

DISCUSSION PAPER SERIES

DP13542

OPTIMAL SELLING MECHANISMS WITH ENDOGENOUS PROPOSAL RIGHTS

Nenad Kos, Sarah Auster and Salvatore Piccolo

INDUSTRIAL ORGANIZATION

OPTIMAL SELLING MECHANISMS WITH ENDOGENOUS PROPOSAL RIGHTS

Nenad Kos, Sarah Auster and Salvatore Piccolo

Discussion Paper DP13542
Published 21 February 2019
Submitted 19 February 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Nenad Kos, Sarah Auster and Salvatore Piccolo

OPTIMAL SELLING MECHANISMS WITH ENDOGENOUS PROPOSAL RIGHTS

Abstract

We study a model of optimal pricing where the right to propose a mechanism is determined endogenously: a privately informed buyer covertly invests to increase the probability of offering a mechanism. We establish the existence of equilibrium and show that higher types get to propose a mechanism more often than lower types allowing the seller to learn from the trading process. In any equilibrium, the seller either offers the price he would have offered if he was always the one to make an offer or randomises over prices. Pure strategy equilibria may fail to exist, even when types are continuously distributed. A full characterization of equilibria is provided in the model with two types, where notably the seller's profit is shown to be non-monotonic in the share of high-value buyers.

JEL Classification: C72, D82, D83

Keywords: mechanism design, Optimal Pricing, bargaining power

Nenad Kos - nenad.kos@unibocconi.it
Bocconi University and CEPR

Sarah Auster - sarah.auster@unibocconi.it
Bocconi University, IGIER

Salvatore Piccolo - salvatore.piccolo@unibg.it
University of Bergamo

Optimal Selling Mechanisms with Endogenous Proposal Rights

Sarah Auster* Nenad Kos † Salvatore Piccolo ‡

January 18, 2019

Abstract

We study a model of optimal pricing where the right to propose a mechanism is determined endogenously: a privately informed buyer covertly invests to increase the probability of offering a mechanism. We establish the existence of equilibrium and show that higher types get to propose a mechanism more often than lower types, allowing the seller to learn from the trading process. In any equilibrium, the seller either offers the price he would have offered if he was always the one to make an offer or randomises over prices. Pure strategy equilibria may fail to exist, even when types are continuously distributed. A full characterization of equilibria is provided in the model with two types, where notably the seller's profit is shown to be non-monotonic in the share of high-value buyers.

JEL Code: C72, D82, D83

Keywords: Mechanism Design, Optimal Pricing, Bargaining Power

*Bocconi University, Department of Decision Sciences and IGIER. E-mail: sarah.auster@unibocconi.it

†Bocconi University, Department of Economics and IGIER. E-mail: nenad.kos@unibocconi.it

‡University of Bergamo, E-mail: salvapiccolo@gmail.com. We would like to thank Alp Atakan, Mehmet Ekmekci, Dino Gerardi, Sander Heinsalu, Toomas Hinnosaar, Johannes Hörner, Stephan Lauer-
mann, Ole Jann, Markus Reisinger, Thomas Wiseman and seminar audiences at Frankfurt School of
Management, Essex, LSE, 2016 CSEF-IGIER symposium, Workshop on Economic Theory, Graz 2016,
Workshop on Advances in Industrial Organization, Bergamo 2017, and University of Ljubljana. This
paper subsumes the results previously appearing in Investing in Bargaining Power.”

Contents

1	Introduction	1
2	The Model	5
3	Equilibrium Analysis	6
4	Binary Types	11
5	Probabilistic Bargaining Model	15
6	Concluding Remarks	16
A	Appendix	17

1 Introduction

The question of how economic agents trade has spurred an immense literature with powerful insight. The main focus lies on the features of the optimal mechanism under informational asymmetries and how they are shaped by scarcity and competition. The latter two have long been recognised as driving forces of bargaining power. The literature has extensively explored how optimal mechanisms and the principal’s payoff depend both on the number of buyers, as well as on the presence of competing sellers.¹ For instance, it is well known that a seller can extract almost full surplus with a mechanism as simple as a second price auction if the number of participating buyers is large. Similar ideas have been discussed already by Adam Smith, who argued that masters have greater bargaining power than workers because they are fewer in numbers; [Smith \(1776\)](#).² Little attention has, however, been dedicated to the question of how bargaining power is distributed in contracting problems when no side has a clear competitive advantage. In such situations, agents have incentives to find ways that improve their bargaining position: potential buyers search the internet for information on how to get a good deal from car salesmen, companies and individuals hire lawyers or other experts when selling or buying valuable assets, parties in a troubled marriage hire divorce lawyers to negotiate on their behalf, just to name a few. Although it is recognised that who determines the mechanism not only matters for the division of the surplus but also for many other features of the equilibrium allocation,³ by and large, the right to propose a mechanism has been treated as exogenous. In this paper, we consider the design of optimal mechanisms in a model with endogenous proposal rights.

We introduce a bilateral trade problem between a seller and a buyer with the novel feature that the proposer of the mechanism is endogenously selected. More precisely, the buyer, privately informed about his valuation of the good, can make a covert investment in order to improve his bargaining position—e.g., the resources spent to find and hire the right lawyer (or other representatives), the amount of attention and time he devotes to specialised courses on ‘the art of negotiation’ (see [Cialdini and Garde \(1987\)](#), [Fisher et al. \(2011\)](#)), etc. Higher investment, naturally, translates into a higher probability of proposing a mechanism. Though investments are not observable, obtaining the right to make a proposal is a signal about the other party’s investment. Our focus lies on the

¹For competition in mechanism design, see [McAfee \(1993\)](#), [Peters and Severinov \(1997\)](#) and for a recent review of the literature [Pai \(2010\)](#).

²“It is not, however, difficult to foresee which of the two parties must, upon all ordinary occasions, have the advantage in the dispute, and force the other into a compliance with their terms. The masters, being fewer in number, can combine much more easily...”

³For instance, [Hagedorn and Manovskii \(2008\)](#) showed that the distribution of bargaining power between workers and firms is an important determinant for the outcomes of search models regarding wages and unemployment volatility.

buyer's investment, since he is the privately informed agent. The party with the proposal power determines the mechanism, which in our bilateral trade environment takes the form of a (distribution over) posted price(s). After the mechanism is executed the game ends; that is, the proposer commits to the mechanism.

We do not claim that this is precisely how trading interactions unfold. Rather, our environment should be interpreted as a reduced form model capturing two salient features. First, if a buyer invests, he is able to obtain a higher share of the surplus, which is captured by an increase in the probability that he gets to make an offer. Second, agents might be able to learn about their opponents during the trading process. Specifically, if a certain type of buyer invests less 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 making an offer, therefore, conveys information about the buyer's willingness to pay. Our objective is to analyze how these two elements interplay, and what novel insights their interaction delivers.

The driving force of our paper is the observation that higher types of the buyer have higher incentives to invest. A buyer with a higher valuation offers more surplus to be shared and therefore has more incentives to fight for it. Hence, the higher the buyer's valuation, the more likely the buyer is to make an offer, establishing a novel source of bargaining power.

The effect of the buyer's valuation on bargaining power has important implications on the seller's optimal pricing policy. Taking into account that the seller is more likely to be able to make an offer against lower types, the seller revises his posterior (conditional on being able to make an offer) towards lower valuations. Despite the fact that this favours lower prices, our main result shows that in any equilibrium the ex-ante optimal price—the price the seller would charge if he were always the one to make an offer—belongs to the support of the seller's price distribution. Moreover, the ex-ante optimal price is the highest price in that support. As an immediate consequence, the only candidate pure-strategy equilibrium is the one where the seller offers the ex-ante optimal price. Due to the seller's updating, however, such equilibrium may not exist. In that case, in equilibrium the seller randomises across a set of prices—a feature that is non-generic when bargaining power is exogenously fixed.

Whether (or not) the pure-strategy equilibrium exists in a given environment is easy to check. Roughly speaking, such equilibrium does not exist when there is a price sufficiently below the ex-ante optimal price that yields only a slightly lower profit than the ex-ante optimal price. Since after updating the seller assigns relatively more weight to types between the two prices, the profit from the lower price increases with respect to that of

the ex-ante optimal one. We construct an example in which, despite valuations being distributed on an interval according to a distribution with no gaps or mass points, all equilibria are in mixed strategies. In the equilibrium we describe, the seller randomises over two prices.

We provide a full characterization of equilibria for the case where the buyer's valuation is binary. In this situation, the seller optimally offers a price equal to one of the two valuations. The equilibrium is generically unique and is most readily described as a function of the ex-ante probability with which the buyer's valuation is high. When this probability is low, the ex-ante optimal price is equal to the low valuation of the buyer. Since the seller's posterior belief about the high type can only decrease with respect to the prior, as argued above, this price remains optimal after updating. When the seller's prior favors the high price ex ante just slightly, the buyer's equilibrium investment strategy makes the seller indifferent between both prices after updating. In equilibrium the seller offers both the low and the high price with positive probability. Finally, when the probability that the buyer's type is high is sufficiently large, the seller always offers the high price.

A striking feature of the binary model is that the seller's expected payoff is non-monotonic in the proportion of high types. It is flat for low priors, decreasing in the intermediate region and increasing for high priors. When the prior is low, both types of the buyer trade with probability one, whether they or the seller make the offer. The benefit of making an offer themselves is therefore equal to the price the seller would charge if he were to make the offer. Since this price is the same for both types, so are their incentives to invest. Hence, both types of the buyer invest the same amount, making it irrelevant for the seller whether he faces a high or a low type. The most intriguing region is the intermediate one, where the seller randomises over prices. Since the seller is indifferent between the two prices, his expected payoff can be evaluated at the low price. This price is accepted by both types, implying that the seller's payoff conditionally on being able to make an offer is constant in the discussed region. However, as the probability of the high type increases, the high type's investment increases and the seller faces the high type more often. As a result, the probability that the seller gets to make an offer falls. Taken together, this implies that the seller's expected payoff is strictly decreasing in the share of high types. Finally, in the region where the seller offers the high price, his payoff is increasing. Since the low-value buyer rejects the high price, the seller benefits only from the high type.

The above comparative static result implies that if the seller could choose to deal with one of two buyers, he might be more inclined to approach a buyer who is more likely to

be a low type. While the surplus to be shared is smaller, the buyer spends fewer resources towards securing that surplus. Thus, the seller might prefer the easy bargain.

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 results suggest 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 and yet captures some of the arresting features of our framework.

Related Literature: [Myerson \(1981\)](#) studied a mechanism design model in which the seller is always the one to choose a mechanism. Among other things, he shows that if the seller is selling to only one buyer, he can always restrict himself to a posted price. Moreover, the seller can extract nearly full surplus as the number of buyers grows. The seller's bargaining power and, through that, the ability to extract surplus are put to the test in the models of competition in mechanism design; for an outline of the literature see [Pai \(2010\)](#). We explore an alternative way of reducing the seller's bargaining power, namely by allowing the buyer to fight for his right to propose a mechanism himself.

To the best of our knowledge, we are the first to consider a model of endogenous proposal rights in the context of optimal pricing with commitment (mechanism design). Models where agents compete for proposal rights have been studied in the context of bargaining without commitment; see, for instance, [Yildirim \(2007\)](#), [Yildirim \(2010\)](#), [Board and Zwiebel \(2012\)](#) and [Ali \(2015\)](#). As opposed to our work, these papers analyse environments where the surplus is of commonly known size.⁴ Under the same assumption [Karagozoglu and Rachmilevitch \(2018\)](#) study a model where agents pay a cost to prepare for negotiations. In their framework, costly preparations do not affect the allocation of proposal rights, but are instead necessary to remain at the negotiation table.

Other papers analyse trading models with the feature that an agent partially learns from the interaction. In [Lauermann and Wolinsky \(2017\)](#), a privately informed seller gets to select a number of buyers that participate in a first price auction for an object. The fact that they have been selected reveals something about the common value of the object, similarly as being able to make an offer reveals something to the seller in our model. However, the objective is rather different from ours. While they study equilibria

⁴A somewhat different, but related, literature—see [Crawford \(1979\)](#), [Evans \(1997\)](#) and [Pérez-Castrillo and Wettstein \(2001\)](#)—studies the effects of competition on coalitional bargaining games.

of a particular mechanism (the first price auction) and how they aggregate information, we study optimal pricing. [Heinsalu \(2017\)](#) analyses a model with endogenous entry into a competitive market with adverse selection. Entry is informative about the entrant's private information and, therefore, reflected in the market price.

2 The Model

A buyer and a seller want to trade a good. The buyer's valuation v is drawn from a cumulative distribution function F on $[0, 1]$; the seller's is commonly known and normalised to zero. Upon observing his valuation, the buyer can make an investment to increase the probability with which he gets to make an offer.⁵ More precisely, the buyer chooses an investment $e \in \mathbb{R}_+$ at a cost ke , after which he gets to make an offer with probability $\rho(e)$, where $\rho : \mathbb{R}_+ \rightarrow (0, 1)$ is a concave and increasing function.⁶ The seller gets to make an offer with the complementary probability. After the offer is accepted/rejected the game ends.

Remark 1. An example to keep in mind is that of a Tullock contest in which agents' chances to make an offer are proportional to their effort investments; [Tullock \(1980\)](#). In particular, $\rho(e) = e/(e + e_s)$, where $e_s > 0$ is the seller's (fixed) effort.

The buyer's pure strategy is a tuple of functions (e, p_b, x_b) , where $e : [0, 1] \rightarrow \mathbb{R}_+$ maps the buyer's types into the investment choices, $p_b : [0, 1] \rightarrow \mathbb{R}_+$ specifies the price the buyer offers when he gets to make an offer (in the remainder of the text we suppress the qualifier "when he gets to make an offer"), and $x_b : [0, 1] \times \mathbb{R}_+ \rightarrow \{0, 1\}$ specifies whether a type accepts (1) or rejects (0) a given price. The seller's pure strategy is a tuple (p_s, x_s) , where $p_s \in \mathbb{R}_+$ is the price the seller offers and $x_s : \mathbb{R}_+ \rightarrow \{0, 1\}$ is the seller's acceptance strategy.

The equilibrium concept we employ is sequential equilibrium; equilibrium for short. While a sequential equilibrium may not be well defined for games with a continuum of types and actions in general, such issues will not arise in our framework. In fact, the only possible out of equilibrium continuation games are the ones reached after unexpected prices. Given our private values assumption, out of equilibrium beliefs in such information sets are irrelevant.

⁵In an earlier version of the paper, we considered the case where also the seller invests. Since the buyer knows the seller's valuation and, hence, does not learn, we abstract from the seller's investment choice here.

⁶By assuming that the range of ρ is $(0, 1)$ we avoid rather uninteresting problems brought upon by the possibility of out of equilibrium beliefs.

The outlined model, as always, is a simplified representation of reality. We do not claim that the trading interaction proceeds precisely as specified above. However, the model captures two important features that we would like to explore: 1.) A larger investment 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. The existing literature commonly uses the probability that an agent makes a take-it-or-leave-it offer as a proxy for bargaining power (see for example Zingales (1995) and Inderst (2001)). 2.) By observing a signal about the opponent's investment, an agent may learn 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 the buyer making (possibly) different investment decisions.

3 Equilibrium Analysis

It is useful to start with some preliminary observations. If the buyer gets to make an offer, he optimally offers price $p_b = 0$ and the seller accepts. The buyer's strategy can therefore be reduced to his choice of investment e and his decision to accept/reject. The buyer's optimal investment depends on the seller's pricing strategy. The seller's randomisation over prices can be described by a cumulative function H , where $H(p)$ is the probability that the seller offers a price weakly smaller than p . Given H , the buyer solves the problem

$$\max_{e \in [0, \infty)} \rho(e)v + (1 - \rho(e)) \int_0^v (v - p)dH(p). \quad (1)$$

After investing e the buyer gets to make an offer with probability $\rho(e)$ and obtains a payoff v . With the complementary probability the seller makes an offer, which the buyer accepts as long as his valuation is not below the price. The following lemma characterises the solution to the buyer's problem and shows that the optimal investment is non-decreasing in the buyer's valuation. The proofs of this and all following results are collected in the appendix

Lemma 1. *If the buyer expects the seller to randomise across prices according to H , his optimal investment is uniquely described by the function*

$$e^*(v; H) = (\rho')^{-1} \left(\frac{k}{v(1 - H(v)) + \int_0^v pdH(p)} \right),$$

for all v such that $\rho'(0) (v(1 - H(v)) + \int_0^v tdH) \geq k$ and $e^*(v; H) = 0$ otherwise. For all H , the function $e^*(v; H)$ is non-decreasing in v .

The buyer's investment function is equal to zero for types close to zero, increasing up to the lowest type who always trades when the seller makes an offer, and constant thereafter; see Figure 1.⁷ This is easy to see when the buyer expects the seller to offer a (deterministic) price p . For types below the seller's price p , the loss of not being able to make an offer is equal to their type (the surplus they get when making an offer themselves), and as such strictly increasing in v .⁸ Types above p , on the other hand, trade even when the seller makes the offer. The loss they incur when the seller makes the offer is the price they have to pay. This loss is constant for all types above the price, as are therefore the incentives to invest. Notice further that the buyer's expected payoff is strictly concave in e , which means that each type has a unique best response. In equilibrium the buyer thus plays a pure strategy.




Figure 1: Buyer's optimal investment when expecting the seller to offer price p .

The monotonicity of the buyer's investment implies that higher types are more likely to be able to make an offer and, conversely, that the seller is more likely to make an offer against the lower types. Given the buyer's investment strategy $e(v)$, the probability that type v makes an offer is $\rho(e(v))$. Using this probability, the seller can compute the posterior distribution over the buyer's valuations, conditionally on being able to make the offer himself.

$$G(v; e) = \frac{\int_0^v (1 - \rho(e(t))) dF(t)}{\int_0^1 (1 - \rho(e(t))) dF(t)}. \quad (2)$$

Notice that F first-order stochastically dominates G . While being able to make an offer is good news for the seller, he understands that he is more likely to win against lower

⁷In the situation we describe, k is sufficiently small such that some types want to invest. Instead, when k is large, the optimal investment is zero for all types.

⁸For types close to zero, the gain from making an offer may not be sufficient to compensate the cost of investment, so they prefer not to invest.

types of the buyer. Winning, therefore, depresses his belief. Given the buyer's investment strategy e , the seller then solves the problem

$$\max_{p \in \mathbb{R}_+} (1 - G(p; e))p. \quad (3)$$

The revised beliefs dictate the optimal price, leading to a rather interesting fixed point problem. The buyer conjectures the price (or the distribution over prices) that the seller will charge and chooses his investment in response. The seller updates his beliefs and charges a price that is optimal with respect to his posterior. In equilibrium, the buyer's conjecture about the seller's price distribution must coincide with the seller's optimal price distribution. The following lemma shows that this problem has a fixed point.

Lemma 2. *An equilibrium of the game exists.*

The main obstacle to showing the existence is presented by the discontinuity of the seller's payoff in the price. If there is a mass point in the buyer's distribution over valuations, the seller's payoff jumps up as the price approaches the mass point from above. Here we rely on Reny's (1999) result on the existence of mixed strategy equilibria in discontinuous games with infinite strategy spaces. Before applying the result we simplify the game by using some of the above derived properties of equilibria. In particular, in any sequential equilibrium the buyer accepts a price offer if and only if the price is at least as high as his valuation and the seller if it is non-negative; we assume that the agent accepts an offer if indifferent. Moreover, knowing that the seller will accept any non-negative price, the buyer optimally offers price 0. The auxiliary game assumes such behavior. Leaving each type of the buyer to decide the investment and the seller to choose the price. After proving the existence of an equilibrium in the auxiliary game, we show that it embeds as a sequential equilibrium in the original game.

Equilibrium pricing. Having established the existence of an equilibrium, we turn attention to the question of the seller's pricing strategy. If the seller were always the one to make an offer, as is commonly assumed, or if he neglected learning, he would offer the price that maximises $p(1 - F(p))$; for short the ex-ante optimal price, denoted by p^* . We abstract from the non-generic cases where $p(1 - F(p))$ has multiple maximisers. In what follows, we show that p^* is in the support of the seller's randomisation over prices in every equilibrium. Moreover, it is the upper bound of the support.

Proposition 1. *Fix an equilibrium. Then p^* is the maximum of the support of the seller's distribution over prices. As a consequence, every equilibrium takes one of the following two forms:*

- *it is a pure-strategy equilibrium in which the seller offers price p^* ;*
- *it is a mixed-strategy equilibrium in which the seller randomises over a set of prices of which p^* the highest.*

Proposition 1 shows that in any equilibrium the seller finds it optimal to offer the ex-ante optimal price p^* , and that the seller will never offer a price higher than p^* . One might be tempted to conclude that the latter property is a consequence of the fact that the seller's posterior first-order stochastically dominates the prior distribution. This is, however, not enough to claim that the optimal price under the posterior is not higher than p^* . A more direct argument is needed. Let \bar{p} be the maximum of the support of the distribution describing the seller's randomisation and, towards a contradiction, assume $\bar{p} > p^*$. In this case, the buyer's best-reply investment is strictly increasing in his type on the interval $[p^*, \bar{p}]$. As compared to the prior, the seller is therefore more likely to make an offer against types in (p^*, \bar{p}) than against types above \bar{p} . This implies that after updating, the price \bar{p} becomes less attractive relative to p^* . Given that offering p^* dominates offering \bar{p} ex-ante (i.e. in the absence of any possibility for the buyer to invest), it follows that p^* must dominate \bar{p} also ex-post (i.e. when investment is possible), contradicting the assumption that \bar{p} is in the support of the seller's equilibrium price distribution.

What might be more surprising is that the upper bound of the set of equilibrium prices cannot be smaller than p^* . This has to do with the fact that all types above the upper bound \bar{p} make the same investment (see Lemma 1). As a consequence, the relative posterior probabilities across these types are the same as under the prior and, by the same token, the maximizer over the prices in $[\bar{p}, 1]$ is the same before and after updating. Now if $\bar{p} < p^*$, the fact that p^* is optimal ex-ante implies that it also yields the highest profit among the prices in $[\bar{p}, 1]$ after the revision of beliefs. Hence, p^* would yield a strictly higher profit than \bar{p} , again contradicting the assumption that \bar{p} belongs to the set of prices over which the seller randomises.

Given these observations, the only possible pure-strategy equilibrium is the one where learning plays no role and the seller offers the ex-ante optimal price p^* . Existence of such an equilibrium can be checked by computing the buyer's best response to price p^* , as specified in Lemma 1, and verifying that the seller has no incentives to deviate to a lower price after updating his beliefs.⁹ It is then easy to see when the existence of a pure-strategy equilibrium fails. Namely, this happens when based on the prior there is a price p' sufficiently below p^* that achieves a profit sufficiently close to that generated by

⁹The pure strategy equilibrium always exists if the cost of investment is sufficiently large so that when the buyer expects the sellers to offer p^* , zero investment is optimal for all types. The seller does not learn from having the opportunity to make an offer and, consequently, offers price p^* .

p^* . In such a case, if the buyer expects p^* , the seller's posterior shifts sufficiently towards lower types in order to make the profit from offering p' larger than that from offering p^* . The following example shows that a pure-strategy equilibrium can fail to exist even in environments where F has a density on an interval.

Example 1. Let the function $\rho(e) = e/(e + e_s)$ be a Tullock function with $e_s > 0$ and consider a probability distribution described by the CDF

$$F(v) = \begin{cases} av & \text{if } 0 \leq v \leq d, \\ c + bv & \text{if } v > d, \end{cases}$$

where $b = (1 - ad)/(1 - d)$ and $c = -(1 - a)d/d$. Note that F is continuous on its support and the density is constant on each of the two subintervals. We specify the following parameters: $a = 1.5, d = 0.4, e_s = 0.5, k = 0.4$, with the ex-ante optimal price being $p^* = 0.5$. When the buyer expects price 0.5 and invests optimally, the seller's expected payoff evaluated at his posterior belief is no longer maximised at 0.5 but instead at ≈ 0.326 . Hence, there is no pure-strategy equilibrium. However, there is an equilibrium where the seller randomises across two prices: the seller offers $p \approx 0.332$ with probability ≈ 0.189 and the ex-ante optimal price $p^* = 0.5$ with the complementary probability. This is illustrated in Figure 2. \square




Figure 2: Seller's expected payoff before updating (solid curve) and in equilibrium (dashed curve)

The example illustrates that even when the distribution of the buyer's valuation has no atoms in a mixed-strategy equilibrium the seller need not randomise over an interval of prices. Since we impose little restriction on the type distribution, we cannot exclude the possibility of multiple equilibria. Further insights can be gained by exploring the environment with binary types which we do in the following section.

4 Binary Types

Suppose the buyer's valuation belongs to the set $\{v_L, v_H\}$, with $0 < v_L < v_H$, and let μ denote the ex-ante probability that the buyer's valuation is high. We assume that the marginal cost of investment is low enough for the low type to invest: $k \leq \rho'(0)v_L$.¹⁰

First, notice that in equilibrium the seller only offers prices equal to v_L or v_H with positive probability. Proposition 1 then implies that when the ex-ante optimal price is v_L —equivalently, when $\mu < \frac{v_L}{v_H}$ —the only possible equilibrium is the one in which the seller offers the price v_L .

The rest of the analysis focusses on the case where the ex-ante optimal price is v_H , i.e. $\mu > v_L/v_H$. By Lemma 2 we know an equilibrium exists. Let σ be the probability that the buyer attaches to the seller charging the price v_H . The buyer's optimal investment is then:

$$e^*(v) = \begin{cases} (\rho')^{-1}\left(\frac{k}{\sigma v_H + (1-\sigma)v_L}\right) & \text{if } v = v_H, \\ (\rho')^{-1}\left(\frac{k}{v_L}\right) & \text{if } v = v_L. \end{cases}$$

Whenever the seller makes an offer, the low type's payoff is 0—either the price is too high to trade or the seller extracts the full surplus. Consequently, the low type's payoff does not depend on the probability of the seller charging the high price. The high type's payoff, on the other hand, does. The larger is σ , the higher is the expected price the high type has to pay when the seller makes the offer. Concavity of ρ implies that $(\rho')^{-1}$ is a decreasing function and, hence, that the high type's investment $e^*(v_H)$ is increasing in σ . In other words, the higher is the expected price the seller charges, the larger is the buyer's benefit of making the offer himself and, thus, the incentives to invest.

When σ is strictly positive, the high type invests more than the low type and the seller updates his beliefs accordingly.¹¹ The seller's posterior belief about the buyer's type being high, conditional on the seller making an offer, is

$$\hat{\mu} := \frac{\mu \left(1 - \rho \left((\rho')^{-1} \left(\frac{k}{\sigma v_H + (1-\sigma)v_L} \right) \right) \right)}{\mu \left(1 - \rho \left((\rho')^{-1} \left(\frac{k}{\sigma v_H + (1-\sigma)v_L} \right) \right) \right) + (1 - \mu) \left(1 - \rho \left((\rho')^{-1} \left(\frac{k}{v_L} \right) \right) \right)}.$$

The seller's posterior $\hat{\mu}$ is strictly decreasing in σ and takes value μ (equal to the prior) when $\sigma = 0$. Given the restriction $\mu > v_L/v_H$, two cases need to be considered. If $\hat{\mu}$ evaluated at $\sigma = 1$ is weakly greater than v_L/v_H , the ex-ante optimal price $p^* = v_H$ is optimal after updating for all values of σ . In this case the only equilibrium is the

¹⁰The case analysed here is the most interesting one. Considering other configurations of k would increase the number of cases without adding much substance.

¹¹When the seller offers the low price with certainty, both types trade with probability one. Their benefit from making an offer is the price they do not pay to the seller (v_L) which is the same for both. In consequence the two types invest the same amount.

pure-strategy equilibrium, where the seller offers the price v_H . On the other hand, if $\hat{\mu}$ evaluated at $\sigma = 1$ is smaller than v_L/v_H , there is a unique value of σ at which the seller's posterior $\hat{\mu}$ coincides with v_L/v_H , rendering him indifferent between the low and the high price. This value of σ is the randomisation strategy of the seller in the mixed-strategy equilibrium that obtains. Letting m be the value of μ at which $\hat{\mu}$ evaluated at $\sigma = 1$ is exactly v_L/v_H , the discussion can be summarised as follows.

Proposition 2. *In the environment with two types, the equilibrium is generically unique. The seller charges price v_H with probability*

$$\sigma = \begin{cases} 0 & \text{if } \mu < \frac{v_L}{v_H}, \\ \sigma^* & \text{if } \mu \in (\frac{v_L}{v_H}, m), \\ 1 & \text{if } \mu \geq m, \end{cases}$$

where σ^* is the unique solution to

$$\frac{1}{1 + \frac{1-\mu}{\mu} \frac{1-\rho((\rho')^{-1}(\frac{k}{v_L}))}{1-\rho((\rho')^{-1}(\frac{k}{\sigma v_H + (1-\sigma)v_L}))}} = \frac{v_L}{v_H}, \quad (4)$$

and v_L with the remaining probability.

In a situation where the seller always makes the offer, the case where he is indifferent between the high and the low price is non-generic. Proposition 2 shows that when proposal rights are endogenous instead, there is always an interval of values of μ where either of the two prices is optimal for the seller in equilibrium. Moreover, the seller's equilibrium probability of offering the high price as a function of the prior μ has two jumps, taking a value in $(0, 1)$ in the intermediate region.

As we have seen, the seller revises his belief towards the low type when he gets to make an offer. When μ is below v_L/v_H so that the ex-ante optimal price is v_L , this price remains optimal after updating. Instead, when μ is just above v_L/v_H , the prior favors the high price only slightly. If the buyer expects the high price and chooses his investment accordingly, the seller's posterior would drop below v_L/v_H , making the low price optimal. However, if the buyer expects the seller to charge the low price, both types of the buyer would invest the same amount and the seller's posterior would stay above v_L/v_H , making the high price optimal. The seller, thus, randomises over the two prices. Finally, when μ is higher than m , the posterior belief $\hat{\mu}$ is above v_L/v_H , regardless of the probability with which the seller offers the high price. In consequence, the price v_H is optimal after updating and, hence, the pure-strategy equilibrium with the high price obtains.

Comparative Statics. It is of interest to study how the seller's payoff changes with the proportion of the high types.

Proposition 3. Let $u_s(\mu)$ denote the seller's equilibrium payoff as a function of the prior μ . Then,

$$u'_s(\mu) \begin{cases} = 0 & \text{if } \mu < v_L/v_H \\ < 0 & \text{if } \mu \in (v_L/v_H, m) \\ > 0 & \text{if } \mu \geq m. \end{cases}$$




Figure 3: Seller's equilibrium expected payoff

The striking property of the seller's payoff is that it is non-monotonic in the probability of the high type, in particular, that it is decreasing on the interval $[v_L/v_H, m]$ (illustrated in Figure 3). A closer inspection of the structure of the equilibrium 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 both types of the buyer make the same investment. His payoff in this region is, therefore, independent of the prior.

To see why the seller's payoff is decreasing in the intermediate region, recall that in the relevant region of priors the seller is indifferent between the two price offers. His payoff from offering the low price is

$$\left[\mu \left(1 - \rho \left((\rho')^{-1} \left(\frac{k}{\sigma v_H + (1 - \sigma)v_L} \right) \right) \right) + (1 - \mu) \left(1 - \rho \left((\rho')^{-1} \left(\frac{k}{v_L} \right) \right) \right) \right] v_L,$$

where the term in the square brackets is the probability that he makes an offer and v_L is his conditional payoff in that event—the seller is randomizing over the two prices and the price v_L is accepted with certainty. The conditional payoff is thus independent of the seller's prior. As to the probability of the seller making an offer, it is useful to recall that in the mixed-strategy equilibrium, the seller's posterior belief $\hat{\mu}$ is fixed to v_L/v_H . When the prior rises, this entails that the high-type buyer distinguishes himself from the low type by increasing his investment, thereby keeping the seller's posterior belief constant.




Figure 4: Seller's probability of making an offer and posterior belief in equilibrium

Hence, the seller's probability of making an offer is falling for two reasons: as μ increases, 1) the probability that the high type makes an offer increases, and 2) the seller faces the high type more often, as we illustrate in Figure 4. This, together with the fact that the seller's payoff conditional on making the offer is constant, implies that the seller's expected payoff is decreasing in the prior probability of the buyer's type being high.

For high priors, $\mu > m$, the seller optimally offers the separating price v_H . The high type's investment is constant in μ , which means the probability of the seller making an offer against the high type is constant as well. Since the high price v_H is only accepted by the high type, this implies that above m the seller's payoff is strictly increasing in μ .

Taken together, the seller's expected payoff is non-monotonic in the prior: it is constant for low priors, decreasing for intermediate priors and increasing for high priors. Whether for high priors the seller's payoff increases above the value it attains for low priors depends on the parameters of the problem. The seller's expected payoff is smaller at $\mu = 1$ than at $\mu = 0$ when the difference in the buyer's valuation is small relative to the difference in probabilities with which the seller gets to make the offer. This suggests that the case for the seller's payoff being decreasing in the fraction of high types could have been made in an environment with perfect information. Our result is stronger: even when the seller's payoff at $\mu = 1$ is larger than at $\mu = 0$, there is an intermediate region of priors under which the seller is strictly worse off than when he faces the low-value buyer with certainty.

The non-monotonicity of the seller's payoff stands in stark contrast to models where the probability of each agent making an offer is exogenously determined and fixed over types; more about this in following section. The fact that the seller's payoff can be decreasing in the probability of the high type has the following interesting application. If the seller could approach one of two potential buyers (or enter one of two markets)—the first consumer almost certainly to have a low valuation, the second more likely to have a high valuation—the seller might just prefer to trade with the agent who is likely to value the good less. There is less surplus to be split between the two parties, but the buyer will not fight aggressively for it either. The seller prefers, so to say, to pick the low hanging fruit.

5 Probabilistic Bargaining Model

In the probabilistic bargaining model, a buyer and a seller get to make an offer with a fixed probability, which stands as a proxy for their bargaining power. The probability of making an offer is exogenously fixed and independent of the agents' valuations. The elegance and simplicity of such models make them a popular tool in economics and finance; see, for example, [Zingales \(1995\)](#) and [Inderst \(2001\)](#), though probabilistic bargaining goes back at least to [Rubinstein and Wolinsky \(1985\)](#). Yet our results show that probabilistic bargaining models might be simplifying too much. If agents can affect their bargaining power, the probability of making an offer should increase in the type. Moreover, a fixed probability of making an offer does not allow for learning, nor can it deliver the non-monotonicity of the seller's payoff. A natural question arises whether probabilistic bargaining models can be generalised enough to capture the salient features of endogenously generated bargaining power, while preserving sufficient tractability to be interesting for wider applications.

We propose a modification of the probabilistic bargaining model, where the buyer's probability of making an offer increases in his valuation. In particular, we assume that type v of the buyer makes an offer with probability $\rho_0(v)$, where $\rho_0(v)$ is a strictly increasing function. The seller makes an offer with the complementary probability. Under this assumption, the opportunity to make an offer carries information for the seller and thus affects his pricing decision.

For the case of binary types, we saw that the seller offers the high price v_H if his posterior belief conditional on making an offer is higher than v_L/v_H (see [Section 4](#)). That is:

$$\frac{\mu(1 - \rho_0(v_H))}{\mu(1 - \rho_0(v_H)) + (1 - \mu)(1 - \rho_0(v_L))} \geq \frac{v_L}{v_H}. \quad (5)$$

Letting m_0 denote the value of μ at which condition (5) is satisfied with equality, the seller optimally offers the high price v_H if and only if $\mu \geq m_0$. Given the assumption $\rho_0(v_H) > \rho_0(v_L)$, the seller revises his beliefs downwards upon getting the right to propose. Therefore, he is willing to offer the high price only if his prior favors the high price sufficiently, implying that the threshold m_0 is strictly greater than v_L/v_H . We can therefore write the seller's expected payoff as a function of the prior μ .

$$u_s = \begin{cases} [\mu(1 - \rho_0(v_H)) + (1 - \mu)(1 - \rho_0(v_L))]v_L & \text{if } \mu \leq m_0 \\ \mu(1 - \rho_0(v_H))v_H & \text{if } \mu > m_0. \end{cases}$$

For low priors, the seller's payoff is decreasing in μ . The seller offers the low price and both types accept it. The seller, however, suffers from the incidence of high types, as they

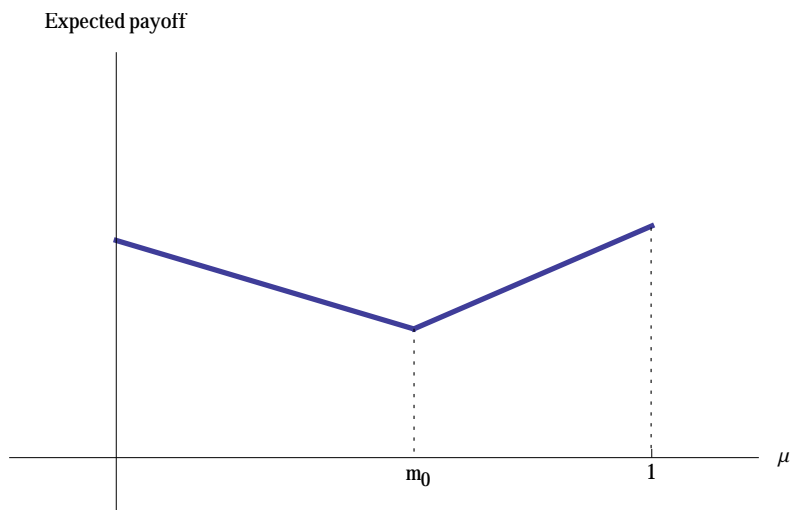


Figure 5: Seller's equilibrium payoff: random proposal model

decrease his probability of making an offer. In the region of high priors, the seller charges the high price, which is only accepted by the high type. Though the occurrence of high types decreases the seller's chances of making an offer, the effect is overpowered by the fact that the seller's price offer is accepted more often. The seller's expected payoff is, hence, increasing in the probability of the buyer's value being high.

The version of the probabilistic bargaining model with different probabilities approximates the model with endogenous proposal rights relatively well and captures the non-monotonicity of the seller's payoff. Of course, it does not match the latter model exactly. The most notable difference is that in the model with endogenous proposal rights, both types make the same investment when they expect the seller to offer the low price. As a consequence, for priors below v_L/v_H , the seller's payoff is constant rather than decreasing in μ . More generally, in the model with investments, the difference between the probabilities of the high and the low type making an offer increases in the prior. This accounts, among other things, for the difference in the seller's payoff below the threshold v_L/v_H . Despite this caveat, we believe that a researcher who wishes to use a more manageable tool will be well suited with the probabilistic model of bargaining proposed here.

6 Concluding Remarks

We introduce a model of bilateral trade in which the right to propose a mechanism is endogenous. The buyer, after learning his value for the object, makes an investment which determines the probability of choosing a mechanism; the seller selects the mechanism with the remaining probability.

We show that buyers with higher valuations invest (weakly) more resources. More

precisely, the investment is flat for very low types, who find it too costly to invest, and very high types, who always trade given the seller's equilibrium pricing strategy. These type have the same gain from obtaining the proposal right, namely, not paying the seller's expected price. The fact that the investment is increasing in the buyer's valuation prompts the seller to revise his beliefs towards lower types upon receiving the opportunity to select a mechanism. This leads to an interesting implication regarding the equilibria of the game: the only possible pure-strategy equilibrium is the one where the seller's learning does not lead him to revise the price, that is, where the seller offers the ex-ante optimal price. If the latter does not constitute an equilibrium, the seller must randomise over prices in equilibrium, with the ex-ante optimal price being the maximum of the support of the seller's randomisation.

We provide a full characterization of equilibria for the environment in which the buyer has two possible values. Interestingly, the seller's payoff is non-monotonic in the share of high-value buyers. Therefore, if the seller were to choose between trading with a buyer that is almost certainly of low value and a buyer who is more likely to be of high value, he might choose the one with the smaller expected value. While high-value buyers offer a higher potential surplus, they also negotiate the terms of trade more aggressively.

Flexibility of our model offers several avenues for future research. In a separate work, we analyse the model where the investments are observable, and thus serve as a signalling device. As a consequence, high investments can become a double-edged sword for the buyer, on the one hand, enabling him to propose the mechanism more often, but on the other, making him seem too eager to trade, thereby revealing his valuation to the seller. Some further directions for research could be to study how the presence of two-sided private information or multiple agents affects the investment and the pricing problem.

A Appendix

Proof of Lemma 1. The first-order condition of (1) is given by

$$\rho'(e) \left(v(1 - H(v)) + \int_0^v p dH(p) \right) \leq k.$$

When this condition holds as an equality for $e > 0$, the solution is interior; otherwise we have a corner solution with $e = 0$. Let \underline{v} be the value of $v \in [0, 1]$ for which the above condition holds as an equality at $e = 0$. Then the optimal investment function is described by $e^*(v) = 0$ for all $v \leq \underline{v}$ and

$$e^*(v) = (\rho')^{-1} \left(\frac{k}{v(1 - H(v)) + \int_0^v p dH(p)} \right)$$

otherwise. Concavity of ρ implies that ρ' and its inverse are decreasing. Together with the fact that $v(1 - H(v)) + \int_0^v p dH(p)$ is non-decreasing in v , it follows that $e^*(v)$ is non-decreasing in v . \square

Proof of Lemma 2. The proof proceeds in two steps: first we argue existence of a Bayesian Nash equilibrium in an auxiliary game, then we show the existence of a perfect Bayesian equilibrium in the original game.

In the auxiliary game we fix the buyer's price offer to $p_b(v) = 0, \forall v \in [0, 1]$ and seller's acceptance strategy to $x_s(p) = 1, \forall p \geq 0$ (accept). Similarly, we set the buyer's acceptance strategy to $x_b(v, p) = 1$ if $p \leq v$ and $x_b(v, p) = 0$ otherwise. The seller's strategy is then the price offer p_s when he is asked to make an offer, and the buyer's strategy is the investment choice e . Given a pair (e, p_s) the buyer's and seller's payoffs are, respectively

$$\rho(e)v + (1 - \rho(e))1_{v \geq p_s}(v - p_s) - ke,$$

and

$$(1 - \rho(e))1_{v \geq p_s}p_s.$$

In what follows we verify conditions needed for Theorem 1 in [Carbonell-Nicolau and McLean \(2017\)](#); which itself builds on existence results in [Reny \(1999\)](#) and [Monteiro and Page Jr \(2007\)](#). Uniform payoff security of the game is implied by the continuity of the seller's payoff in e and the buyer's payoff in p . In addition, the sum of the two players' payoffs, $\rho(e)v + (1 - \rho(e))1_{v \geq p_s}v - ke$, is upper semicontinuous for every v in the profile of strategies (e, p_s) due to the assumption that the buyer buys the good when indifferent. Finally, since only the buyer is privately informed, the absolute continuity of the distribution is automatically satisfied. The only remaining requirement is compactness of the strategy sets. The seller's optimal price is never above 1, therefore we can restrict his set of prices to $[0, 1]$. Likewise, we confine the buyer's investment to the set $[0, 1/k]$; provided that prices are non-negative, the most the buyer can gain from making an offer himself is his highest possible valuation: 1. Restricting e to $[0, 1/k]$ and p_s to $[0, 1]$ renders the strategy sets compact. Theorem 1 of [Carbonell-Nicolau and McLean \(2017\)](#), then, delivers the existence of equilibrium.

Having established the existence of an equilibrium of the auxiliary game, we return to the original game. The seller's strategy consists of the price he charges when given the opportunity and the acceptance decision for any price he may face. The buyer's strategy consists of the investment, the price he charges if he gets to make an offer and the acceptance decision when the seller is offering the price. The seller's assumed acceptance strategy in the continuation game where he faces the buyer's price offer is clearly optimal. The seller's strategy is, therefore, a best response to the buyer's behavior. Likewise, one can argue that if the buyer had a profitable deviation, then he would have a profitable

deviation where he plays the assumed acceptance strategy and offers price 0 when called upon. Such a profitable deviation, as we know, does not exist. \square

Proof of Proposition 1. Fix an equilibrium and let the seller's randomisation over prices be described by the cumulative distribution function H . Let \bar{p} denote the maximum of the support of H . We want to argue that $\bar{p} = p^*$.

In equilibrium the probability with which buyer type v gets to make an offer (see Lemma 1) is

$$\rho(e^*(v; H)) = \rho \left(\max \left\{ 0, (\rho')^{-1} \left(\frac{k}{v(1 - H(v)) + \int_0^v p dH(p)} \right) \right\} \right).$$

Given this probability, the seller's posterior when making an offer is

$$G(v; e^*(\cdot; H)) = \frac{\int_0^v (1 - \rho(e^*(t; H))) dF(t)}{\int_0^1 (1 - \rho(e^*(t; H))) dF(t)}.$$

Any price in the support of H must maximise

$$p(1 - G(p; e^*(\cdot; H))) = p \frac{\int_p^1 (1 - \rho(e^*(t; H))) dF(t)}{\int_0^1 (1 - \rho(e^*(t; H))) dF(t)},$$

and therefore $p \int_p^1 (1 - \rho(e^*(t; H))) dF(t)$. By Lemma 1, $e^*(v; H)$ is non-decreasing in v , which means that $\rho(e^*(v; H))$ is non-decreasing in v , and therefore that $1 - \rho(e^*(v; H))$ is non-increasing in v . Towards a contradiction, suppose now $\bar{p} \neq p^*$. Then

$$\begin{aligned} p^* \int_{p^*}^1 (1 - \rho(e^*(t; H))) dF(t) &\geq p^* \int_{p^*}^1 (1 - \rho(e^*(\bar{p}; H))) dF(x), \\ &= (1 - \rho(e^*(\bar{p}; H))) p^* (1 - F(p^*)), \\ &> (1 - \rho(e^*(\bar{p}; H))) \bar{p} (1 - F(\bar{p})), \\ &= \bar{p} \int_{\bar{p}}^1 (1 - \rho(e^*(t; H))) dF(t), \end{aligned}$$

where the first inequality is due to the fact that $1 - \rho(e^*(v; H))$ is non-increasing in v and constant above \bar{p} , and the second is due to the ex-ante optimality of p^* . Given the inequalities, the price \bar{p} cannot maximise $p(1 - G(p; e^*(\cdot; H)))$ and, hence, cannot belong to the support of H . This proves that $\bar{p} = p^*$ holds.

Since p^* always belongs to the support of H , the only possible pure-strategy equilibrium is the one where the seller offers p^* . \square

Proof of Proposition 2. The low-type buyer solves the problem

$$\max_{e_b} \rho(e_b)v_L - ke_b.$$

Let e_L^* denote the solution of this problem. The assumption $k \leq \rho'(0)v_L$ assures that the solution is uniquely determined by the first-order condition $\rho'(e_L^*)v_L = k$. The high-type buyer solves the problem

$$\max_{e_b} v_H - (1 - \rho(e_b))[\sigma v_H + (1 - \sigma)v_L] - ke_b.$$

The solution of this problem depends on the seller's strategy σ and will be denoted by $e_H^*(\sigma)$, implicitly defined by the first-order condition $\rho'(e_H^*(\sigma))[\sigma v_H + (1 - \sigma)v_L] = k$. Notice that $e_H^*(\sigma)$ is strictly increasing on $[0, 1]$ with $e_H^*(0) = e_L^*$.

We will characterise the equilibrium by distinguishing three cases:

- $\sigma = 0$: suppose in equilibrium the seller offers price v_L with probability one. In this case the investment choice of both buyers is the same, so the seller does not learn in equilibrium. Price v_L is then optimal if and only if $\mu \leq v_L/v_H$. Hence, if and only if $\mu \leq v_L/v_H$, there exists a pure-strategy equilibrium where the seller offer the low price v_L .
- $\sigma = 1$: suppose that in equilibrium the seller offers price v_H with probability one. We have $e_H^*(1) > e_L^*$, hence the high-type buyer invests strictly more than the low-type buyer. The seller's posterior when he makes the offer is

$$\hat{\mu}|_{\sigma=1} = \frac{\mu(1 - \rho(e_H^*(1)))}{\mu(1 - \rho(e_H^*(1))) + (1 - \mu)(1 - \rho(e_L^*))}.$$

The posterior $\hat{\mu}|_{\sigma=1}$ is strictly increasing in μ . The parameter m is defined as the value of μ at which $\hat{\mu}|_{\sigma=1} = v_L/v_H$. Hence, $\hat{\mu}|_{\sigma=1}$ is greater than v_L/v_H if and only if $\mu \geq m$. It follows that there exists an equilibrium where the seller offers the high price v_H with probability one if and only if $\mu \geq m$. By $\rho(e_L^*) < \rho(e_H^*(1))$, we have $m > v_L/v_H$.

- $\sigma \in (0, 1)$: suppose now the seller randomises across prices in equilibrium. This requires that the high-type buyer's investment, $e_H^*(\sigma)$, is such that after updating the seller is indifferent between both prices. That is:

$$\frac{\mu(1 - \rho(e_H^*(\sigma)))}{\mu(1 - \rho(e_H^*(\sigma))) + (1 - \mu)(1 - \rho(e_L^*))} = v_L/v_H \quad (6)$$

Since $e_H^*(\sigma)$ is continuous and strictly increasing on $[0, 1]$, the term on the left-hand side of (6) is continuous and strictly decreasing on $[0, 1]$. At $\sigma = 0$, the left hand side is equal to μ . It follows that (6) is solved by some $\sigma \in (0, 1)$ if and only if $\mu \in (v_L/v_H, m)$, in which case the solution is unique. Hence, there

exists an equilibrium where the seller randomises over both prices if and only if $\mu \in (v_L/v_H, m)$. Substituting for $e_H^*(\sigma)$ and e_L^* in (6) yields expression (4), which implicitly defines the seller's equilibrium randomisation, σ^* .

□

Proof of Proposition 3. Given the buyer's equilibrium investment strategy, as described in Proposition 2, the seller's payoff is given by the product of the probability with which he makes an offer and his conditional payoff in the event he does. His expected payoff as a function of the prior belief μ can thus be written as

$$u_s(\mu) = \begin{cases} (1 - \rho(e_L^*))v_L & \text{if } \mu < \frac{v_L}{v_H} \\ [\mu(1 - \rho(e_H^*(\sigma^*))) + (1 - \mu)(1 - \rho(e_L^*))]v_L & \text{if } \mu \in (\frac{v_L}{v_H}, m) \\ \mu(1 - \rho(e_H^*(1)))v_H & \text{if } \mu \geq m. \end{cases}$$

where σ^* is defined by (4).

- For $\mu \leq v_L/v_H$, the payoff function u_s does not depend on μ and hence we have $u'_s(\mu) = 0$.
- For $\mu \in (v_L/v_H, m)$, taking the first derivative of the seller's payoff with respect to μ yields

$$u'_s(\mu) = \left[-(\rho(e_H^*(\sigma^*)) - \rho(e_L^*)) - \mu\rho'(e_H^*(\sigma^*))e_H^*(\sigma^*)\frac{d\sigma^*}{d\mu} \right] v_L.$$

The first term in the square bracket is strictly negative. Since ρ and e_H^* are increasing functions, the second term is negative if $d\sigma^*/d\mu$ is positive. Given that the seller's posterior on the left-hand side of (4) is increasing in μ and decreasing in σ , $d\sigma^*/d\mu > 0$ is indeed satisfied. Hence, the second term in the square bracket is negative as well and the seller's expected payoff is strictly decreasing in μ in the parameter region $\mu \in (v_L/v_H, \hat{m})$.

- Finally, for $\mu \geq m$ we have $u'_s(\mu) = (1 - \rho(e_H^*(1)))v_H$, which is strictly positive.

□

References

- ALI, S. N. (2015): "Recognition for sale," *Journal of Economic Theory*, 155, 16–29.
- BOARD, S. AND J. ZWIEBEL (2012): "Endogenous competitive bargaining," Tech. rep., mimeo.

- CARBONELL-NICOLAU, O. AND R. P. MCLEAN (2017): “On the existence of Nash equilibrium in Bayesian games,” *Mathematics of Operations Research*, 43, 100–129.
- CIALDINI, R. B. AND N. GARDE (1987): *Influence*, vol. 3, A. Michel.
- CRAWFORD, V. P. (1979): “A procedure for generating Pareto-efficient egalitarian-equivalent allocations,” *Econometrica: Journal of the Econometric Society*, 49–60.
- EVANS, R. (1997): “Coalitional bargaining with competition to make offers,” *Games and Economic Behavior*, 19, 211–220.
- FISHER, R., W. L. URY, AND B. PATTON (2011): *Getting to yes: Negotiating agreement without giving in*, Penguin.
- HAGEDORN, M. AND I. MANOVSKII (2008): “The cyclical behavior of equilibrium unemployment and vacancies revisited,” *American Economic Review*, 98, 1692–1706.
- HEINSALU, S. (2017): “Investing to Access an Adverse Selection Market,” .
- INDERST, R. (2001): “Screening in a matching market,” *The Review of Economic Studies*, 68, 849–868.
- KARAGOZOGLU, E. AND S. RACHMILEVITCH (2018): “Costly Preparations in Bargaining,” Working Paper.
- 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.
- LAUERMANN, S. AND A. WOLINSKY (2017): “Bidder solicitation, adverse selection, and the failure of competition,” *American Economic Review*, 107, 1399–1429.
- MCAFEE, R. P. (1993): “Mechanism design by competing sellers,” *Econometrica: Journal of the econometric society*, 1281–1312.
- MONTEIRO, P. K. AND F. H. PAGE JR (2007): “Uniform payoff security and Nash equilibrium in compact games,” *Journal of Economic Theory*, 134, 566–575.
- MÜNSTER, J. AND M. REISINGER (2015): “Sequencing bilateral negotiations with externalities,” .
- MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- PAI, M. M. (2010): “Competition in mechanisms,” *ACM SIGecom Exchanges*, 9, 7.
- PÉREZ-CASTRILLO, D. AND D. WETTSTEIN (2001): “Bidding for the surplus: a non-cooperative approach to the Shapley value,” *Journal of Economic Theory*, 100, 274–294.

- PETERS, M. AND S. SEVERINOV (1997): “Competition among sellers who offer auctions instead of prices,” *Journal of Economic Theory*, 75, 141–179.
- RENY, P. J. (1999): “On the existence of pure and mixed strategy Nash equilibria in discontinuous games,” *Econometrica*, 67, 1029–1056.
- RUBINSTEIN, A. AND A. WOLINSKY (1985): “Equilibrium in a market with sequential bargaining,” *Econometrica: Journal of the Econometric Society*, 1133–1150.
- SMITH, A. (1776): *An Inquiry into the Nature and Causes of the Wealth of Nations*, (1776).
- TULLOCK, G. (1980): *Efficient rent reeking*, TX: Texas A&M University Press, 97–112.
- 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.