $\{0,1\}$, as desired. □

Thus, even if the seller could observe his investment decision, in equilibrium the buyer would still choose the maximally spread out distribution $(1 - \mu^*)H_0 + \mu^* H_1$ as in Theorem 2, and the seller would respond by posting price 1. Therefore, the equilibrium outcome is the same regardless of the observability of investment, and so the hold-up inefficiency arises also in this case.

Finally, observe that the buyer does not extract any information rents from the seller also in this case. The buyer prefers to rely on his outside options but rather relies on the outside options only to repay his investment choice.

5.2 Contractible Investment

Finally, I study what happens in the model if the investment choice of the buyer is both contractible, and observable. In particular, I assume that the Designer is able to offer to buyer a *complete, contingent contract* prior to his investment decision, which specifies a mechanism for each possible distribution chosen.

Consider the following timing of events. First, the Designer offers a contract to the buyer. A *contract* is a menu $\{(q_G, p_G)_{G \in \mathcal{D}[0,1]}\}$ that associates with each distribution $G \in \mathcal{F} \subseteq \mathcal{D}[0, 1]$ a direct revelation mechanism $(q_G, p_G) : \operatorname{supp} G \to [0, 1]^2$. Having been offered the contract $\{(q_G, p_G)\}_{G \in \mathcal{D}[0,1]}$, the buyer decides whether to accept it or not. If he accepts, he next chooses a distribution over outside options — say, F— which the Designer observes. Finally, the buyer privately observes his outside option, $b \in \operatorname{supp} F$, and reports a type b' to the Designer, who then conducts trade according to the rules specified by the mechanism, $(q_F(b'), p_F(b'))$.

I will restrict attention to subgame perfect equilibria in truthful reporting. That is to say, I will assume that the contract only includes incentive compatible mechanisms, and the buyer reports his type truthfully, regardless of his investment choice. In equilibrium, the Designer recognizes that the buyer will respond optimally to the contract he's offered, that is, he will invest in the distribution, and hence choose the associated mechanism, that grants him the largest expected payoff. Say that contract $\{(q_G, p_G)_{G \in \mathcal{D}[0,1]}\}$ *induces* distribution F if the mechanism associated with F, (q_F, p_F), gives the buyer the largest expected payoff within the contract:

$$F \in \text{argmax}_{G \in \mathcal{D}[0,1]} \int_0^1 \big(1 - b + q_G(b)(b - p_G(b)\big) \mathrm{d}G(b) - C(G).$$

Distribution F is *inducible* if there exists a contract inducing F. Notice that F is inducible if and only if $C(F) \leqslant \mu^* + c(\mu^*)$. Indeed, if a contract inducing F exists, then also the simple contract in which the Designer posts price 0 if F is chosen, and price 1 otherwise, induces F.[8] On the other hand, if such simple contract does not induce F, then no other contract can. Therefore, F is inducible if and only if the above simple contract induces F. For that to be the case, the payoff from choosing F, $1 - C(F)$, should be at least as large as what the buyer obtains by investing in any other distribution :

$$1 - C(F) \geqslant \max_{G \in \mathcal{D}[0,1]} \int_0^1 (1-b) \mathrm{d}G(b) - C(G) = \max_{\mu \in [0,1]} 1 - \mu - c(\mu).$$

To find the optimal contract, it is enough to answer two questions. First, which (inducible) distribution should the Designer target? Second, which mechanism should be associated with the target distribution? Indeed, the contract can be completed by simply punishing any other distribution with a posted price mechanism at 1.

In the model of Section 3, the Designer's objective function was a weighted average of the buyer's and seller's expected payoffs, where the buyer's payoff was net of the investment cost. The idea behind this modelling choice is that the investment cost is sunk when the buyer and the seller negotiate. However, when the parties sign a contract prior to the investment decision, the cost of it is not

[8] Since I will restrict the attention to ex-post individually rational mechanisms, the posted price mechanism at 1 is the harshest punishment available to the Designer to punish undesirable investment choices. In principle, the contract may specify also other punishments, for instance a large reimbursements from the buyer to the seller could be triggered if distribution F is not chosen. In that case the distribution F should still grant a payoff of at least $1 - \mu^* - c(\mu^*)$ for the buyer to enter the contract.

sunk yet, and so it should be included in the Designer's objective function.[9] The Designer's problem is

$$\max_{F,(q_F,p_F)} \quad \alpha\int_0^1 q_F(b)p_F(b)\mathrm{d}F(b) + (1-\alpha)\bigg(\int_0^1 \Big(1-b+q_F(b)\big(p_F(b)-b\big)\Big)\mathrm{d}F(b) - C(F)\bigg),$$

$$\text{s.t.} \int_0^1 \big(1-b+q_F(b)(b-p_F(b))\big)\mathrm{d}F(b) - C(F) \geqslant \max_{\mu\in[0,1]} 1-\mu-c(\mu),$$

$$\forall b,b' \in \operatorname{supp} F,\quad 1-b+q_F(b)(b-p_F(b)) \geqslant 1-b+q(b')(b-p(b')),$$

$$\forall b \in \operatorname{supp} F,\quad q(b)>0 \implies 0\leqslant p(b)\leqslant b.$$

Theorem 4. *Suppose that investment is both observable and contractible.*

- *If the Designer favors the buyer ($\alpha<\frac{1}{2}$), then she induces distribution H_1 by posting price 0.*
- *If the Designer favors the seller ($\alpha>\frac{1}{2}$), then she induces distribution H_1 by posting price $\mu^*+c(\mu^*)$.*

Off the equilibrium path, the optimal contract specifies a posted price mechanism at 1.

Proof. First, notice that by construction both mechanisms induce distribution H_1. Since trade is socially efficient in this model, posting any price that ensures trade with probability 1 maximizes the (unweighted) sum of the buyer and seller's payoffs. Let V_B^0 and U_S^0 be the buyer's and seller's payoff under the posted price 0. Let $\tilde{V}_B$ and $\tilde{U}_S$ be the buyer's and seller's payoff under an alternative ex-ante individually rational contract (possibly not inducing distribution H_1). Since trade is socially efficient,

$$U_S^0+V_B^0 \geqslant \tilde{U}_S+\tilde{V}_B.$$

Suppose first that $\alpha<\frac{1}{2}$. Since $U_S^0-\tilde{U}_S \geqslant V_B^0-\tilde{V}_B$, it follows that $\alpha(U_S^0-\tilde{U}_S)\geqslant \alpha(V_B^0-\tilde{V}_B)$, and since $\alpha<\frac{1}{2}$, $(1-\alpha)(U_S^0-\tilde{U}_S)>\alpha(V_B^0-\tilde{V}_B)$. This proves that posting a price of 0 and inducing distribution H_1 is optimal if the Designer favors the seller.

[9]This said, the inclusion of $C(F)$ in the objective function of the Designer does not impact the results.

Now suppose that $\alpha > \frac{1}{2}$, and let V_B^*, U_S^*, be the payoffs of the buyer and the seller under posted price $\mu^* + c(\mu^*)$. Analogously to the previous paragraph, $V_B^* - \tilde{V}_B \geqslant U_B^* - \tilde{U}_B$, so that $(1-\alpha)(V_B^* - \tilde{V}_B) \geqslant (1-\alpha)U_B^* - \tilde{U}_B\alpha(U_B^* - \tilde{U}_B)$, as desired. □

The optimal contract awards all the trade surplus to the buyer, if he's the preferred player. This is not surprising, as it is also what happens when investment is neither contractible nor observable. On the other hand, if the Designer favors the seller, she still has to award some trade surplus to the buyer, to avoid him looking for an alternative option.

While the analysis of this Section has been conducted assuming that the buyer could generate any distribution in $\mathcal{D}[0,1]$, the result of Theorem 4 still holds if the available set is $\mathcal{F} \subseteq \mathcal{D}[0,1]$, once the least cost of generating mean μ, $c(\mu)$, is appropriately redefined.

6 Discussion

How robust are the conclusions of Section 4 to weakening Assumptions 1-3? First, I show that nothing substantially changes if any of the Assumptions is weakened and the Designer cares more about the buyer than the seller. The proof of Theorem 2 makes it clear that neither the properties of C nor the domain of available distribution $\mathcal{F}$ matter for proving that the equilibrium price is 0. As for the choice of distribution, any distribution minimizing the cost would, together with the 0 price, form an equilibrium. If Assumption 2 holds, or, more generally, if H_1 is the only distribution with 0 cost, then there would be a unique equilibrium, $(H_1, 0)$. Either way, all equilibria are payoff equivalent: the buyer's payoff is 1, and the designer's payoff is $(1-\alpha)$.

For the rest of this Section, assume that the designer favors the buyer ($\alpha > \frac{1}{2}$). Suppose now that Assumption 1 is violated, so that the buyer has available only a (compact and convex) subset $\mathcal{F}$ of $\mathcal{D}[0,1]$, but continue assuming that C is strictly decreasing in first order stochastic dominance and mean preserving spread (Assumptions 2 and 3). The mechanics behind Lemma 1 wouldn't work without the full domain assumption, as the buyer may be unable to concentrate the probability of trade in a single type, 1. Still, if Assumption 3 holds, the best reply of the buyer to any price $p \in [0,1]$ would be undominated in the mean preserving spread order. Indeed, remember that the buyer's best reply to p solves the program

$$\max_{F \in \mathcal{F}} 1 - p + \int_0^p F(b)\mathrm{d}b - C(F).$$

If F is a mean preserving spread of G, and $F, G \in \mathcal{F}$, then $\int_0^p \big(F(b) - G(b)\big)\mathrm{d}b \geqslant 0$. Therefore,

$$1 - p + \int_0^p F(b)\mathrm{d}b - C(F) > 1 - p + \int_0^p G(b)\mathrm{d}b - C(G),$$

so that G cannot be a best reply to p.

Lemma 3. *Suppose that the set of available distributions is $\mathcal{F}$ and that the cost function C satisfies Assumption 3. If (p^*, F^*) is an equilibrium, then the only mean preserving spread of F^* is F^* itself.*

Corollary 2. *Suppose that Assumption 3 holds. If the distribution concentrated at 0 is available to the buyer, $H_0 \in \mathcal{F}$, and the designer favors the seller, then the only equilibrium is $(0, (1-\mu^*)H_0 + \mu^* H_1)$.*

After this discussion, it should also be clear that if Assumption 2 is violated, but Assumptions 1 and 3 are maintained, the equilibrium would still be $((1 - \mu^*)H_0 + \mu^* H_1, 1))$, where $\mu^* \in \operatorname{argmax}_{\mu \in [0,1]} 1 - \mu - c(\mu)$. Indeed, by Assumption 3 the buyer's best reply to any price p belongs to the set of maximally spread out distributions

$$\{F \in \mathcal{D}[0,1] : \exists \mu \in [0,1], F = (1-\mu)H_0 + \mu H_1\}.$$

Hence, in any equilibrium the designer posts price 1, so that the buyer still has to choose the mean μ^* maximizing $1 - \mu - c(\mu)$. However, assuming that the cost function be strictly decreasing in first order stochastic dominance is far from useless, as it implies (by Lemma 1) that in any equilibrium the price be 1. This is particularly relevant when Assumption 3 is weakened.

Indeed, suppose that Assumptions 1 and 2 are satisfied, but the cost function C is only *weakly* decreasing in mean preserving spread: if $F \gtrsim_{MPS} G$ then $C(F) \leqslant C(G)$ (possibly with equality). Because of Lemma 1, in any equilibrium the price would be 1. Consequently, the buyer's best reply can still be found with the two-step procedure described in Section 4.2: finding an optimal average outside option, μ^*, then choosing the distribution with least cost among those with cost $c(\mu^*)$.

Clearly, the two-point distribution $(1-\mu^*)H_0 + \mu^* H_1$ would be cost minimizing among the CDFs with average μ^*, so that $(1, (1-\mu^*)H_0 + \mu^* H_1)$ would still be an equilibrium. Suppose there is another distribution $(1, F)$. For this to happen, it must be that (1) $\mathbb{E}_F b = \mu^*$, (2) $C(F) = c(\mu^*) = C((1-\mu^*)H_0 + \mu^* H_1)$, and (3) posting price 1 is a best reply to F. Since F is a mean preserving contraction of $(1-\mu^*)H_0 + \mu^* H_1$, the probability of trade in the equilibrium $(1, F)$ would be lower than in the equilibrium $(1, (1-\mu^*)H_0 + \mu^* H_1)$. Indeed trade happens if and only if the outside option of the buyer is 1, but since F is a mean preserving contraction of $(1-\mu^*)H_0 + \mu^* H_1$, the probability that buyer's type under F would necessarily be lower than that under the two-point distribution. This discussion is summarized in the following Theorem.

Theorem 5. *Suppose that Assumptions 1 and 2 are satisfied, and that the cost function C is weakly decreasing in mean preserving spread. Then in any equilibrium $(\hat{p}, \hat{F})$, the price is 1 ($\hat{p} = 1$) and the distribution of outside options $\hat{F}$ satisfies $\mathbb{E}_{\hat{F}} b = \mu^*$ and $C(\hat{F}) = c(\mu^*)$, where $\mu^* \in \operatorname{argmax}_\mu 1 - \mu - c(\mu)$. Moreover, $(1, (1-\mu^*)H_0 + \mu^* H_1)$ is another equilibrium, which Pareto dominates $(1, \hat{F})$.*

As an application of this result, consider the mean based cost function $\mathcal{L}$ of Example 3. Since $\mathcal{L}$ is strictly decreasing in FSD and weakly decreasing in MPS, Theorem 5 applies, and the least cost of generating average μ is $\ell(\mu)$. The set of equilibria contains all pairs $(1, F)$ such that (1) $\mathbb{E}_F b = \mu^* \in \operatorname{argmax}_{\mu \in [0,1]} 1 - \mu - \ell(\mu)$ and (2) 1 is a best reply of the designer to F.

However, it is hard to further weaken Assumption 3 without risking the existence of an equilibrium. Indeed, if Assumptions 1 and 2 are satisfied, so that the buyer's choice in any equilibrium would have to be a CDF F with average μ^* and cost $C(F) = c(\mu^*)$. If, for example, C is *strictly increasing* in mean preserving spread, then such distribution would be the one concentrated at μ^*, H^*_μ. But the designer's best reply to H^*_μ would be μ^*, not 1! This implies that no (pure) equilibrium exists in this context.

Theorem 6. *Suppose that Assumptions 1 and 2 hold, and that C is strictly increasing in Mean Preserving Spread. Then no pure equilibrium exists.*

7 Conclusion

This paper studies the impact that investment in outside options can have on the outcomes of a negotiation. In the model, one of the two traders (the buyer) chooses a distribution generating a random outside option before negotiating with the other trader (the seller). The main result shows that, despite privately observing the realized value of the outside option, and even if the seller cannot observe the distribution chosen, in equilibrium the buyer is unable to extract information rents. Indeed, if the buyer has more bargaining power, then in equilibrium he does not invest in generating outside options, and the price equals the opportunity cost of the seller. Alternatively, if the seller has more bargaining power, then the price equals the willingness to pay of the buyer, and the buyer either sees all his surplus extracted, if the parties trade, or else he exercises his outside option. This result is robust to a substantial weakening of the assumptions used to derive it.

The results of the paper point to some intriguing questions. First, in the model only one party (the buyer, specifically) acquires throgh effort an outside option, while the other party can only trade with him. It is often more realistic to immagine that both parties in a relationship look around for alternative opportunities to improve on their bargaining position. Second, the benchmark given by the optimal contract of Section 5.2 is computed assuming that the effort of the buyer be both *contractible* and *observable* to the designer, who can then implement any punishment deemed appropriate in case of a deviation by the buyer. An alternative possibility regards a moral hazard framework, in which the effort of the buyer is indeed contractible, but not observable: the designer offers the contract prior to the investment choice, without the possibility of verifying it. In this case, the contract would only specify one mechanism, and the designer would decide which mechanism in order to induce a particular distribution chosen by the buyer. The problem is selecting a mechanism to influence the decision of an agent that generates his private information is also studied by Mensch (2020). The main difference with my framework is that in his model there is an underlying prior distribution of a *state of the world*, and the agent acquires a signal to learn about it.

A Appendix

In this Appendix, I prove Theorem 1, that states that for each $F \in \mathcal{D}[0,1]$ and each $\alpha \in [0,1]$, there is a posted price mechanism which is a best reply to F. The proof of the Theorem relies on the well-known convex analysis fact that a continuous and linear functional W on a compact and convex set $\mathcal{M}$ attains its maximum value at an extreme point of $\mathcal{M}$. However, some care must be put before applying this fact to the Designer's problem since her objective function is not linear in the choice variable, the mechanism. So the main proof consists in four steps. First, redefine the space of mechanisms that the Designer chooses from, so to make her objective function linear. Second, define an appropriate Banach space containing all the mechanisms the Designer could choose, and such that the Designer's objective is also continuous in the norm topology. Third, checking that the set of (direct revelation) mechanisms that are incentive compatible and ex-post individually rational is a compact and convex subset of the Banach space found above. And finally, show that the extreme points of the set of feasible mechanisms coincide with the posted price mechanisms. Throughout the proof, fix $F \in \mathcal{D}[0,1]$, the CDF the Designer is responding to.[10] I will refer to elements of $\operatorname{supp} F$ as *types*.

Step 1. Let (q,p) be a direct revelation mechanism on $\operatorname{supp} F$, and define the function $t : \operatorname{supp} F \to [0,1]$ by $t(b) = q(b)p(b)$. From now on, a *mechanism* is a pair (q,t). For technical reasons, it is useful to extend the mechanism (q,t) to the whole interval $[0,1]$, by declaring it 0 outside of $\operatorname{supp} F$. So say that a mechanism $(q,t) : [0,1] \to [0,1]^2$ is F-feasible if it is Incentive Compatible,

$$\forall b, b' \in \operatorname{supp} F, \quad 1 - b + q(b)b - t(b) \geqslant 1 - b + q(b')b - t(b'),$$

ex-post Individually Rational,

$$\forall b \in \operatorname{supp} F, \quad t(b) \in [0, q(b)b],$$

and identically 0 on $[0,1] \backslash \operatorname{supp} F$,

$$\forall x \notin F, \quad q(b) = t(b) = 0.$$

Let $\mathcal{M}(F)$ be the set of F-feasible mechanisms. The Designer's problem can then be written as

$$\max_{(q,t) \in \mathcal{M}(F)} \alpha \int_0^1 t(b) \mathrm{d}F(b) + (1-\alpha) \int_0^1 \big(1 - b + t(b) - q(b)b\big) \mathrm{d}F(b).$$

Step 2. Let $(\mathcal{B}[0,1]^2, \|\cdot\|_1)$ be the (Banach) space of all pairs of bounded real functions on $[0,1]$ endowed with the norm

$$\|(f,g)\|_1 = \int_0^1 |f| \mathrm{d}\lambda + \int_0^1 |g| \mathrm{d}\lambda,$$

where λ is the Lebesgue measure. Therefore, the functional $W_\alpha : \mathcal{B}[0,1]^2 \to \mathbb{R}$ defined by

$$(f,g) \mapsto W_\alpha(f,g) = \alpha \int_0^1 g(b) \mathrm{d}F(b) + (1-\alpha) \int_0^1 \big(1 - b + g(b) - f(b)b\big) \mathrm{d}F(b)$$

is linear and continuous in $\|\cdot\|_1$.

[10]This outlined argument is the standard one used when F admits a density– see Borgers (2015), Chapter 2. The innovation is to extend this argument to *fully general* cumulative distribution functions, not just discrete or continuous ones.

Step 3. I claim that $\mathcal{M}(F)$ is compact and convex. Convexity is easy: if $(q,t),(\tilde{q},\tilde{t}) \in \mathcal{M}(F)$, and $\xi \in [0,1]$, then for all $b,b' \in \operatorname{supp} F$,

$$\xi\big(q(b)b - t(b)\big) + (1-\xi)\big(\tilde{q}(b)b - \tilde{t}(b)\big) \geqslant \xi\big(q(b')b - t(b')\big) + (1-\xi)\big(\tilde{q}(b')b - \tilde{t}(b')\big),$$

that is, $\xi\big(q,t\big) + (1-\xi)\big(\tilde{q},\tilde{t}\big)$ is Incentive Compatible,

$$0 \leqslant \xi t(b) + (1-\xi)\tilde{t}(b) \leqslant \xi q(b)b + (1-\xi)\tilde{q}(b)b,$$

that is, it is ex-post Individually Rational, and is clearly 0 outside of $\operatorname{supp} F$. Thus, $\xi(q,t\big) + (1-\xi)\big(\tilde{q},\tilde{t}\big) \in \mathcal{M}(F)$, as desired.

In order to show that $\mathcal{M}(F)$ is compact, I will use Helly's selection theorem, which states that a uniformly bounded sequence of monotone functions admits a pointwise convergent subsequence. To apply Helly's theorem, I will first show that any F-feasible mechanism is non-decreasing on the support of F. Then, I will consider a modified version of $\mathcal{M}(F)$, which will be denoted by $\widehat{\mathcal{M}}(F)$, such that each memeber of this set is non-decreasing on the whole interval $[0,1]$, and coincides with some F-feasible mechanism on $\operatorname{supp} F$. By Helly selection theorem, $\widehat{\mathcal{M}}(F)$ is compact, and, by constructing a continuous mapping between $\widehat{\mathcal{M}}(F)$ and $\mathcal{M}(F)$ I show that also the latter is compact.

So pick an F-feasible mechanism (q,t). I claim that both q and t are non-decreasing on $\operatorname{supp} F$. Take $b,b' \in \operatorname{supp} F$, with $b \geqslant b'$. By Incentive Compatibility,

$$q(b)b - t(b) \geqslant q(b')b - t(b'), \quad \text{and} \quad q(b')b' - t(b') \geqslant q(b)b' - t(b),$$

so adding the two conditions,

$$(b-b')(q(b) - q(b')) \geqslant 0,$$

so that q is non-decreasing. Rearranging the Incentive Compatibility condition of b', I obtain

$$t(b) - t(b') \geqslant b'\big(q(b) - q(b')\big) \geqslant 0,$$

so that also t is non-decreasing on $\operatorname{supp} F$.

Define the set $\widehat{\mathcal{M}}(F)$ by

$$\widehat{\mathcal{M}}(F) = \left\{ (\hat{q},\hat{t}) : \exists (q,t) \in \mathcal{M}(F), (\hat{q},\hat{t})(b) = \begin{cases} (q,t)(b) & b \in \operatorname{supp} F, \\ \sup\{(q,t)(b') : b' \in \operatorname{supp} F \cap [0,b]\} & b \notin \operatorname{supp} F. \end{cases} \right\}$$

In other words, if $(\hat{q}, \hat{t}) \in \widehat{\mathcal{M}}(F)$, then $(\hat{q}, \hat{t})$ coincides with some F-feasible mechanism (q, t) on $\operatorname{supp} F$, and assumes the "last" value that (q, t) took, outside of $\operatorname{supp} F$. Notice that $\widehat{\mathcal{M}}(F)$ only contains pairs of non-decreasing and uniformly bounded functions. I claim that $\widehat{\mathcal{M}}(F)$ is compact. Take a sequence $(\hat{q}_n, \hat{t}_n)_{n\in\mathbb{N}} \in (\mathcal{B}[0,1]^2)^{\mathbb{N}}$ such that $(\hat{q}_n, \hat{t}_n) \in \widehat{\mathcal{M}}(F)$ for each n. By Helly's selection theorem there exists a pointwise converging subsequence $(\hat{q}_{n_k}, \hat{t}_{n_k})_{n_k}$, with limit $(\bar{q}, \bar{t})$. By the Dominated Convergent theorem, $(\bar{q}, \bar{t})$ is also the norm limit of $(\hat{q}_{n_k}, \hat{t}_{n_k})_{n_k}$. Standard argument then establish that $(\bar{q}, \bar{t}) \in \widehat{\mathcal{M}}(F)$, which is therefore compact.

Next, consider the mapping $\Gamma : \widehat{\mathcal{M}}(F) \to (\mathcal{B}[0,1])^2$, defined by

$$\Gamma(\hat{q}, \hat{t})(b) = \begin{cases} (\hat{q}, \hat{t})(b) & \text{if } b \in \operatorname{supp} F, \\ 0 & \text{else.} \end{cases}$$

By construction, $\Gamma(\hat{q}, \hat{t})$ is an F-feasible mechanism: $\Gamma(\hat{q}, \hat{t}) \in \mathcal{M}(F)$. The mapping Γ is also Lipschitz continuous, with Lipschitz constant equal to 1. Indeed, take $(\hat{q}, \hat{t}) \in \widehat{\mathcal{M}}(F)$. Then

$$\|\Gamma(\hat{q}, \hat{t})\|_1 = \int_{\operatorname{supp} F} |\hat{q}| \mathrm{d}\lambda + \int_{\operatorname{supp} F} |\hat{t}| \mathrm{d}\lambda \leqslant \int_0^1 |\hat{q}| \mathrm{d}\lambda + \int_0^1 |\hat{t}| \mathrm{d}\lambda = \|(\hat{q}, \hat{t})\|_1,$$

where the first equality follows form $\Gamma(\hat{q}, \hat{t})$ being 0 outside of $\operatorname{supp} F$. Thus, Γ is (Lipschitz) continuous, so it maps compact sets into compact sets, which proves that $\mathcal{M}(F) = \Gamma(\widehat{\mathcal{M}}(\mathcal{F}))$ is compact.

Step 4. The last step consists into proving the following statement.

Claim. *An F-feasible mechanism (q, t) is an extreme point of $\mathcal{M}(F)$ if and only if there exists $x \in \operatorname{supp} F \cup \{0\}$ such that (q, t) is a posted price mechanism (PPM) at x.*

To prove the "if" direction, consider a posted price mechanism (q, t) at $x \in \operatorname{supp} F\{0\}$. Suppose that $(q, t) = \frac{1}{2}(\bar{q}, \bar{t}) + \frac{1}{2}(\underline{q}, \underline{t})$, for some $(\bar{q}, \bar{t})$, $(\underline{q}, \underline{t})$ in $\mathcal{M}(F)$. Fix $b \in [0,1]$. Since $q(b)$ is either 0 or 1, then $\bar{q} = \underline{q} = q$. Moreover, if $\bar{t}(b)$ and $\underline{t}(b)$ do not coincide with $t(b)$, then either one– say, $\bar{t}(b)$– must be larger than

$t(b)$, violating ex-post Individual Rationality. Thus, (q,t) is an extreme point of $\mathcal{M}(F)$.

For the converse direction, fix $(q,t) \in \mathcal{M}(F)$ and suppose there is no $x \in \operatorname{supp} F \cup \{0\}$ such that (q,t) is a posted price mechanism at x. There are two possible cases.

Case I. The mechanism (q,t) is a posted price mechanism, but the price posted is not an element of the support of F, nor is 0. That is, there is $\tilde{x} \in \operatorname{supp} F$, and $y \notin \operatorname{supp} F$ such that

$$q(b) = \begin{cases} 0 & b \in [0,\tilde{x}) \cup (\operatorname{supp} F)^c, \\ 1 & b \in [\tilde{x}, 1]; \end{cases} \quad \text{and} \quad t(b) = \begin{cases} 0 & b \in [0,\tilde{x}) \cup (\operatorname{supp} F)^c, \\ y & b \in [\tilde{x}, 1]. \end{cases}$$

Since (q,t) is incentive compatible, it must be that $[y,\tilde{x}) \cap \operatorname{supp} F = \varnothing$, as otherwise a type in that interval would have an incetive to lie and report $\tilde{x}$ instead of his true type. Pick $\epsilon \in (0, \min\{y - \sup\{b \in \operatorname{supp} F : b < \tilde{x}\}, \tilde{x} - y\})$,[11] and consider the following posted price mechanisms, $(\bar{q},\bar{t})$ and $(\underline{q},\underline{t})$, defined by $\bar{q} = \underline{q} = q$ and

$$\underline{t}(b) = \begin{cases} 0 & b \in [0,\tilde{x}) \cup (\operatorname{supp} F)^c, \\ y-\epsilon & b \in [\tilde{x}, 1]. \end{cases} \quad \text{and} \quad \bar{t}(b) = \begin{cases} 0 & b \in [0,\tilde{x}) \cup (\operatorname{supp} F)^c, \\ y+\epsilon & b \in [\tilde{x}, 1]. \end{cases}$$

By construction, both $(\underline{q},\underline{t})$ and $(\bar{q},\bar{t})$ are F-feasible, and $\frac{1}{2}(\underline{q},\underline{t}) + \frac{1}{2}(\bar{q},\bar{t}) = (q,t)$, so that (q,t) is not an extreme point of $\mathcal{M}(F)$.

Case II. The mechanism (q,t) is not a posted price mechanism, and, in particular, there exists $b \in \operatorname{supp} F$ such that $q(b) \in (0,1)$. Define the following types:

- $b'_1 = \inf\{b \in \operatorname{supp} F : q(b) = 1\}$ and $b''_1 = \sup\{b \in \operatorname{supp} F : q(b) < 1\}$;
- $b'_{\frac{1}{2}} = \inf\{b \in \operatorname{supp} F : q(b) \geqslant \frac{1}{2}\}$ and $b''_{\frac{1}{2}} = \sup\{b \in \operatorname{supp} F : q(b) < \frac{1}{2}\}$.

(Of course, if the support of F has no gaps $b'_1 = b''_1$ and $b'_{\frac{1}{2}} = b''_{\frac{1}{2}}$.) As a matter of notation, let U be the expected payoff that the highest type obtains in the

[11] The interval in which ϵ lies is non-empty, since the support of F is a closed set, and therefore y, being outside of it, cannot be equal to $\sup\{b \in \operatorname{supp} F : b < \tilde{x}\}$.

mechanism (q,t), that is, $U = 1-\bar{b}_F + q(\bar{b}_F)\bar{b}_F - t(\bar{b}_F)$, where $\bar{b}_F = \max\{\operatorname{supp} F\}$. Notice that there may or may not be a type b which trades with probability 1, i.e. the set $[b'_1, 1] \cap \operatorname{supp} F$ may or may not be empty. If there is such a type, however, then he must trade at price $1-U$ by incentive compatibility.

I will now define two mechanisms, $(\underline{q}, \underline{t})$ and $(\bar{q}, \bar{t})$ with midpoint (q,t), that is, $(q,t) = \frac{1}{2}(\underline{b}, \underline{t}) + \frac{1}{2}(\bar{b}, \bar{t})$. For $\kappa \geqslant 0$, define

$$(\underline{q},\, \underline{t})(b) = \begin{cases} (0,\, 0) & b \in (\operatorname{supp} F)^c \cup [0, b''_{\frac{1}{2}}), \\ (2q(b) - 1,\, U + \kappa - 1 + 2t(b)) & b \in [b''_{\frac{1}{2}}, b'_1) \cap \operatorname{supp} F, \\ (1,\, 1 + \kappa - U) & b \in [b'_1, 1] \cap \operatorname{supp} F, \end{cases}$$

and

$$(\bar{q},\, \bar{t})(b) = \begin{cases} (0,\, 0) & b \in (\operatorname{supp} F)^c, \\ (2q(b),\, 2t(b)) & b \in [0, b''_{\frac{1}{2}}) \cap \operatorname{supp} F, \\ (1,\, 1 - U - \kappa) & b \in [b''_{\frac{1}{2}}, 1] \cap \operatorname{supp} F. \end{cases}$$

By construction, $\frac{1}{2}(\bar{q}, \bar{t}) + \frac{1}{2}(\underline{q}, \underline{t}) = (q,t)$. After considerable tedium, one can show that there exists κ that makes both $(\bar{q}, \bar{t})$ and $(\underline{q}, \underline{t})$ incentive compatible and ex-post individually rational, so that (q,t) is not an extreme point of $\mathcal{M}(F)$, as desired.

References

Dilip Abreu and Faruk Gul. Bargaining and Reputation. *Econometrica*, 68:85–117, 2000.

Dirk Bergemann and Juuso Välimäki. Information in Mechanism Design. In Richard Blundell, Whitney K. Newey, and Torsten Persson, editors, *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, volume 1 of *Econometric Society Monographs*, page 186–221. Cambridge University Press, 2006. doi: 10.1017/CBO9781139052269.007.

Simon Board and Marek Pycia. Outside Options and the Failure of the Coase Conjecture. *American Economic Review*, 104:656–671, 2014.

Tilman Borgers. *An Introduction to the Theory of Mechanism Design*. Oxford University Press, New York, NY, 2015.

Yeon-Koo Che and József Sákovics. A Dynamic Theory of Hold-up. *Econometrica*, 72:1063–1103, 2004.

Kai Lai Chung. *A Course in Probability Theory*. Academic Press, London, UK, 3rd edition, 2001.

Ronald Henry Coase. Durability and Monopoly. *Journal of Law and Economics*, 15:143–149, 1972.

Olivier Compte and Philippe Jehiel. On the role of Outside Options in Bargaining with Obstinate Parties. *Econometrica*, 70:1477–1517, 2002.

Daniele Condorelli and Balázs Szentes. Information Design in the Holdup Problem. *Journal of Political Economy*, 128(1):681–709, 2020.

David De Meza and Ben Lockwood. Does Asset Ownership Always Motivate Managers? Outside Options and the Property Rights Theory of the Firm. *Quarterly Journal of Economics*, 113:361–386, 1998.

Drew Fudenberg, David K. Levine, and Jean Tirole. Incomplete Information Bargaining with Outside Opportunities. *Quarterly Journal of Economics*, pages 37–50, 1987.

Alex Gershkov, Benny Moldovanu, Philipp Strack, and Mengxi Zhang. A Theory of Auctions with Endogenous Valuations. *Journal of Political Economy*, 129 (4):1011–1051, 2021.

Sanford J. Grossman and Oliver D. Hart. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy*, 94: 691–719, 1986.

Faruk Gul. Unobservable Investment and the Hold-up Problem. *Econometrica*, 69:343–376, 2001.

Faruk Gul, Hugo Sonnenschein, and Robert Wilson. Foundations of Dynamic Monopoly and the Coase Conjecture. *Journal of Economic Theory*, 39:155–190, 1986.

Oliver D. Hart and John Moore. Property Rights and the Nature of the Firm. *Journal of Political Economy*, 98:1119–1158, 1990.

Benjamin E. Hermalin and Michael L. Katz. Information and the Hold-up Problem. *RAND Journal of Economics*, 40:405–423, 2009.

Ilwoo Hwang and Fei Li. Transparency of Outside Options in Bargaining. *Journal of Economic Theory*, 167:116–147, 2017.

Benjamin Klein, Robert G. Crawford, and Armen A. Alchian. Vertical Integration, Appropriable Rents and the Competitive Contracting Process. *The Journal of Law and Economics*, 21:297–326, 1978.

Mark Machina. "Expected Utility" Analysis without the Independence Axiom. *Econometrica*, 50:277–323, 1982.

Andrew McClellan. Outside Options and Optimal Bargaining Dynamics. *Working paper*, 2021.

Jeffrey Mensch. Screening Inattentive Agents. *Working Paper*, 2020.

Hodaka Morita and Maroš Servátka. Investment in Outside Options as Opportunistic Behavior: An Experimental Investigation. *Southern Economic Journal*, 85:457–484, 2018.

Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

Roger B. Myerson and Mark A. Satterthwaite. Efficient Mechanisms for Bilateral Trading. *Journal of Economic Theory*, 29:265–281, 1983.

Nicola Persico. Committee Design with Endogenous Information. *The Review of Economic Studies*, 71:165–194, 2004.

James P. Quirk and Rubin Saposnik. Admissibility and Measurable Utility Functions. *The Review of Economic Studies*, 92:140–156, 1962.

Doron Ravid, Anne-Katrin Roesler, and Balazs Szentes. Learning Before Trading: On the Inefficiency of Ignoring Free Information. *Journal of Political Economy*, forthcoming.

Anne-Katrin Roesler and Balazs Szentes. Buyer-Optimal Learning and Monopoly Pricing. *American Economic Review*, 107:2072–2080, 2017.

William P. Rogerson. Contractual Solutions to the Hold-up Problem. *The Review of Economic Studies*, 59:777–793, 1992.

Michael Rothschild and Joseph E. Stiglitz. Increasing Risk: I. A Definition. *Journal of Economic Theory*, 2:225–243, 1970.

Ariel Rubinstein. Perfect Equilibrium in a Bargaining Model. *Econometrica*, 50: 97–109, 1982.

Avner Shaked and John Sutton. Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model. *Econometrica*, 52:1351–1364, 1984.

Xianwen Shi. Optimal Auctions with Information Acquisition. *Games and Economic Behavior*, 74:666–686, 2012.

Asher Wolinsky. Matching, Search and Bargaining. *Journal of Economic Theory*, 42:311–333, 1987.