Competing for Surplus in a Trade Environment

Sarah Auster Nenad Kos Salvatore Piccolo *

PRELIMINARY
May 25, 2017

Abstract

Situations in which agents exert effort or make investments to improve their bargaining position are ubiquitous: potential buyers search the internet for information on how to bargain with car salesmen, companies and individuals hire lawyers or other experts when selling or buying valuable assets, parties in a troubled marriage find divorce lawyers to negotiate on their behalf, just to name a few. We study a model of surplus division with costly effort or investments in the bilateral trade environment. The effort determines the probability with which either agents gets to make an offer and thereby the surplus they are able to secure, however, it may also signal their private information, something that the rival can use to his advantage. We characterize the equilibria of the game and show that the seller's payoff can be non-monotonic in the share of high value buyers. The seller faced with an entry decision might, therefore, find it propitious to enter a market mostly populated with low value costumers. We also develop a model in which the probability of the agents getting an offer is exogenously given—rather then determined by their decisions—but can vary with the type. The latter model is easier to use, yet captures some important features of the more involved model.

JEL Code:
Keywords: Bilateral Trade, Bargaining Power, Contests

*Auster: Bocconi University, Department of Decision Sciences and IGIER. E-mail: sarah.auster@unibocconi.it. Kos: Bocconi University, Department of Economics and IGIER. E-mail: nenad.kos@unibocconi.it, Piccolo: University of Bergamo, E-mail: salvapiccolo@gmail.com. We would like to thank Dino Gerardi, Toomas Hinnosaar, Johannes Horner, Markus Reisinger, Thomas Wiseman and seminar audiences at Frankfurt School of Management, 2016 CSEF-IGIER symposium and Workshop on Economic Theory, Graz 2016.

Contents

1 Introduction

Negotiations typically require a good amount of preparation. Cialdini and Garde (1987) provide an interesting account of tactics car-salesmen use to close sales; they also outline how to safeguard against them. In other cases, bargaining is too complicated to be undertaken on one's own: negotiating details of a business deal is often best left to lawyers, or at least undertaken under their advisory. The same is true when navigating through the unfortunate maze of a divorce.

The amount of resources people invest in improving their bargaining position depends not only on their individual characteristics, such as their innate ability to negotiate and the value attached to the deal they bargain over, but also on the rival's expected strength, and the possible information that these actions may signal. Highlighting the forces that shape these investments and the underlying trade-offs is an important step towards a better understanding of how people bargain and why they choose to do so in certain environments (markets). How much should people invest to improve their bargaining strength? What type of information do these pre-negotiation investments signal? How does this information affect the final terms of a deal and the parties' welfare? What type of environments attract better negotiators? These are just some of the questions we are trying to answer.

We address these questions by studying a model in which parties can exert costly effort in order to improve their bargaining position—e.g., the amount of resources spent to find and hire the right lawyer (or other representatives) negotiating on their behalf, the amount of attention and time they devote to specialized courses on 'the art of negotiation' (see Cialdini and Garde (1987), Fisher et al. (2011)), etc. To make our point in the simplest possible fashion, we consider a stylized bilateral trading problem between a seller and a buyer who is privately informed about his valuation. Both agents can exert effort towards winning the right to make an offer, a higher effort resulting in a higher probability of winning. The agent who receives the chance to make an offer proposes a price at which the object is traded, which the other agent accepts or rejects, thus ending the game. Agent's cannot observe each other's effort decision but can infer information about the opponent's behavior—and hence potentially about the opponent's type—by observing which party gets to make an offer.

We do not claim that this is precisely how bargaining unfolds, rather, our environment should be interpreted as a reduced form model capturing two salient features. First, if agents invest effort, they are able to obtain a higher share of surplus. This is captured through an increase in the probability that they win the opportunity to make an offer.

Second, agents learn about their opponents during the bargaining process. Specifically, if a certain type of buyer exerts less effort compared to others, then the seller is more likely to secure the chance to make an offer against that type. The revelation that the seller is the one who gets to make an offer can, therefore, be a signal of the buyer's willingness to pay. Our objective is to analyze how these two elements interplay, and what novel insights their interaction delivers.

The model is as follows. Two agents—a seller and a buyer—are negotiating a trade: the seller's value is commonly known and normalized to zero, the buyer's value is either low or high, but always above the seller's. Both agents choose between exerting and not exerting effort. Effort comes at a cost and determines the probability with which either party gets to make an offer through a contest success function. Whenever the buyer gets to make an offer, he optimally offers a price equal to zero, since the seller's valuation is commonly known. On the other hand, if the seller can make an offer, he decides between offering a low price that both types of the buyer accept and a high price that only the buyer with the high valuation accepts. Whether the low or the high price is optimal depends on the seller's posterior belief, determined by the equilibrium effort decisions of both trading parties.

The analysis starts with the observation that the high valuation buyer always has at least as high an incentive to exert effort as the low valuation buyer—a result akin of monotonicity results in mechanism design. The difference in incentives is strict if and only if the buyer expects the seller to offer the high price with strictly positive probability. When the high type exerts effort more often than the low, the seller is more likely to secure the opportunity to make an offer, thus, the right to make an offer depresses the seller's belief about the buyer's valuation being high. As a consequence, conditionally on making an offer, the seller's posterior is never above his prior.[1] .

We begin our analysis by considering a simplified version of the model in which the seller's effort is fixed. In this case, the seller's only choice is which price to offer. Whenever the buyer's effort cost is sufficiently high or low, the buyer has a dominant effort choice regardless of his type and the seller does not learn in equilibrium. His optimal pricing strategy is therefore determined by his prior belief. Of greater interest is the intermediate region of the effort costs where the low type buyer optimally exerts no effort, while the high type's effort choice depends on the seller's behavior. The analysis is split further depending on the seller's prior. When the seller's prior is low he offers the low price, neither type exerts effort, and consequently, the seller learns nothing. The most

[1] On the other hand, when the buyer gets to make an offer, the seller's posterior is weakly greater than his prior. Due to the private value assumption, however, the seller's posterior belief in this case is irrelevant.

interesting case arises for intermediate values of the prior. In this region, if the seller were to naively follow the prior he would offer the high price. This would incentivize the high type buyer to exert effort and the seller's posterior would fall. Subsequently the seller would prefer to offer the low price. The seller offering the low price, on the other hand, would make the high type buyer best respond by refraining from effort. The seller's posterior then would coincide with the prior, and the seller would in fact best respond with the high price. To escape such cycling of beliefs, the two players need to randomize: the high type buyer randomizes over the effort choices in such a way that the seller is indifferent between the two prices, the seller in response randomizes over the two prices in order to render the high type buyer indifferent between the two effort choices. Finally, for high priors, the seller offers the high price and the high type buyer exerts effort. The seller learns and revises his belief downward, but the prior is sufficiently high that even after the belief revision the high price is preferable.

We then explore how the seller's payoff varies with the proportion of the high types. We find that the payoff is constant in the fraction of high types for low priors, decreasing for intermediate priors and increasing for high priors. The striking feature that a higher proportion of the high value buyers can be detrimental to the seller's payoff arises as a consequence of the strategic interaction. In the intermediate region of priors, where buyer and seller use a mixed strategy, the probability with which the high type buyer exerts effort increases in the seller's prior: as the prior increases, the high type buyer has to separate himself further from the low type in order to keep the seller's posterior constant. The seller, therefore, gets to make an offer less often. At the same time, the seller is indifferent between the two prices, so his payoff conditionally on making an offer does not change with the prior; in fact, it is equal to the low value. As a result, his total payoff is declining. The result has interesting consequences for entry games. If the seller could choose to enter one of two markets, he might indeed enter a market with a larger proportion of the low value buyers. Roughly speaking, although the buyers with low valuations offer less surplus to be shared, they also put up less of a fight when splitting the said surplus.

In the richer version of the model where the seller can also choose to exert effort, the results are qualitatively similar. There is, however, an additional parameter region where novel behavior arises: the seller randomizes over pairs of efforts and prices–no effort coupled with the low price and effort coupled with the high price. Such mixing is a consequence of the fact that the seller learns less when he exerts effort than when he does not. Intuitively, when the seller does not exert effort, having the right to make an offer strongly suggests that the buyer has not exerted effort himself and is therefore more likely to have a low valuation.

We explore the welfare effects of bargaining and, in particular, discuss the role of the buyer's and seller's effort costs. In equilibrium each party's probability of exerting effort decreases in their own effort cost, as one might expect. We then show that the cost of effort affects welfare through three channels: first, it directly affects the size of the loss incurred in equilibrium through the exertion of effort; second, it changes the relative probability of the informed and uninformed party making an offer; third, it has an effect on how much the seller learns in equilibrium. Through the latter two channels, the total and relative effort costs have important implications for the trading surplus that is realized in equilibrium.

Our results have bearing on the so called probabilistic bargaining models. In such models the buyer and the seller have exogenously fixed probabilities of making an offer. Moreover, these probabilities tend to be interpreted as bargaining powers; see for example Zingales (1995), Inderst (2001), Krasteva and Yildirim (2012) and Münster and Reisinger (2015). Our model implies that in some environments it would be more realistic to assume that bargaining power is increasing in the buyer's valuation. We show that taking the probability of making an offer exogenously fixed and increasing in the buyer's type provides the researcher with a simpler model yet captures some of the arresting features of our framework.

Related Literature: The foundation for models of non-cooperative surplus division under complete information was set by Stahl (1972) and Rubinstein (1982). Stahl (1972) proposed a finite time model in which two agents alternate in proposing a division of a commonly known surplus until an offer is accepted. The model is easily solved through backward induction, namely in the last period any proposal is accepted, given that the outside option is zero. Rubinstein (1982) solves for the infinite horizon version of the game, among other things showing that the solution is unique. A powerful implication of these models is that although the bargaining may proceed infinitely long, rational agents can foresee the outcome and achieve an agreement in the first round. Incomplete information bargaining models flourished after the work of Fudenberg and Tirole (1983) and Gul, Sonnenschein, and Wilson (1986).

Closest to our are models by Yildirim (2007), Yildirim (2010) and Board and Zwiebel (2012). Yildirim (2007) and Yildirim (2010) study a sequential bargaining model in which agents exert efforts to be the proposer of a split of a pie of a fixed and commonly known size. After the proposer is chosen and the proposal is made, the agents vote on whether to accept it. The focus is on how different voting rules affect the bargaining outcome. Board and Zwiebel (2012) study a similar model with two agents who repeatedly compete to make an offer to each other for a fixed pie until one of the offers is accepted. The competition is in the form of a first price auction, and the agents are endowed with a limited bidding capital. They characterize how the bargaining outcomes depend on the size and distribution of the bidding capital. Unlike our model, these papers analyse environments where the surplus is of a commonly known size.

2 The Baseline Model

A seller and a buyer want to trade an object. The seller's valuation is commonly known and normalised to zero, while the buyer's valuation v is either high or low, denoted by v_H and v_L respectively. We assume $v_H > v_L > 0$ and denote the probability that the buyer's valuation is v_H by μ. After learning his type, the buyer decides whether to exert effort; his effort choice is denoted by $e_b \in \{0, 1\}$. Costly effort affects the buyer's probability of making an offer to the seller. Let $k_b e_b$ denote the cost of effort and let ρ_1 (resp. ρ_0) denote the buyer's probability of making an offer to the seller when he does (resp. does not) exert effort. We assume $0 < \rho_0 < \rho_1 < 1$. The seller gets to make an offer with the remaining probability. After the offer is accepted/rejected the game ends.

In the baseline model we keep the seller's effort fixed. The seller potentially learns about the buyer's type through the possibility of making an offer and moreover the extent to which the seller learns may depend on his choice of effort. Shutting down the seller's effort component enables us to distill the direct effect of learning on the equilibrium outcome. Later we relax this assumption in order to see how the effects of the seller's effort interplay with learning.

A pure strategy for the buyer is a tuple (e_b, p_b) for each type, where $p_b \in \mathbb{R}$ denotes the price the buyer proposes when called upon to make an offer. The seller only chooses a price offer $p_s \in \mathbb{R}$.

Remark 1. An example to keep in mind is that of a Tullock contest in which agents' chances to make an offer are commensurate to their effort investments. More precisely, suppose that the buyer can choose between two actions $\{e^L, e^H\}$, where $e^H > e^L > 0$, and the seller's action is fixed to e^s. The probability that agent i gets to make an offer is $\frac{e_i}{e_i+e_j}$, where e_j is the opponents effort. The buyer incurs cost $\tilde{k}_b(e_b - e^L)$ from exerting effort e_b, while the seller incurs no cost. Our model corresponds to setting $\rho_1 \triangleq \frac{e^H}{e^H+e^s}$ and $\rho_0 \triangleq \frac{e^L}{e^L+e^s}$.

We call a profile of strategies and a belief function an equilibrium if the strategies are sequentially rational given the beliefs and the beliefs are updated using Bayes rule whenever possible. In addition, we assume that each agent after their own deviation presumes that the other agent has followed his equilibrium strategy. Two types of off-equilibrium paths are possible. First, the buyer could be the one making an offer after an own deviation in effort. In such a case, he computes the posterior using his deviation strategy and equilibrium prescribed strategies for the opponent. Second, an out of equilibrium price could be announced. The beliefs at such a node, however, are irrelevant due to the private values structure of the model.

We maintain two interpretations for the effort decision in this environment. First, agents often exert effort during or prior to the bargaining. They read a book that (presumably) increases their bargaining skills, i.e. Cialdini and Garde (1987), Trump and Schwartz (2009), attend a course on bargaining, or search the internet for advice. Data shows that people spend on average more than eight hours researching how to bargain before they buy a car in the US.[2] Second, when an asset has significant value agents find specialists to negotiate on their behalf. These can take the form of lawyers and/or negotiation consultants. Marital dissolution, for example, is a multibillion dollar business.[3]

The outlined model, as always, is a simplified representation of reality. We do not claim that bargaining proceeds precisely as specified above. However, the model does capture two important features that we would like to explore: 1.) A larger investment in bargaining enables the agent to secure a higher share of the surplus. This is captured by the probability with which an agent gets to make an offer. In fact, in the literature the probability that an agent makes a take-or-leave-it offer is commonly used as a proxy for bargaining power (see for example Zingales (1995) and Inderst (2001)). 2.) By observing the opponent's bargains stance, an agent may learn something about the opponent's type. The seller faced with an opportunity to make an offer can infer something about the buyer's valuation due to different types of buyer making (possibly) different effort decisions.

2.1 Optimal Bargaining Effort

It is useful to start with some preliminary observations. If the buyer gets to make an offer, he optimally offers price $p_b = 0$. The buyer's strategy can therefore be described by the probability with which he exerts effort. Let this probability be denoted by β_i when the buyer's type is $i = L, H$. On the other hand, if the seller gets to make the offer he either offers the pooling price v_L or the separating price v_H. The seller's strategies can thus be reduced to the probability with which he offers price v_H; denoted by $\sigma \in [0, 1]$.

[2]See http://agameautotrader.com/agame/pdf/2016-car-buyer-journey.pdf

[3] See http://www.huffingtonpost.com/susan-pease-gadoua/divorce-is-big-business-a_b_792271.html

The buyer type v_i's expected payoff from exerting effort $e_b \in \{0, 1\}$ when the seller uses the strategy σ is

$$\rho_{e_b} v_i + (1 - \rho_{e_b}) \mathrm{E}_{\sigma_H}[1_{[p_s \leq v_i]}(v_i - p_s)] - k_b e_b.$$

After making an effort choice e_b the buyer gets to make an offer with probability ρ_{e_b} and obtains a payoff equal to v_i. With the complementary probability the seller makes an offer, which the buyer accepts as long as his valuation is not below the price.

Of a particular interest will be the buyer's benefit from exerting effort, which can be written as

$$\Delta u_b(v_i) = \underbrace{\Delta\rho}_{\triangleq \rho_1 - \rho_0} \left(v_i - \mathrm{E}_{\sigma_H}[1_{[p_s \leq v_i]}(v_i - p)]\right) - k_b. \tag{1}$$

The term $\Delta\rho$ measures the increase in the probability of the buyer making an offer due to exerting effort, while the second factor is the increase in the buyer's payoff when he makes the offer instead of the seller. In this case the buyer obtains a payoff equal to his valuation v_i, whereas if the seller makes the offer the buyer's payoff is $v_i - p$ if he accepts and 0 otherwise.

An important consequence of the above derived benefits is that the high type buyer's incentives to exert effort are at least as high as those of the low type.

Lemma 1. *For any strategy σ, we have $\Delta u_b(v_H) \geq \Delta u_b(v_L)$, where the inequality holds with equality if and only if $p_s = v_L$.*

The above result follows after a closer inspection of the buyer's benefits from exerting effort. For the low type buyer expression (1) simplifies to

$$\Delta u_b(v_L) = \Delta\rho v_L - k_b. \tag{2}$$

The increase in the buyer's payoff from making an offer is v_L: if the seller gets to make an offer, the low type buyer's payoff is zero; he either rejects the seller's offer or pays a price equal to his valuation. The low type buyer, therefore, benefits from exerting effort if and only if $k_b \leq \Delta\rho v_L$. As a result, the low type buyer's optimal strategy is independent of the price the seller charges, implying that he has a dominant strategy.[4] If his cost k_b is sufficiently small he exerts effort, otherwise he does not.

[4] Strictly speaking, low type buyer's strategy is conditionally dominant, since we are already restricting the seller to the prices in $\{v_L, v_H\}$

The high type buyer's benefits from switching to exerting effort is

$$\Delta u_b(v_H) = \Delta\rho \mathrm{E}_{\sigma_H}[p_s] - k_b. \tag{3}$$

The high type buyer's benefit from making an offer is the price he does not have to pay. Thus, the higher the expected price charged by the seller, the larger are the benefits from exerting effort. Lemma 1 now follows from the observation that $\mathrm{E}_{\sigma_H}[p_s] \geq v_L$. Further insight can be gained in the case of the seller charging the price v_L. When the seller charges the low price, trade occurs with certainty and both types of buyer receive their respective values. Their benefit from making an offer is, therefore, the price they do not have to pay; in this case v_L. Consequently, when the seller charges the low price, both types of buyer have the same incentive to exert effort.

The high type's behavior is more intricate. Since $\mathrm{E}_{\sigma_H}[p_s]$ belongs to the interval $[v_L, v_H]$, we have

$$\Delta\rho v_L - k_b \leq \Delta u_b(v_H) \leq \Delta\rho v_H - k_b.$$

The high type buyer has two dominance regions: he always exerts effort when the cost of effort is sufficiently small; in particular $k_b < \Delta\rho v_L$. In that case, the cost is so small that the buyer finds it worth to expend effort even if he expects the seller to charge price v_L. On the other hand, when the cost of effort is so high that even if the seller charges the high price exerting effort is not optimal, $\Delta\rho v_H < k_b$, not exerting effort is the dominant action for the high type buyer. In the intermediate range of effort, the high type's behavior depends on the seller's expected price; see Figure 1. The above analysis can be summarised as follows.

Lemma 2. *Exerting effort is conditionally dominant for both types of buyer if $k_b < \Delta\rho v_L$, whereas not exerting effort is conditionally dominant for the low type buyer if $k_b > \Delta\rho v_L$ and for the high type buyer if $k_b > \Delta\rho v_H$.*

Both types optimally choose not to exert effort when the cost of effort is sufficiently high and to exert effort when the cost is sufficiently low. When the marginal cost of effort is in the intermediate region,

$$\Delta\rho v_L \leq k_b \leq \Delta\rho v_H,$$

the low type buyer optimally refrains from exerting effort, whereas the high type's optimal effort decision depends on the expected price the seller charges when making the offer.

Lemma 1 implies that the buyer's equilibrium choice of effort is monotonic: the low type never exerts effort with a higher probability than the high type. This has consequences on the seller's behavior. If the seller believes that the buyer type v_i exerts effort

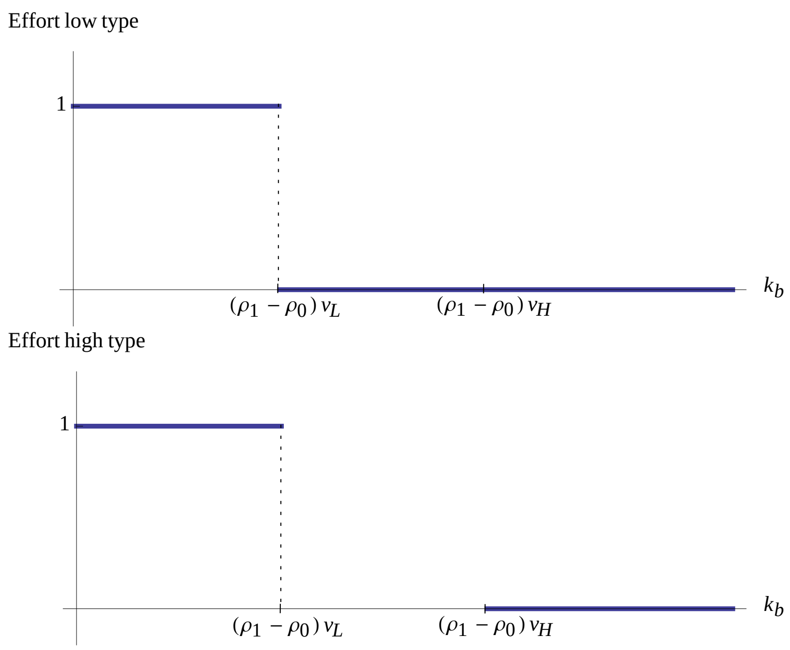

Figure 1: Conditional dominance regions for low and high type buyer

with probability β_i $(i = L, H)$ and the seller gets to make the offer himself, his posterior that the buyer is of high type is

$$\hat{\mu} \triangleq \frac{\mu\left[\beta_H(1-\rho_1) + (1-\beta_H)(1-\rho_0)\right]}{\mu\left[\beta_H(1-\rho_1) + (1-\beta_H)(1-\rho_0)\right] + (1-\mu)\left[\beta_L(1-\rho_1) + (1-\beta_L)(1-\rho_0)\right]},$$

where $\beta_H(1-\rho_1)+(1-\beta_H)(1-\rho_0)$ is the probability that the seller wins against the high type, and $\beta_L(1-\rho_1)+(1-\beta_L)(1-\rho_0)$ the probability that he wins against the low type buyer. Since Lemma 1 implies that the high type buyer is more likely to exert effort, that is $\beta_H \geq \beta_L$, he is also the one who is more likely to make an offer. In other words, the seller is more likely make an offer when he is facing the low type. His posterior, denoted $\hat{\mu}$, is, therefore, at most as high as his prior.

Lemma 3. *In equilibrium $\hat{\mu} \leq \mu$.*

While being able to make an offer is good news for the seller, he is under no illusion and understands that he is more likely to win against the low type buyer. Winning, therefore, depresses his belief. The reader should note that the above lemma does not contradict the fact that beliefs are a martingale. Indeed, when the buyer is making an offer, the seller revises his beliefs upwards. Seller's beliefs when the buyer is making an offer are, however, of no interest to us due to the private-value nature of our environment.

The above result echoes Coasian dynamics where the prices fall over time – a result closely related to the fact that the seller's belief is declining; see Fudenberg and Tirole

(1991). The reason for the declining beliefs, of course, is different. In the Coasian dynamics high types are more eager to accept a price, therefore the liklihood of encountering the high type in the market is falling over time. In our model, the seller is less likely to be able to make an offer against the high type.

Equilibria. In what follows we characterize equilibria. We split the analysis into cases depending on the prior. Before proceeding we should point out that the threshold v_L/v_H plays a prominent role in the analysis. Namely, when the seller offers the low price, he obtains the payoff v_L, whereas, if he offers the high price his expected payoff is $\hat{\mu}v_H$. The seller, thus, optimally offers the pooling price if and only if his posterior is low enough: $\hat{\mu} \leq \frac{v_L}{v_H}$.

We start with the case where the seller's prior is low: $\mu \leq v_L/v_H$.

Proposition 1. *If $\mu \leq \frac{v_L}{v_H}$, there is a generically unique equilibrium with the property that both types of buyer undertake the same effort choice and the seller offers price v_L.*

The seller's decision to offer price v_L is a consequence of Lemma 3: if his prior is below the threshold v_L/v_H, his posterior will be below the threshold too. Moreover, once the seller is charging the low price, Lemma 1 implies that both types of buyer have the same incentive to exert effort and, therefore, make the same effort choice. The seller, in turn, learns nothing from gaining the opportunity to make an offer. In particular, his posterior coincides with his prior.[5]

Next, we turn attention to the case when the seller's prior is high: $\mu > v_L/v_H$. Lemma 2 established that both types of buyer make the same effort choice when the cost of effort is extreme—either very high or very low. In those cases, the two types of buyer have the same dominant strategy and the seller cannot infer anything from the bargaining process. He, therefore, offers the high price.

Proposition 2. *Let $k_b \leq \Delta\rho v_L$ or $k_b \geq \Delta\rho v_H$ and $\mu > v_L/v_H$. In the (generically) unique equilibrium both types of the buyer make the same effort choice and the seller offers price v_H.*

To summarize, if the buyer's effort cost is sufficiently small or large and/or the seller's prior on the buyer's type is sufficiently low, there is no learning in equilibrium.

The seller learns. Hereafter we focus on the environment in which the high type buyer's optimal effort choice depends on the seller's behavior, in particular, on the case

[5] Genericity in the above proposition refers to the case where both types of buyer are indifferent between exerting effort and not exerting effort when the seller is expected to offer the low price.

when the buyer's marginal cost is intermediate:

$$\Delta\rho v_L < k_b < \Delta\rho v_H.$$

As a reminder, in this region the low type buyer prefers to exert low effort, while the high type buyer's optimal effort decision depends on the price he expects the seller to charge. The potential difference in the strategies of the two types of buyer gives rise to a possibility for the seller to learn about his competitor. To this end, we will maintain the assumption $\mu > \frac{v_L}{v_H}$, so that if the seller were to lear nothing, he would optimally propose the separating price v_H.

Before stating the result, it is useful to define additional notation. Let

$$m \equiv \frac{(1-\rho_0)\frac{v_L}{v_H}}{(1-\rho_0)\frac{v_L}{v_H} + (1-\rho_1)(1-\frac{v_L}{v_H})}, \tag{4}$$

be the prior belief such that if the high type buyer exerts effort and the low type does not, the seller's posterior is precisely v_L/v_H. As a consequence, for all priors above m, the seller's posterior will be above the threshold v_L/v_H, regardless of the buyer's behavior, and the seller optimally offers price v_H.

Proposition 3. *Let $k_b \in (\Delta\rho v_L, \Delta\rho v_H)$ and $\mu > \frac{v_L}{v_H}$. There is a unique equilibrium:*

- *when $\mu < m$, the high type buyer exerts effort with probability $\beta_H = \frac{(1-\rho_0)(\mu v_H - v_L)}{\Delta\rho\mu(v_H - v_L)}$ and the low type exerts no effort; the seller offers price v_H with probability $\sigma = \frac{k_b - \Delta\rho v_L}{\Delta\rho(v_H - v_L)}$;*
- *when $\mu \geq m$, the high type buyer exerts effort and the low type does not; the seller offers price v_H.*

The equilibrium exhibits the most interesting behavior when the prior is only slightly above the threshold $\frac{v_L}{v_H}$; more precisely, when the prior is between $\frac{v_L}{v_H}$ and m. If the seller was to learn nothing during bargaining, he would offer the high price v_H. In anticipation of the seller's behavior, the high type buyer would find it optimal to exert effort, while the low type buyer would not. The seller should then revise his beliefs downward and optimally offer the low price v_L. Foreseeing the seller offering the pooling price neither type of the buyer would have an incentive to exert effort. But then the seller would not learn, and, completing the circle, would want to offer the separating price.

To prevent such cycling of beliefs, the high type buyer must randomize between exerting and not exerting effort in such a way that the seller is indifferent between the two prices, that is, in such a way that the seller's posterior is precisely v_L/v_H; this pins down β_H. The probability of high type exerting effort is strictly increasing in the prior on the relevant region. Namely, the higher the prior the bigger must be the difference between the low and the high type's effort to bring the seller's posterior down to v_L/v_H. The seller, on the other hand, offers the high price with the probability that makes the high type buyer indifferent between the two efforts. The high type's indifference delivers the seller's probability of charging the high price, σ. Notice that the latter is independent of the prior; after all, it is determined by the high type's indifference.

For the priors above m the seller's posterior will be above the threshold v_L/v_H irrespective of the buyer's behavior. The seller, therefore, offers the high price, which sets a strong incentive for the high type buyer to make the offer himself. Consequently, he exerts effort with certainty. This covers the second bullet point in the above result.

2.2 Comparative Statics

Of particular interest is how the seller's payoff varies with his prior μ.

Corollary 1. *The seller's equilibrium expected payoff is*

$$u_s = \begin{cases} (1-\rho_0)v_L & \text{if } \mu < \frac{v_L}{v_H} \\ (1-\mu)(1-\rho_0)\frac{v_L v_H}{v_H - v_L} & \text{if } \frac{v_L}{v_H} \leq \mu \leq m \\ \mu(1-\rho_1)v_H & \text{if } m < \mu. \end{cases}$$

The striking property of the seller's payoff is that it is decreasing in the probability of the high type on the interval $[v_L/v_H, m]$. The implied non-monotonicity stands in stark contrast to models where the probability of each of the two agents making an offer is exogenously determined and fixed over types; more about this in Section 4.2.

A closer inspection of the structure of the equilibria reveals what lies behind the peculiar behavior of the seller's payoff. For low priors, $\mu < v_L/v_H$, the seller invariably offers the pooling price v_L and neither type of buyer exerts effort. His payoff is price v_L times the probability that the seller makes an offer conditionally on the buyer not exerting effort, $1-\rho_0$, and as such independent of the prior.

To see why the seller's payoff is decreasing when his prior is in the intermediate region, recall that in the relevant region of priors the seller is indifferent between the two price

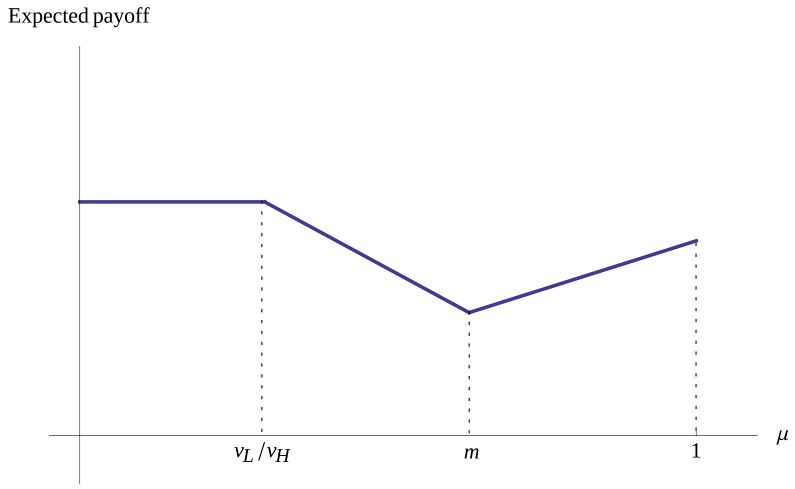

Figure 2: Seller's equilibrium expected payoff

offers. His payoff from offering the low price is

$$u_S(\mu) = [\mu(\beta_H(1-\rho_1) + (1-\beta_H)(1-\rho_0)) + (1-\mu)(1-\rho_0)]\, v_L.$$

The term in the brackets is the probability that the seller gets to make an offer, while v_L is the seller's payoff conditionally on making an offer. Notably, the seller's payoff conditionally on making the offer is independent of his posterior or the prior. Since he is randomizing over the two price and the price v_L is accepted with certainty, his conditional payoff is v_L. As to probability of the seller making an offer, when the prior rises, the high type buyer must exert effort with higher probability to distinguish himself from the low type, and thereby bring the seller's posterior down to the threshold belief. With other words, β_H is increasing in μ. The probability that the seller gets to make an offer is, therefore, falling for two reasons: 1) as μ increases, the high type is exerting effort with higher probability, and 2) the seller faces the high type more often. The seller's probability of making an offer, consequently, decreases in μ, while his payoff conditionally on making the offer is constant. Together this implies that an increase in the seller's prior leads to a decrease in his expected payoff.

In the third region, where $\mu > m$, the seller optimally offers the separating price v_H, which is accepted only by the high type buyer. Here the high type buyer exerts effort, implying that the seller's probability of making an offer is $1-\rho_1$. In contrast to the first two parameter region, the seller's payoff in the third parameter region is strictly increasing in the fraction of high type buyers. This follows from the fact that the probability with which the seller gets to make an offer against the high type, and consequently his payoff conditionally on facing the high type, is constant in μ, while his chances of meeting a high type and having price v_H accepted are strictly increasing.

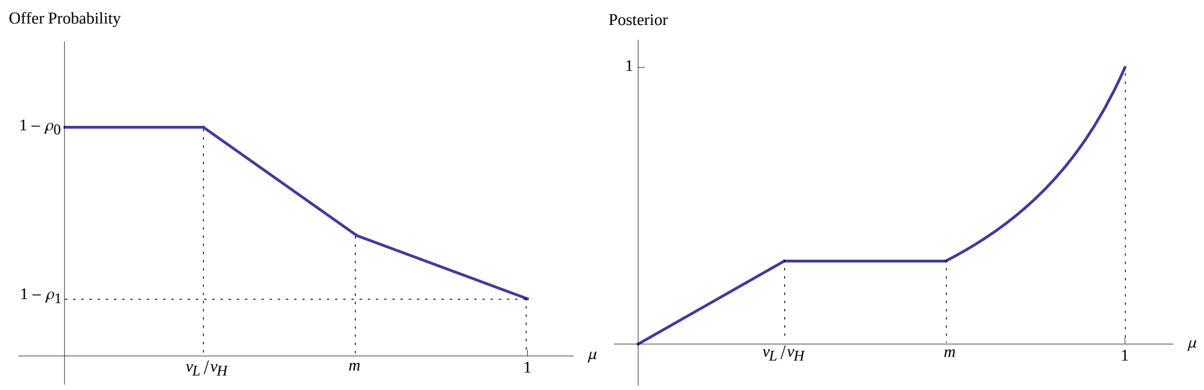

Figure 3: Seller's probability of making an offer and posterior belief in equilibrium

The seller's expected payoff is therefore non-monotonic in the prior: it is constant for the low priors, decreasing for the intermediate priors and finally increasing for the high priors. Whether the seller's expected payoff increases above the value it attains for low priors depends on the parameters of the problem. The seller is better off facing the low type buyer with certainty ($\mu = 0$) than the high type buyer with certainty ($\mu = 1$) if $(1 - \rho_0)v_L \geq (1 - \rho_1)v_H$. A high type buyer promises higher gains from trade but the probability that the seller can appropriate these gains are smaller than the respective probability when facing a low type buyer. A seller therefore obtains a higher expected payoff when faced with a low type buyer if the difference in the buyer's valuation is small relative to the difference in probabilities with which the seller gets to make the offer.

Finally, notice that the case for the seller's payoff being decreasing could have been made even in the environment with perfect information. This is the case when the seller's payoff at $\mu = 1$ is smaller than at $\mu = 0$. However, our result is stronger than that, even when the seller's payoff from facing the high type with certainty is larger than when he is facing the low type, there is an intermediate region of priors under which the seller is worse off than facing the low type.

The non-monotonicity in the seller's expected payoff stands in stark contrast to models where the probability of each of the two agents making an offer is exogenously determined and fixed over types. This has potentially interesting implications for applied work. Suppose that the seller has to decide on entering one of two markets: the first populated almost exclusively by low value buyers, the second consisting of a fair share of high value custumers. The above non-monotonicity result shows that the seller might indeed prefer to enter the market with low value custumers.

3 Two Sided Effort

We now augment the analysis with the possibility of the seller exerting effort $e_s \in \{0,1\}$ at a cost $k_s e_s$. The seller's effort will in addition to affecting the probability with which the seller gets to make an offer play a prominent role in the seller's learning about the buyer's valuation. We show that this interaction gives rise to novel and interesting equilibrium behavior

We assume that the probability of making an offer as a function of the buyer's and seller's effort choice takes the following simple form:

$$\begin{cases} \rho & \text{if } e_i > e_{-i}, \\ \frac{1}{2} & \text{if } e_i = e_{-i}, \\ 1-\rho & \text{if } e_i < e_{-i}, \end{cases}$$

where $\rho \in (\frac{1}{2}, 1)$. The effect of effort on the probability of making an offer is symmetric for the buyer and the seller, while differences in bargaining skills are captured by the relative costs of effort. As in the previous section, a canonical example of our environment is a Tullock contest, though here both agents choose between two efforts.

3.1 Equilibrium Analysis

The seller optimally chooses a price p_s from the set $\{v_L, v_H\}$. His strategy can, therefore, be described by a probability distribution σ over pairs (e_s, p_s) with $e_s \in \{0,1\}$ and $p_s \in \{v_L, v_H\}$. The buyer optimally offers a price equal to zero, thus his decision problem boils down to the choice of effort.

Letting E_σ denote the expectation operator with respect to σ, we can write the high type buyer's benefit from exerting effort as

$$\begin{aligned} \Delta u_b(v_H) &\triangleq \Pr_\sigma[e_s = 1]\mathrm{E}_\sigma\left[\left(\rho - \frac{1}{2}\right)p_s - k_b | e_s = 1\right] + \Pr_\sigma[e_s = 0]\mathrm{E}_\sigma\left[\left(\rho - \frac{1}{2}\right)p_s - k_b \mid e_s = 0\right] \\ &= \left(\rho - \frac{1}{2}\right)\mathrm{E}_\sigma[p_s] - k_b. \end{aligned}$$

The term $\rho - \frac{1}{2}$ is the increase in probability that the buyer gets to make an offer when switching to exerting effort. This increase does not depend on whether or not the seller exerts effort due to $\frac{1}{2} - (1-\rho) = \rho - \frac{1}{2}$. For the low type buyer we have:

$$\Delta u_b(v_L) \triangleq \left(\rho - \frac{1}{2}\right)v_L - k_b.$$

Since the buyer's benefit of exerting effort does not depend on the seller's effort choice, some of the lemmata that we developed in Section 2 carry over to this environment. In particular, the high type buyer has at least as high an incentive to exert effort as the low type. This implies that (generically) the high type buyer exerts effort with at least as high probability as the low type. Another implication, corresponding to Lemma 3, is that the seller's posterior is never above his prior.

The seller's benefit of exerting effort when planning to offer a price p_s:

$$\Delta u_s = (\rho - \frac{1}{2})[\mu p_s + (1-\mu)1_{[p=v_L]}p_s] - k_s. \quad (5)$$

The term $\rho - \frac{1}{2}$ is the increase in probability that the seller gets to make the offer when he shifts to exerting effort. This increase does not depend on whether the buyer exerts effort or not. The seller's optimal effort choice is, therefore, determined by the price he intends to propose when making the offer. At the same time, it does not directly depend on the buyer's behavior. The seller's equilibrium price, and through that his effort, however, will depend on the buyer's conduct.

The seller's choice of effort affects how much he learns about the buyer and, as a result, his optimal pricing strategy. Indeed, the effect of the seller's effort on his posterior is an important aspect of the analysis to follow.

Lemma 4. *When the high type buyer exerts effort with a higher probability than the low type buyer, the seller's posterior $\hat{\mu}$ is smaller when $e_s = 0$ than when $e_s = 1$.*

Exerting effort diminishes the extent to which the seller updates his belief, i.e., learns. All else equal exerting effort favors the separating price v_H. More precisely, if after not exerting effort the seller optimally offers v_H, he prefers price v_H also after exerting effort. This has important implications for the equilibrium characterization that follows. We will restrict attention to the case where the two types of buyer do not have the same dominant effort choice, so that learning is indeed possible. This requires $k_b \in ((\rho-1/2)v_L, (\rho-1/2)v_H)$. There are two parameters—k_s and μ—that govern the structure of equilibria. To mirror the analysis where the seller did not have the choice of effort we fix the cost of effort, k_s, and describe how the equilibrium varies with μ.

First, if $k_s > (\rho - \frac{1}{2})v_H$, the benefit from exerting effort is negative regardless of the seller's prior and his choice of prices. Even if the seller was to face the high type with certainty expending effort would be too costly. The seller, therefore, optimally chooses not to exert effort for any μ. The equilibrium analysis of Section 2 where the seller did not have an effort choice applies after setting $\rho_1 = \rho$ and $\rho_0 = \frac{1}{2}$.

On the other hand, if $k_s < (\rho - \frac{1}{2})v_L$, the seller exerts effort for all μ. If, $\mu < \frac{v_L}{v_H}$, the seller optimally offers the low price as in Section 2 and consequently finds it propitious to exert effort. If $\mu > v_L/v_H$, the benefit of exerting effort is at least $(\rho - \frac{1}{2})v_L - k_s$ for any randomization over prices. Therefore, for high priors too, the seller optimally exerts effort. The equilibrium analysis of Section 2 can be applied with the parameters $\rho_1 = \frac{1}{2}$ and $\rho_0 = 1 - \rho$. These cases are summarised in the following proposition.

Proposition 4. *Equilibrium is generically unique and satisfies the following properties:*

- *If $k_s \leq (\rho - 1/2)v_L$, the seller exerts effort and the seller's pricing strategy as well as the buyer's effort strategy are described by Propositions 1- 2 with $\rho_1 = 1/2$ and $\rho_0 = 1 - \rho$.*
- *If $k_s \geq (\rho - 1/2)v_H$, the seller exerts no effort and the seller's pricing strategy as well as the buyer's effort strategy are described by Propositions 1- 2 with $\rho_1 = \rho$ and $\rho_0 = 1/2$.*

More elaborate, and interesting, is the analysis of the case where the seller does have an incentive to change his effort choice as the prior changes. Before we proceed, a closer inspection of the seller's benefit of exerting effort, $\Delta u_s(p_s)$, is needed. First, if the seller plans to charge the price v_H, his incentive to exert effort increases with his prior μ. Second, for $\mu > \frac{v_L}{v_H}$ we have $\Delta u_s(v_H) > \Delta u_s(v_L)$, so his incentive to expend effort increases with the price he intends to charge. Thus, if the seller is not willing to exert effort at a prior above the threshold $\frac{v_L}{v_H}$ when planning to offer the high price v_H, he will not expend effort for any lower price. Let $m^H(k_s)$ be the prior at which the seller is indifferent between exerting effort and not when he charges the high price: $(\rho - \frac{1}{2})m^H(k_s)v_H - k_s = 0$, or differently

$$m^H(k_s) = \frac{k_s}{(\rho - \frac{1}{2})v_H}. \tag{6}$$

The above reasoning implies that irrespective of the buyer's behavior the seller prefers not to exert effort for priors below $m^H(k_s)$. Therefore, for $\mu \leq m^H(k_s)$ the analysis from Section 2 applies with $\rho_1 = \rho$ and $\rho_0 = \frac{1}{2}$.

Recall that, in the environment without the seller's effort choice, the seller's pricing decision was to offer the low price for low priors, randomize over the two prices at intermediate priors and offer the high price for high priors. Given $\rho_1 = \rho$ and $\rho_0 = \frac{1}{2}$, the threshold between the last two regions, as determined in (4), is given by

$$m_0 = \frac{\frac{1}{2}\frac{v_L}{v_H}}{\frac{1}{2}\frac{v_L}{v_H} + (1-\rho)(1 - \frac{v_L}{v_H})}.$$

The parameter m_0 is the seller's prior such that, conditionally on not exerting effort, the seller's posterior when the two types of buyers completely separate in their effort choice is precisely $\frac{v_L}{v_H}$. We can now break the analysis into the cases where the threshold $m^H(k_s)$ falls below or above m_0.

Suppose $m^H(k_s) > m_0$. As we just argued, for priors below $m^H(k_s)$ the analysis from Section 2 applies with the seller not exerting effort. Given $m^H(k_s) > m_0$, the threshold $m^H(k_s)$ then falls into the parameter region where, in the equilibrium characterized in Proposition 3, the high type buyer exerts effort with probability one and the seller offers the high price. When we allow the seller to make an effort decision, at the prior $m^H(k_s)$ the seller optimally switches to expending effort. His posterior jumps upwards (see Lemma 3), implying that the seller indeed prefers to offer the high price. Hence, for $\mu \geq m^H(k_s)$ the seller exerts effort and offers price v_H, while for the remaining values of μ the equilibrium is as characterized in Proposition 3. The following proposition summarizes this analysis.

Proposition 5. *Let $k_b, k_s \in ((\rho - 1/2)v_L, (\rho - 1/2)v_H)$ and $m^H(k_s) > m_0$. There is a generically unique equilibrium with the following properties:*

- *if $\mu \leq \frac{v_L}{v_H}$, nobody exerts effort and the seller offers the low price;*
- *if $\frac{v_L}{v_H} < \mu \leq m_0$, the high type buyer randomizes over the two effort choices, the seller does not exert effort and randomizes over the two prices;*
- *if $m_0 < \mu < m^H(k_s)$, the high type buyer exerts effort, while the seller does not; the seller offers price v_H;*
- *it $m^H(k_s) \leq \mu$, the seller and the high type exert effort, and moreover, the seller offers price v_H.*

The low type buyer never exerts effort.

It remains to investigate the case where $m^H(k_s)$ is smaller than m_0. In this case, the threshold $m^H(k_s)$ falls into the intermediate region of priors with respect to the equilibrium characterized in Proposition 3.[6] For priors just below $m^H(k_s)$ the seller thus exerts no effort and randomizes across prices. At the threshold $m^H(k_s)$ the seller prefers switching to positive effort if contemplating to offer the high price but not if planning to offer the low price. If the seller was to exert effort and offer the high price with probability one, the high type buyer would optimally exert effort as well. This is an equilibrium if

[6] Since we are interested in the case $k_s > (\rho - 1/2)v_L$, the condition $m^H(k_s) > v_L/v_H$ is always satisfied.

the seller has no incentives to switch to not exerting effort and offering the low price. That is

$$\mu\frac{1}{2}v_H - k_s \geq \left(\mu(1-\rho) + (1-\mu)\frac{1}{2}\right) v_L.$$

Letting $\tilde{m}(k)$ denote the prior at which the above condition is satisfied with equality, the described equilibrium exists for all $\mu \geq \tilde{m}(k)$. It turns out that, under the assumed parameter conditions, the threshold $\tilde{m}(k_s)$ is strictly larger than $m^H(k_s)$. This implies that there is an additional parameter region, where there is no equilibrium in which the seller has a fixed effort. This brings about a rather interesting new type of behavior.

Proposition 6. *Assume $k_b, k_s \in ((\rho - 1/2)v_L, (\rho - 1/2)v_H)$ as well as $m^H(k_s) < m_0$. There is a generically unique equilibrium with the following properties:*

- *if $\mu \leq \frac{v_L}{v_H}$, nobody exerts effort and the seller offers the low price;*
- *if $\frac{v_L}{v_H} < \mu \leq m^H(k_s)$, the high type buyer randomizes over the two effort choices, the seller does not exert effort and randomizes over the two prices;*
- *if $m^H(k_s) < \mu < \tilde{m}(k_s)$, the high type buyer randomizes over the two effort choices, while the seller randomizes over the two pairs $(0, v_L)$ and $(1, v_H)$;*
- *for $\tilde{m}(k_s) \leq \mu$, the high type buyer and the seller exert effort and the seller offers the high price;*

where $\tilde{m}(k_s) = \frac{k_s + \frac{1}{2}v_L}{\frac{1}{2}v_H + (\rho - \frac{1}{2})v_L}$. In all cases the low type buyer does not exert effort.

Proposition 6 shows that when $m^H(k_s) < \mu < \tilde{m}(k_s)$, there is an equilibrium where the seller randomizes across the effort-price pairs $(0, v_L)$ and $(1, v_H)$. Such randomization can be sustained due to Lemma 4: when the seller exerts effort, his posterior is higher than when he does not. Indeed, in the relevant region of priors no effort is optimally followed by the low price, while effort is optimally followed by the high price. The seller's randomization preserves the expected price, leaving the high type buyer indifferent between the two effort choices. The high type buyer, in response, adjusts his randomization to make the seller indifferent. In contrast to the case when μ is just below $m^H(k_s)$, the probability with which the buyer exerts effort does not make the seller indifferent between prices for a given effort but instead between the two strategies $(0, v_L)$ and $(1, v_H)$. The threshold $\tilde{m}(k_s)$ is the highest prior at which the seller can indeed be made indifferent between the two effort-price pairs.

3.2 Comparative Statics

First we explore how the seller's payoff varies as a function of his prior μ. We focus on the case described in Proposition 6.

Corollary 2. *Assume $k_b, k_s \in ((\rho - 1/2)v_L, (\rho - 1/2)v_H)$ and $m^H(k_s) < m_0$. The seller's equilibrium expected payoff is*

$$u_s = \begin{cases} \frac{1}{2}v_L & \text{if} \quad \mu \leq \frac{v_L}{v_H} \\ \frac{1}{2}(1-\mu)\frac{v_H v_L}{v_H - v_L} & \text{if} \quad \frac{v_L}{v_H} \leq \mu \leq m^H(k_s) \\ \frac{(\frac{1}{2} - \mu\rho)v_H + k_s}{v_H - v_L} v_L & \text{if} \quad m^H(k_s) \leq \mu \leq \tilde{m}(k_s) \\ \frac{1}{2}\mu v_H - k_s & \text{if} \quad \tilde{m}(k_s) \leq \mu. \end{cases}$$

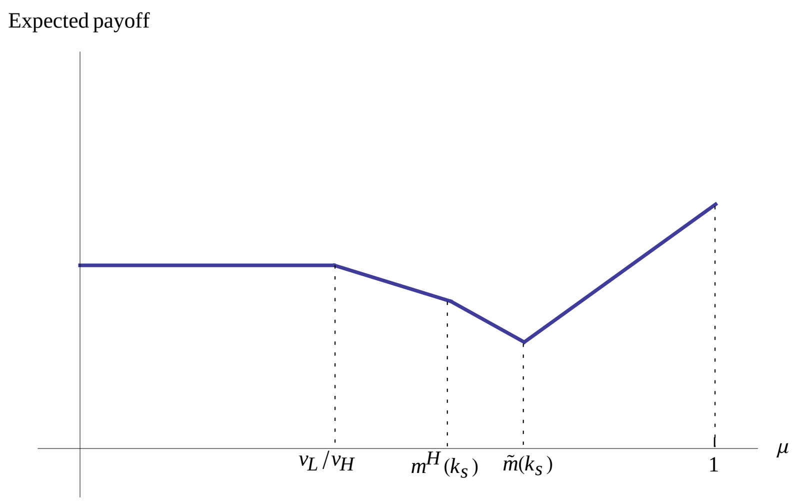

Figure 4: Seller's equilibrium expected payoff: two-sided effort

For priors below v_L/v_H nobody exerts effort, so the seller gets to make the offer v_L half of the time. Right above the threshold, the players play a mixed strategy equilibrium where the seller exerts no effort but, due to the high type buyer's randomization, he is indifferent between the two prices. Thus, conditionally on making an offer, the seller's payoff is constant, while his probability of making an offer strictly decreases in μ. In consequence, the seller's expected payoff is decreasing.

At the threshold $m^H(k_s)$, the seller's benefit of exerting effort when offering the pooling price v_H becomes positive. In order to keep him indifferent between this option and exerting no effort and offering v_L, the high type buyer's probability of exerting effort must increases in μ faster. This exhibits downward pressure on the seller's posterior and thereby reduces the value of offering the separating price v_H. Notably, the seller's posterior after exerting no effort now decreases in his prior: the high type buyer's probability of exerting effort increases sufficiently fast so that the learning effect overpowers the direct effect of the increase in μ, resulting in a lower posterior; the seller's posterior is depicted in Figure 5. As in the previous region of μ, the seller's expected payoff is decreasing in his prior, since his probability of making an offer decreases in μ. Finally, when μ reaches the threshold $\tilde{m}(k_s)$, the seller's prior is sufficiently high so that he and the high type buyer exert effort with probability one. Conditionally on the buyer being the high type, the seller wins half of the time and offers price v_H.

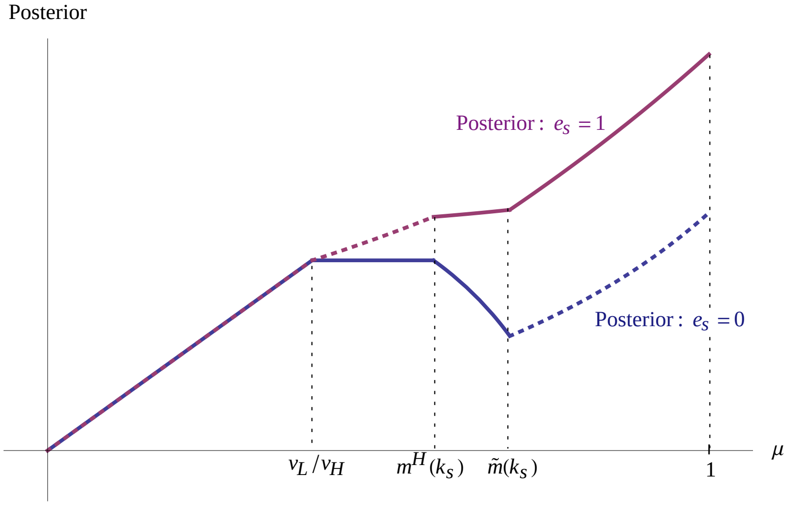

Figure 5: Seller's equilibrium posterior after exerting and not exerting effort

4 Discussion

4.1 Effort Costs and Welfare

In what follows we will discuss how total surplus depends on the effort cost parameters k_b and k_s. The two parameters affect the probability with which either party exerts effort in equilibrium and through that welfare. First, the seller's and buyer's equilibrium efforts are non-increasing in their own costs, k_s and k_b, as can be verified from Propositions 1-3. However, since the probability with which agents exert effort in equilibrium varies with the cost parameters in a non-continuous way, effects on welfare are rather erratic. Instead of giving a full account, we focus on three important channels through which effort costs influence welfare.

The first channel is the direct effect of an increase in k_s or k_b on the total cost of effort that is incurred in equilibrium. As long as an increase in the cost of effort does not lead to a change of the equilibrium effort choice of the respective agent, such increase clearly lowers total surplus. A higher effort cost might, however, deter an agent from exerting effort, thus lowering the incurred cost of effort and potentially increasing surplus.

The second channel is due to one-sidedness of the private information in our model. If the buyer was always the one to make an offer, trade would take place with probability one and the gains from trade would be maximized. In contrast, when the seller makes an offer, he might find it optimal to offer the separating price v_H, thereby forgoing the possibility of trade with the low type buyer. By implication, the realized gains from trade cannot decrease when the probability that the buyer makes the offer increases; all else equal.

An increase in the buyer's probability of making an offer can come either from the buyer exerting more effort or the seller exerting less. In the latter case, Lemma 4 shows that the seller's posterior is lower when he exerts no effort. This makes him more likely to charge the pooling price v_L, which, holding the buyer's effort fixed, increases the trading surplus. We therefore have a third 'learning channel': a higher cost of effort for the seller reduces his incentives to exert effort and can thereby increase the degree to which he learns about the buyer's type conditionally on making an offer, which in turn affects his optimal pricing strategy.

With this in mind, consider the effect of the buyer's effort cost k_b on total surplus. We will focus on the cases where k_b is sufficiently small or large so that both types of buyer have the same dominant effort choice. Given that the seller does not learn in equilibrium, we can concentrate on the first two channels. When the buyer's effort cost is sufficiently high, $k_b > (\rho - 1/2)v_H$, the buyer optimally exerts no effort and thus incurs no cost. On the other hand, when the buyer's effort cost is small, specifically when $k_b < (\rho - 1/2)v_L$, the buyer chooses to exert effort, incurring a cost that is increasing in k_b and vanishes as k_b tends to zero. To assess total surplus, the effect of k_b on realized effort costs has to be compared to the effect on trading surplus through the change in the probability with which the buyer makes an offer. Given that this probability is increasing in the buyer's effort and therefore decreasing in k_b, the effect on trading surplus is non-positive. Recalling that the effect on realised effort costs vanishes as k_b tends to zero, we can conclude that, in situations where the seller offers v_H with a positive probability,[7] welfare is strictly higher when k_b is very small than when it is large.

As for the seller's marginal cost of effort, for values of k_s sufficiently low or high, $k_s < (\rho - 1/2)v_L$ and $k_s > (\rho - 1/2)$ repsectively, the effect on the cost incurred in

[7] That is, $\mu > v_L/v_H$.

equilibrium is very similar to the one discussed for the buyer. For low values of k_s the cost is increasing in k_s, for high values it is zero. Who makes an offer, however, has now the opposite implication for welfare: a high value of k_s implies that the seller expends no effort, making it more likely for the buyer to make an offer and for trade to take place. In addition, high values of k_s enhance welfare through the learning effect. By Lemma Lemma 4, for a given strategy of the buyer, the seller revises his belief downward further when he does not exert effort. As a consequence, the conditional probability that the seller offers the low price in equilibrium is no smaller when k_s is large than when it is small.[8] In contrast to the case of k_b, a sufficiently large value of k_s is better for welfare than a value of k_s close to zero. Namely, while the difference in effort costs incurred in equilibrium vanishes as k_s tends to zero, a high value of k_s makes it more likely that the buyer makes an offer and less likely that trade is forgone in case the seller makes an offer.

Finally, we find that an increase in k_s can lead to an increase in total surplus, even when the seller's equilibrium effort choice does not change. What is even more perplexing, not only welfare but also the seller's payoff can increase with k_s. This effect arises in the parameter region where the seller randomizes across the pairs of strategies $(0, v_L)$ and $(1, v_H)$; see Proposition 6. In this region, the buyer's payoff does not depend on the seller's cost parameter since the seller's randomization between $(0, v_L)$ and $(1, v_H)$ is independent of k_s.[9] The seller's expected payoff is

$$u_s = \frac{(\frac{1}{2} - \mu\rho)v_H + k_s}{v_H - v_L} v_L,$$

where the term on the right-hand side is increasing in k_s; see Corollary 2. As k_s increases, exerting effort and offering the high price becomes less attractive for the seller. In consequence, the high type buyer can exert effort with a smaller probability, without violating the seller's indifference condition. This implies that the payoff associated to exerting no effort and offering the low price must be higher, as it does not depend on k_s and the seller gets to make an offer with a strictly higher probability. In equilibrium this payoff must be equal to the one the seller obtains when exerting effort and offering the high price. The positive effect on the probability of making an offer outweights the higher effort cost the seller incurs. Hence, the seller's expected payoff (and in consequence total surplus) strictly increases in k_s.

[8]Formally, the threshold for the mixed strategy equilibrium when the seller exerts no effort, m_0, is strictly greater than the respective threshold when the seller exerts efforts, defined by

$$m_1 = \frac{\rho \frac{v_L}{v_H}}{\rho \frac{v_L}{v_H} + \frac{1}{2}(1 - \frac{v_L}{v_H})}.$$

[9]The seller randomizes to keep the buyer indifferent, who clearly does not care about k_s.

One implication of the above result is that there are situations in which the seller would like to commit to having a higher cost of effort, e.g. by only having access to lawyers that are more expensive without being more efficient.

4.2 Probabilistic Bargaining Model

The probabilistic bargaining model is a model in which the buyer and the seller get to make an offer with a fixed probability. Unlike in the above proposed model, the probability of making an offer is exogenously given and independent of the agents' valuations. Such models are commonly used in economics and finance; for examples see Inderst (2001) and Zingales (1995), though probabilistic bargaining goes back at least to Rubinstein and Wolinsky (1985).

To simplify the comparison with our model, we first present a version of a probabilistic model of bargaining. Suppose that the buyer's and the seller's valuations are as above. When they meet, the seller gets to make a take-it-or-leave-it price offer with probability ρ and the seller with the remaining probability. If the offer is rejected the game ends.

Notice that the seller learns nothing about the buyer when making an offer. He, therefore, makes the offer on the basis of his prior distribution. The probability of being able to make an offer, ρ, captures the buyer's bargaining power. The only strategy of interest is the seller's pricing strategy. Since the seller's posterior is equal to his prior, he offers the high price if his prior is above v_L/v_H and the low price otherwise. This model corresponds to the model presented in previous sections only in the case where not exerting effort is the dominant strategy for both players. The cases where both players have dominant strategies, but one or both players exert effort, are quite similar too with the difference that in our model some surplus is burnt due to the cost of effort. The interesting properties, like learning and non-monotonicity of the seller's payoff, however, cannot be replicated with a simple probabilistic model.

Our model shows that higher value buyers have a higher propensity to exert effort and through that higher bargaining power. To capture this feature we propose a richer version of the probabilistic bargaining model in which the high type buyer gets to make an offer with probability ρ_H and the low type buyer with probability ρ_L, where $\rho_H > \rho_L$.

Proposition 7. *The seller optimally offers the high price if*

$$\mu \geq \underbrace{\frac{(1-\rho_L)\frac{v_L}{v_H}}{(1-\rho_H)(1-\frac{v_L}{v_H})+(1-\rho_L)\frac{v_L}{v_H}}}_{=\mu^*}$$

and the low price otherwise.

It is easy to see that μ^* (the right hand side of the above inequality) is larger than v_L/v_H. Being given the opportunity to make an offer is a negative signal for the seller; he has an easier time making an offer against the low type. Winning the chance to make an offer, therefore, lowers his belief. Consequently, he is willing to offer the high price only when his prior sufficiently favors the high type.

To highlight the comparison between the probabilistic bargaining model with varying probabilities of making offers and our model with efforts, we examine how the seller's payoff varies with the prior.

Corollary 3. *The seller's equilibrium expected payoff is*

$$u_s = \begin{cases} [\mu(1-\rho_H)+(1-\mu)(1-\rho_L)]v_L & \text{if} \quad \mu \leq \mu^* \\ \mu(1-\rho_H)v_H & \text{if} \quad \mu > \mu^*. \end{cases}$$

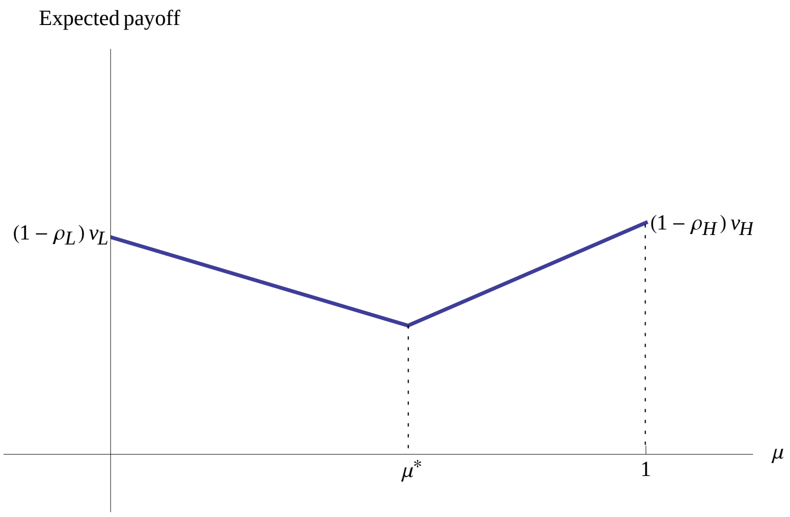

Figure 6: Seller's equilibrium payoff: random proposal model

The assumption $\rho_H > \rho_L$ implies that the seller's payoff is decreasing as long as $\mu \leq \mu^*$ and increasing for $\mu > \mu^*$. In the region of low priors the seller offers price v_L and both types accept it. The seller, however, suffers from incidence of high types, as this decreases his probability of making an offer. This property is in contrast with the above model where the probability of making an offer is endogenously determined. In that model, when expecting the low price both types of buyer make the same effort choice, and therefore, win with the same probability against the seller.

Also when the seller is charging the high price, the occurrence of high types decreases his chances of making an offer. The effect is, however, overpowered by the fact that the seller's price offer is accepted more often. The sum of the two effects is an unequivocal benefit of facing more high types for the seller.

The version of the probabilistic bargaining model–with different probabilities for different types–matches the model with efforts much more closely, in particular, it captures the non-monotonicity of the seller's payoff. It does not match it exactly, though, due to the differences in the region of priors below v_L/v_H where through endogenous choice the two types of buyer end with the same probability to make an offer.

Despite the differences between the two models, we find that a researcher who wishes to use a more manageable format will be well suited with the probabilistic model. Probabilistic model where ρ_H and ρ_L depend on μ would of course make another step towards the model with efforts.

5 Appendix

Proof of Lemma 1. The result follows from comparison of (3) and (2), and the fact that the seller never offers a price below v_L due to $E[p] \geq v_L$. □

Proof of Lemma 2. The result is argued in the text preceding the statement of the lemma. □

Proof of Lemma 3. The seller's posterior can be written as

$$\begin{aligned}\hat{\mu} &= \frac{\mu\left[\beta_H(1-\rho_1)+(1-\beta_H)(1-\rho_0)\right]}{\mu\left[\beta_H(1-\rho_1)+(1-\beta_H)(1-\rho_0)\right]+(1-\mu)\left[\beta_L(1-\rho_1)+(1-\beta_L)(1-\rho_0)\right]}\\ &\leq \frac{\mu\left[\beta_H(1-\rho_1)+(1-\beta_H)(1-\rho_0)\right]}{\mu\left[\beta_H(1-\rho_1)+(1-\beta_H)(1-\rho_0)\right]+(1-\mu)\left[\beta_H(1-\rho_1)+(1-\beta_H)(1-\rho_0)\right]}\\ &= \mu,\end{aligned}$$

where the inequality follows from $\beta_H \geq \beta_L$ and $\rho_1 > \rho_0$. □

Proof of Proposition 1. Due to Lemma 3, we know that $\mu < v_L/v_H$ implies $\hat{\mu} < v_L/v_H$. Conditional on making an offer, the seller thus strictly prefers the pooling price v_L. Given this, we have $\Delta u_b(v_L) = \Delta u_b(v_H)$, so that the low type buyer optimally exerts effort if and only if the high type optimally exerts effort. □

Proof of Proposition 2. Under the stated conditions, both types of buyer undertake the same effort choice. This implies that when the seller gets to make the offer his posterior is equal to his prior μ. If μ is smaller (greater) than v_L/v_H, the payoff associated to the pooling price, v_L, is greater (smaller) than the payoff associated to the separating price, μv_H. □

Proof of Proposition 3. Consider first the case $\mu \geq m$. Given $k_b \in ((\rho_1 - \rho_0)v_L, (\rho_1 - \rho_0)v_H)$, for $E_\sigma[p_s] = v_H$ we have $\Delta u_b(v_L) < 0$ and $\Delta u_b(v_H) > 0$. The low type buyer optimally exerts no effort, while the high type buyer optimally does. For the seller offering v_H is optimal as long as his posterior conditional on making an offer is greater than v_L/v_H. Given the equilibrium strategies of the two types of buyers, this requires

$$\frac{\mu\rho_0}{\mu\rho_0 + (1-\mu)\rho_1} \geq \frac{v_L}{v_H}.$$

Solving this inequality for μ, we obtain

$$\mu \geq \underbrace{\frac{(1-\rho_0)v_L}{(1-\rho_0)v_L + (1-\rho_1)(v_H - v_L)}}_{=m}.$$

Under the condition $\mu \geq m$, we thus have a pure strategy equilibrium where the high type exerts effort, the low type does not, and the seller offers price v_H.

Consider next the case $\mu \in (v_L/v_H, m)$. By the argument above and the fact that the seller cannot offer price v_L with certainty, there is no pure strategy equilibrium. The seller must now randomize between offering the pooling and the separating price. This requires that his posterior when he gets to make the offer is equal to v_L/v_H. Setting $\beta_L = 0$, the condition $\hat{\mu} = v_L/v_H$ becomes

$$\frac{\mu[\beta_H(1-\rho_1) + (1-\beta_H)(1-\rho_0)]}{\mu[\beta_H(1-\rho_1) + (1-\beta_H)(1-\rho_0)] + (1-\mu)(1-\rho_0)} = \frac{v_L}{v_H}$$

Solving this equality for β_H, we obtain $\beta_H = \frac{(1-\rho_0)(\mu v_H - v_L)}{\Delta\rho\mu(v_H - v_L)}$. The value of β_H lies in the interval $(0,1)$ if and only if $\mu \in (v_L/v_H, m)$. Finally, in equilibrium it must be optimal for the high type buyer to randomize over exerting and not exerting effort. This requires $\Delta u_b(v_H) = 0$, or equivalently

$$(\rho_1 - \rho_0)[\sigma v_H + (1-\sigma)v_L] = k_b$$

Solving the above equality for σ yields the expression in (7). The value of σ as in (7) belongs to the interval $(0, 1)$ if and only if $k_b \in ((\rho_1 - \rho_0)v_l, (\rho_1 - \rho_0)v_H)$. Under the stated conditions, the equilibrium as characterized in Proposition 3 thus exists and it is unique. □

Proof of Lemma 4. Setting $\beta_L \leq \beta_H$, the difference in the seller's posterior between $e_s = 1$ and $e_s = 0$ is given by

$$\begin{aligned}
& \frac{\mu\left[\beta_H \frac{1}{2} + (1-\beta_H)\rho\right]}{\mu\left[\beta_H \frac{1}{2} + (1-\beta_H)\rho\right] + (1-\mu)\left[\beta_L \frac{1}{2} + (1-\beta_L)\rho\right]} \\
& - \frac{\mu\left[\beta_H(1-\rho) + (1-\beta_H)\frac{1}{2}\right]}{\mu\left[\beta_H(1-\rho) + (1-\beta_H)\frac{1}{2}\right] + (1-\mu)\left[\beta_L(1-\rho) + (1-\beta_L)\frac{1}{2}\right]} \\
= & \frac{\mu(1-\mu)\left(\rho - \frac{1}{2}\right)^2 (\beta_H - \beta_L)}{\left(\mu\left[\beta_H \frac{1}{2} + (1-\beta_H)\rho\right] + (1-\mu)\rho\right)\left(\mu\left[\beta_H(1-\rho) + (1-\beta_H)\frac{1}{2}\right] + (1-\mu)\frac{1}{2}\right)} \\
\geq & \ 0
\end{aligned}$$

The inequality is strict when $\beta_H > \beta_L$. □

Proof of Proposition 5. Recall that $\underline{m}(k_s) = \frac{k_s}{(\rho - \frac{1}{2})v_H}$ is the prior at which the seller is indifferent between the two effort choices when planning to offer price v_H. Since we assume that $k_s > (\rho - \frac{1}{2})v_L$, we have $\underline{m}(k_s) > \frac{v_L}{v_H}$. In turn

$$\begin{aligned}
(\rho - \frac{1}{2})v_L - k_S &\leq (\rho - \frac{1}{2})\underline{m}(k_s)v_H - k_s \\
&= 0.
\end{aligned}$$

In words, the seller's benefit from exerting effort when planning to offer the low price, v_L, is smaller than the same benefit at the prior $\underline{m}(k_s)$ when planning to offer the high price, v_H. The latter is equal to 0 by the definition of $\underline{m}(k_s)$. To summarize, at the priors below $\underline{m}(k_s)$ the seller optimally refrains from exerting effort irrespective of the price he is planning to charge or the buyers behavior. It is as if the seller had no choice in effort at all. We can thus apply the analysis from Proposition 3 after setting $\rho_1 = \rho$ and $\rho_0 = \frac{1}{2}$. Notice also that $\hat{m}_0$ corresponds to m in the before-mentioned proposition. This takes care of the first three bullet points.

The only case remained to consider is $\mu > \underline{m}(k_s)$. By the definition of $\hat{\mu}_0$ the seller optimally charges the high price for the priors above $\hat{\mu}_0$. The definition of $\underline{m}(k_s)$ and the assumption $\underline{m}(k_s) > \hat{\mu}_0$ then imply that the seller prefers to exert effort for the priors above $\underline{m}(k_s)$. This incentivizes the high type buyer to exert effort too, namely his benefit of exerting effort is

$$(\rho - \frac{1}{2})v_H - k_b > 0,$$

where the inequality follows from the assumption on the buyer's cost of effort. □

Proof of Proposition 6. The proof proceeds in two steps: in the first step we provide a characterization of equilibria under some parameters; in the second step we show how the first step can be used to provide the full characterization.

Step 1. We fist argue that, if $\mu \geq \frac{v_L}{v_H}$, $k_s \in ((\rho - \frac{1}{2})v_L, (\rho - \frac{1}{2})\mu v_H)$ and $k_b \in ((\rho - \frac{1}{2})v_L, (\rho - \frac{1}{2})v_H)$, there exists a unique equilibrium with the following properties:

- if $\mu \geq \overline{m}(k_S)$, the high type buyer and the seller exert effort and the seller offers price v_H;
- if $\mu < \overline{m}(k_s)$, the high type buyer exerts effort with probability

  $$\beta_H = \frac{\rho\mu v_H - \frac{1}{2}v_L - k_s}{\mu(\rho - \frac{1}{2})(v_H - v_L)},$$

  the seller exerts effort and offers price v_H with probability

  $$\sigma = \frac{k_b - (\rho_1 - \rho_0)v_L}{(\rho_1 - \rho_0)(v_H - v_L)}; \tag{7}$$

  with complementary probability he exerts no effort and offers price v_L.

The low type buyer never exerts effort.

To prove the above claim, notice that $k_b \in ((\rho - 1/2)v_L, (\rho - 1/2)v_H)$ implies $\Delta u_b(v_L) < 0$. The low type buyer, therefore, optimally refrains from exerting effort regardless of the seller's behavior.

Consider then the seller's optimal effort choice. Since $k_s \in ((\rho - 1/2)v_L, (\rho - 1/2)\mu v_H)$, we have $\Delta u_s|_{p=v_L} < 0$, implying that $(0, v_L)$ dominates $(1, v_L)$, and $\Delta u_s|_{p_s=v_H} > 0$, further implying that $(0, v_H)$ is dominated by $(1, v_H)$. The seller, thus, chooses between $(0, v_L)$ and $(1, v_H)$. His benefit when switching from $(0, v_L)$ to $(1, v_H)$ is given by

$$\tilde{\Delta} u_s = \mu \left[\beta_H \frac{1}{2} + (1 - \beta_H)\rho\right] v_H - k_s - \left[\mu\beta_H(1 - \rho) + \mu(1 - \beta_H)\frac{1}{2} + (1 - \mu)\frac{1}{2}\right] v_L.$$

This term is decreasing in β_H and equal to zero when β_H takes the value

$$\beta_H = \frac{\rho\mu v_H - \frac{1}{2}v_L - k_s}{\mu\left(\rho - \frac{1}{2}\right)(v_H - v_L)}. \tag{8}$$

We next argue that in equilibrium the seller cannot choose $(0, v_L)$ with probability one. This would make it optimal for the high type buyer not to exert effort ($\beta_H = 0$):

$$\begin{aligned} \Delta u_b(v_H) &= (\rho - 1/2)v_L - k_b \\ &< 0, \end{aligned}$$

where the inequality follows from to the assumption $k_b \in ((\rho - 1/2)v_L, (\rho - 1/2)v_H)$; in which case

$$\begin{aligned} \tilde{\Delta} u_s &= \mu\rho v_H - k_s - 1/2 v_L \\ &> (\rho - 1/2)\mu v_H - k_s \\ &> 0, \end{aligned}$$

where the first inequality follows from the assumption $\mu v_H \geq v_L$ and the second from $k_s < (\rho - \frac{1}{2})\mu v_H$. By implication, the seller would find it optimal to deviate to $(1, v_H)$.

Consider next the possibility of an equilibrium where the seller chooses $(1, v_H)$ with probability one. This makes it optimal for the high type buyer to exert effort:

$$\begin{aligned} \Delta u_b(v_H) &= (\rho - 1/2)v_H - k_b \\ &> 0, \end{aligned}$$

where the inequality follows from $k_b \in ((\rho - 1/2)v_L, (\rho - 1/2)v_H)$. We therefore have $\beta_H = 1$. The seller has no incentives to deviate to $(0, v_L)$ if $\tilde{\Delta} u_s \geq 0$, that is:

$$\frac{1}{2}(\mu v_H - v_L) + \left(\rho - \frac{1}{2}\right)\mu v_L \geq k_s$$

Solving this inequality for μ, we obtain

$$\mu \geq \underbrace{\frac{k_s + \frac{1}{2}v_L}{\frac{1}{2}v_H + (\rho - \frac{1}{2})v_L}}_{=\overline{m}(k_s)}$$

Hence, when $\mu \geq \overline{m}(k_s)$, there exists a pure strategy equilibrium where the seller and the high type buyer exert effort, while the low type does not, and the seller offers price v_H.

When the above inequality is not satisfied, in equilibrium the seller must randomize between $(0, v_L)$ and $(1, v_H)$. This requires that the high type buyer's probability of exerting effort is equal to β_H as in (8). In particular, we need $\beta_H \in [0, 1]$. To establish

$\beta_H \geq 0$, notice that

$$\mu\rho v_H - \frac{1}{2}v_L - k_s \geq (\rho - \frac{1}{2})\mu v_H - k_s$$
$$> 0,$$

where the first inequality follows from the assumption $\mu v_H \geq v_L$ and the second from $k_s < (\rho - \frac{1}{2})\mu v_H$. On the other hand, $\beta_H \leq 1$ follows from $\mu \leq \overline{m}(k_s)$.

Randomizing is optimal for the high type buyer if he is indifferent, i.e. if $\Delta u_b(v_H) = 0$. This requires that the seller chooses $(1, v_H)$ with probability σ_H, as in (7). Taken together, this shows that the equilibrium, as described in Proposition 6, exists and that it is unique; thereby concluding Step 1.

Step 2. Notice that the condition $k_s \in ((\rho - \frac{1}{2})v_L, (\rho - \frac{1}{2})\mu v_H)$ imposed in Step 1 establishes a bound on the prior $\mu > \frac{k_s}{(\rho - \frac{1}{2})v_H} = \underline{m}(k_s)$. Moreover, $k_s > (\rho - \frac{1}{2})v_L$ implies that $\underline{m}(k_s) > \frac{v_L}{v_H}$. Consequently, as in the proof of Proposition 5, the seller optimally eschews exerting effort for priors under $\underline{m}(k_s)$. The first two bullet points of the proposition are then covered by Proposition 3, under the assumption that the seller does not expend effort, that is, imposing $\rho_1 = \rho$ and $\rho_0 = \frac{1}{2}$.

As observed above, for $\mu \geq \underline{m}(k_s)$, the analysis in Step 1 applies. This covers the last two bullet points of the proposition. □

References

Board, S. and J. Zwiebel (2012): "Endogenous competitive bargaining," Tech. rep., mimeo.

Cialdini, R. B. and N. Garde (1987): *Influence*, vol. 3, A. Michel.

Fisher, R., W. L. Ury, and B. Patton (2011): *Getting to yes: Negotiating agreement without giving in*, Penguin.

Fudenberg, D. and J. Tirole (1983): "Sequential bargaining with incomplete information," *The Review of Economic Studies*, 50, 221–247.

——— (1991): "Game theory, 1991," *Cambridge, Massachusetts*, 393, 12.

Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of dynamic monopoly and the Coase conjecture," *Journal of Economic Theory*, 39, 155–190.

Inderst, R. (2001): "Screening in a matching market," *The Review of Economic Studies*, 68, 849–868.

KRASTEVA, S. AND H. YILDIRIM (2012): "Payoff uncertainty, bargaining power, and the strategic sequencing of bilateral negotiations," *The RAND Journal of Economics*, 43, 514–536.

MÜNSTER, J. AND M. REISINGER (2015): "Sequencing bilateral negotiations with externalities," .

RUBINSTEIN, A. (1982): "Perfect equilibrium in a bargaining model," *Econometrica: Journal of the Econometric Society*, 97–109.

RUBINSTEIN, A. AND A. WOLINSKY (1985): "Equilibrium in a market with sequential bargaining," *Econometrica: Journal of the Econometric Society*, 1133–1150.

STAHL, I. (1972): "Bargaining Theory (Stockholm School of Economics)," *Stockholm, Sweden.*

TRUMP, D. J. AND T. SCHWARTZ (2009): *Trump: The art of the deal*, Ballantine Books.

YILDIRIM, H. (2007): "Proposal power and majority rule in multilateral bargaining with costly recognition," *Journal of Economic Theory*, 136, 167–196.

——— (2010): "Distribution of surplus in sequential bargaining with endogenous recognition," *Public choice*, 142, 41.

ZINGALES, L. (1995): "Insider ownership and the decision to go public," *The Review of Economic Studies*, 62, 425–448.