The Bilateral Trade Problem

The Case of Correlated Values

Daniel Motoc [1]

Department of Economics
University of Michigan
04/09/2021

[1] I am beyond grateful to Professor Tilman Börgers for his advice and guidance on this paper. Without his investment and encouragement in my research, this paper would have not been possible. More than anything, I am grateful for how much he has deepened my knowledge as well as interest in economics.

Abstract

Understanding when trade can be expected to occur, and at what prices, is crucial for economists and policymakers who are interested in market design. Therefore, the purpose of this paper is to expand on this research and analyze when two agents with varying amounts of information can be expected to agree to a trade (also known as the Bilateral Trade Problem). I first analyze market equilibria in which the seller observes his value for the good but the buyer does not, similar to the model proposed in Akerlof (1970), but the buyer's and seller's values are correlated. I then propose a model with a two sided information asymmetry in which the buyer and the seller each observe a private signal about the good. Each agent cares not only about his private signal, but about the other's signal when determining how much he values the good (giving rise to correlated values). Finally, throughout the paper I analyze two mechanisms - a fixed price mechanism and a mechanism where the buyer can set the price - and discuss the efficiency of these mechanisms when values are correlated.

1 Introduction

Many real world markets are characterized by information asymmetries between agents, and understanding the effect these asymmetries have on market outcomes and efficiency is an important question for economists and policymakers alike. Consider the healthcare market - doctors know more about treatment than patients, patients know more about their true condition than doctors and insurers, and insurers know more about the provisions of their insurance plans than either doctors or patients. Another salient example is the market for auto services. In most cases, the employees at the auto repair shop know significantly more about the services being offered, their prices, and what the car owner really needs than the owner himself. The predictions of classical models of perfect competition and rational behavior give little insight into how such markets will actually operate, and understanding the role imperfect information plays in economic outcomes has become an important area of research.

In 1945, Friedrich Hayek of the Austrian School of Economics famously proposed that market prices reveal information about the underlying conditions of supply and demand. In particular, he believed that the price system was the most effective mechanism to coordinate the allocation of scarce resources because of the information conveyed to buyers and sellers by market prices. Since then, economists have become increasingly interested in what information is contained in prices, as this is a crucial thing to understand in designing effective markets and trading platforms. To motivate this further, consider the following scenario. Assume you are looking to buy a new home, and after your real estate agent gives you a tour of a beautiful new house, you think to yourself that you would purchase the home for $\$400,000$. But at the end of the tour, the real estate agent mentions that the list price is $\$200,000$. The first question you would ask yourself is, why? What does the market know that I do not? Perhaps there are no good schools nearby, or there is a higher than average crime rate in the neighborhood. Most rational consumers would lower their valuation for the house upon seeing the list price, because the market price has conveyed some information about the property.

Another famous example of this is the Efficient Market Hypothesis, which posits that current stock prices incorporate all existing public knowledge about the stock, and therefore the only way to earn a riskless profit is to trade on private information. Whenever you purchase a stock, you are making a bet that the stock will rise in value. But the counter-party to this trade is making the opposite bet: that the stock will decline in value (ignore for a moment the possibility that they are selling to meet liquidity constraints). A rational agent will then ask: why does the seller think the stock will go down? What does he know that I do not? Even more, what does the market know that I do not?

Much of our knowledge about market design and trading platforms works under the assumption that the buyer's and seller's values for the good are independent. But in order to develop more efficient mechanisms, it is important to understand how they work when values are correlated, i.e. when a buyer and a seller care about the information held by the other agent. Therefore, in this paper I analyze market equilibria in games where the buyer and seller have varying amounts of information regarding the good being transacted and care about the information held by the other agent. I focus on what equilibrium strategies look like for the buyer and seller, when and at what price we can expect trade to occur, and how efficient the mechanisms imposed are. Throughout the paper, I focus on two mechanisms: the fixed price mechanism (where a third party, say the government, imposes a fixed price for the good) and a mechanism in which the buyer can set the price.

There are several notable results in this paper. First, in section 2 I present a model with interdependent values and discrete types. In this model, the correlation between the buyer's and seller's values is not strictly responsible for market outcomes. In particular, the correlation does not impact the expected gains from trade or the likelihood of trade occurring. Second, in section 3 I present a similar model with continuous types. In this model, the likelihood of trade occurring does depend on the correlation, and the price the buyer sets is a function of the correlation. However, the effect of changing the correlation depends on other exogenous parameters. In particular, the price the buyer sets may be either increasing or decreasing in the correlation, and the same is true of the threshold beyond which the buyer will agree to trade. In this model, we can construct different markets in which the same correlation leads to different economic outcomes. However, this model does share the same insight as Akerlof's *The Market for Lemons*: all else equal, as the correlation between the buyer's and seller's values increases, the likelihood of trade decreases. Finally, in section 4 I present a model with a double sided information asymmetry in which both the buyer and seller care about the private information held by the other. This model denotes by α the weight the seller attributes to the buyer's information, and by β the weight the buyer attributes to the seller's information. The main result is that fixed price mechanisms may outperform mechanisms in which the buyer sets the price in high α contexts. In such markets (for example, insurance markets), the marginal cost to the buyer of revealing private information is high, because the seller attributes a large weight to this information. By setting a price, the buyer reveals something about his private information, which the seller may then use to raise his valuation for the good. In such contexts, fixed price mechanisms may be more efficient because neither agent can infer any information from the market price.

1.1 Literature Review

My paper contributes to several strands of literature. First, the models considered in this paper are related to the literature on information economics. For example, the models presented in sections 2 and 3 are similar to that in *Akerlof (1970)*, which analyzes a fictional used car market with many buyers and sellers. The defining feature of this market is that there is uncertainty about the quality of the vehicles being sold - they may either be a *lemon* (low quality) or a *peach* (high quality). The seller knows the quality of the vehicle being sold, but the buyer does not. Akerlof shows that no rational agent will agree to pay the full value of a peach for a car, since there is a chance the car he receives will be a lemon. Thus, the value of peaches exceeds the prevailing equilibrium price in this market, and peaches are driven out. The process continues until all peaches are pushed out and only lemons remain - that is, the information asymmetry degraded the average quality of cars in the market. To take this line of reasoning one step further, rational buyers should be able to predict this process and know that only lemons remain in this market, and so no transactions will occur. This is the so called no-trade theorem. The models in sections 2 and 3 do not consider the quality of the good being sold, but they are similar in that the seller observes his value of the good while the buyer does not.

My paper is also related to the literature on mechanism design. *Milgrom and Stokey (1980)* considers a more complex version of the no-trade theorem. They analyze the effect private market signals would have on trade in a market with many buyers and sellers. They show that in a dynamic rational expectations framework, risk-averse traders who receive private market signals can still not agree to a non-null trade. In dynamic rational expectation models, where the current price vector is a result of voluntary trading in a complete and competitive market, the equilibrium prices must be Pareto-optimal, and changes in prices must only be a result of new information becoming available. Therefore, the main reason that no trade can be expected to occur even in the presence of private information is that it is common knowledge that the equilibrium allocation is both feasible and optimal for all other agents. Therefore, any willingness to accept a trade signals to one agent that the other agent knows something he does not. Thus, once again, in this model where all agents receive a private signal, rational agents will still never agree to a trade. My paper is similar to this paper to the extent that it explores the no trade theorem under specific mechanisms and analyzes when trade can be expected to occur in the presence of information asymmetries.

My paper also contributes to the literature of mechanism design that specifically studies the bilateral trade problem. The bilateral trade problem analyzes a buyer and a seller of a single indivisible good where both agents receive private signals about the value of the good. The question this literature tries to answer is the following: what mechanism, or what set of "rules", is optimal to impose on the buyer and seller in order to promote certain outcomes, such as efficiency. A lot of research has already been done to understand the case of independent private values (that is, the buyer and seller have independent, private valuations for the good). *Myerson and Satterthwaite (1983)* proved one of the most well known results in this field. They show that if the buyer and the seller have independent private values for the good, it is impossible to create a mechanism in which three properties hold: efficiency (whoever values the good more will always end up with it), individual rationality (that both player's expected gain from trade is nonnegative), and incentive compatibility (that truthfully reporting one's own valuation for the good is a dominant strategy). The implication is that the mechanism designer would necessarily sacrifice efficiency if he wants to ensure both the buyer and the seller are better off participating in the mechanism and that they prefer telling the truth to lying.

Much of the follow-up literature to the Myerson-Satterthwaite paper attempts to investigate what mechanisms would be "second-best" - that is, if it is impossible to create a mechanism in which those three properties hold, what is the next best thing? *Hagerty and Rogerson (1987)* proved that in the private value case, the fixed price mechanism and the mechanism in which one side can adjust the price are the only dominant strategy mechanisms (this is the motivation for studying these mechanisms). Furthermore, *Börgers and Li (2019)* defines a new class of mechanisms - called "strategically simple" mechanisms - which do not require a high level of strategic sophistication (the agents do not need to form more than first order beliefs). Thus, the main motivation for studying the two mechanisms studied in this paper is that in the independent private value environment, these mechanisms are robust and simple. Additionally, mechanisms in which one side can change the price are unambiguously more efficient than fixed price mechanisms when values are independent. The question is, is this also true of the interdependent value case?

My paper is therefore also related to the literature on mechanism design with interdependent values. That is, the buyer and seller's valuations for the good depend not only on the private signal they observe, but on the signal the other observes. *Cremer and McLean (1985)* and *Cremer and McLean (1988)* pose a puzzle: they prove that when values are interdependent, all outcomes, including efficient outcomes, can be implemented by efficient trading institutions. However, these mechanisms are neither robust nor simple. Thus, the main purpose for identifying the class of robust and simple mechanisms above was to solve this puzzle. The models in this paper are part of the interdependent value environment, however I simply analyze equilibria and do not ask whether these mechanisms are robust or simple.

The rest of the paper is organized as follows: section 2 studies the two mechanisms under a discrete probability distribution, section 3 develops the model further to allow for a continuous distribution of types, section 4 analyzes a more general model with a double sided information asymmetry, section 5 discusses efficiency of the two mechanisms, section 6 is a discussion of the results, and I conclude in section 7.

2 A Model with an Uninformed Buyer and Discrete Types

2.1 Fixed Price Mechanism

Assume there are two agents who are looking to transact a particular asset at a fixed exogenous price π. Each agent assigns one of two values to the asset, one high and one low. Let F be the probability distribution of asset values among the agents, and assume it is common knowledge. F can be summarized by the following table:

	v_b^L	v_b^H
v_s^L	p	q
v_s^H	r	s

where $p+q+r+s=1$. The buyer and seller must either say yes or no to the transaction, i.e. they choose an action $a_i \in \{Y, N\}$. Furthermore, assume that the seller observes v_s but the buyer does not observe v_b. The seller's strategy is a function $s_s : v_s \to \{Y, N\}$ and the buyer's strategy is simply to say yes or no to trade. Assume the buyer has utility

$$U_b(a_1, a_2, v_b) = \begin{cases} 0 & \text{if } a_1 \text{ or } a_2 = N \\ v_b - \pi & \text{if } a_1 = a_2 = Y \end{cases}$$

and the seller has utility

$$U_s(a_1, a_2, v_s) = \begin{cases} 0 & \text{if } a_1 \text{ or } a_2 = N \\ \pi - v_s & \text{if } a_1 = a_2 = Y \end{cases}$$

Definition 1 (Bayesian Nash Equilibrium)**.** *(s_b^*, s_s^*) form a Bayesian Nash Equilibrium if:*

$$1) \quad \sum_{v_b, v_s} U_b(s_b^*, s_s^*, v_b, v_s) p(v_b, v_s) \geq \sum_{v_b, v_s} U_b(a_b, s_s^*, v_b, v_s) p(v_b, v_s)$$

$$2) \quad \sum_{v_b} U_s(s_b^*, s_s^*, v_b, v_s) p(v_b | v_s) \geq \sum_{v_b} U_s(s_b^*, s_s, v_b, v_s) p(v_b | v_s)$$

for all $a_b \in \{Y, N\}$ and strategies s_s.

Note first that this game has infinitely many Bayesian Nash Equilibria. In particular, when $s_b^* = N$, the seller's expected utility will be 0 no matter what strategy he chooses. Thus, I will focus only on Bayesian Nash Equilibria where the buyer agrees to trade. Another trivial case is when $s_b^* = Y$ and $\pi \geq v_s^H$, as in this case both the buyer and seller will always agree to trade. Therefore, I will restrict attention to the case where the buyer says yes to trade and $v_s^L \leq \pi \leq v_s^H$. Consider the following strategies:

$$s_s^* = \begin{cases} Y & \text{if } \pi \geq v_s \\ N & \text{otherwise} \end{cases}$$

and

$$s_b^* = Y$$

That is, the seller's strategy is to say yes as long as the price is above her observed value, and the buyer says yes to trade.

Theorem 1. *Let $v_s^L \leq \pi \leq v_s^H$. (s_b^*, s_s^*) form a Bayesian Nash Equilibrium if and only if $\frac{p}{q} \geq \frac{\pi - v_b^H}{v_b^L - \pi}$*

Proof. First assume that (s_b^*, s_s^*) form a a Bayesian Nash Equilibrium. From Definition 1, we have:

$$\sum_{v_b, v_s} U_b(s_b^*, s_s^*, v_b, v_s) p(v_b, v_s) \geq 0$$

$$\implies U_b(s_b^*, s_s^*, v_b^L, v_s^L)p + U_b(s_b^*, s_s^*, v_b^L, v_s^H)r + U_b(s_b^*, s_s^*, v_b^H, v_s^L)q + U_b(s_b^*, s_s^*, v_b^H, v_s^H)s \geq 0$$

$$\implies p(v_b^L - \pi) + q(v_b^H - \pi) \geq 0$$

$$\implies \frac{p}{q} \geq \frac{\pi - v_b^H}{v_b^L - \pi}$$

Now assume that $\frac{p}{q} \geq \frac{\pi - v_b^H}{v_b^L - \pi}$. It was shown above that this implies the buyer's expected utility is non-negative. Before calculating the seller's expected utility, we can normalize the above probability distribution to give us the conditional probability distribution. We have:

$$p(v_b^L | v_s^L) = \frac{\frac{p}{q}}{\frac{p}{q} + 1} \qquad p(v_b^H | v_s^L) = \frac{1}{\frac{p}{q} + 1}$$

The seller's expected utility is given by:

$$\sum_{v_b} U_s(s_b^*, s_s^*, v_b, v_s) p(v_b | v_s)$$

$$= U_s(s_b^*, s_s^*, v_b^L, v_s^L)p(v_b^L|v_s^L) + U_s(s_b^*, s_s^*, v_b^H, v_s^L)p(v_b^H|v_s^L)$$

$$= (\pi - v_s^L)\frac{\frac{p}{q}}{\frac{p}{q}+1} + (\pi - v_s^L)\frac{1}{\frac{p}{q}+1}$$

$$= \pi - v_s^L$$

In order to show that $\sum_{v_b} U_s(s_b^*, s_s^*, v_b, v_s)p(v_b|v_s) \geq \sum_{v_b} U_s(s_b^*, s_s, v_b, v_s)p(v_b|v_s)$ for all other strategies s_s, note that the seller's strategy is a threshold strategy. That is, the seller will say yes to trade if $v_s \leq c$ for some price c. If $c < v_s^H$, then the seller will say yes only if $v_s = v_s^L$, and so $\sum_{v_b} U_s(s_b^*, s_s^*, v_b, v_s)p(v_b|v_s) = \pi - v_s^L$ as above. If $c = v_s^H$, we have:

$$\sum_{v_b} U_s(s_b^*, s_s^*, v_b, v_s)p(v_b|v_s)$$

$$= (\pi - v_s^L)(p+q) + (\pi - v_s^H)(r+s)$$

$$= \pi - v_s^L(p+q) - v_s^H(r+s)$$

$$\leq \pi - v_s^L(p+q) - v_s^L(r+s)$$

$$= \pi - v_s^L$$

Therefore, $\sum_{v_b} U_s(s_b^*, s_s^*, v_b, v_s)p(v_b|v_s) \geq \sum_{v_b} U_s(s_b^*, s_s, v_b, v_s)p(v_b|v_s)$ and (s_b^*, s_s^*) constitute a Bayesian Nash Equilibrium. □

2.2 Discussion

The intuition for the condition in theorem 1 is simple. It says that, for a low value of v_s, as long as the probability that the buyer's value is low (p) is sufficiently large relative to the probability that his value is high (q), the buyer will engage in trade. Also note that, if we were to shift the range of the buyer's values (v_b^L, v_b^H) up, we would expect that the buyer would be more willing to engage in trade. Indeed, this is consistent with the intuition from theorem 1. When v_b^H and v_b^L increase, the fraction on the right hand side decreases, and a sufficient decrease may lead the buyer to engage in trade he otherwise would have said no to.

It is also important to realize that the condition in theorem 1 does not depend on r or s, but the correlation between v_b and v_s *does.* The correlation function between v_b and v_s is too complicated to present here for arbitrary v_b and v_s. Therefore, assume $v_b^L = v_s^L = 0$ and $v_b^H = v_s^H = 1$. The correlation is given by [1]

$$\rho(q, r, s) = \frac{-s(s+r-1) - q(r+s)}{\sqrt{(q^2 + s(s-1) + q(2s-1))(r^2 + s(s-1) + r(2s-1))}}$$

Let $\frac{p}{q} = 2$. Then, according to condition 1, as long as $\pi \leq \frac{1}{3}$, trade will occur. The following table illustrates the expected gains from trade and the correlation for various values of r and s.

Correlation, Distributions, and Gains from Trade						
p	q	r	s	E[U$_s$]	E[U$_b$]	ρ
0.5	0.25	0.25	0	π	0.25-0.75π	-0.333
0.5	0.25	0	0.25	π	0.25-0.75π	0.577
0.5	0.25	0.2	0.05	π	0.25-0.75π	-0.126
0.5	0.25	0.05	0.2	π	0.25-0.75π	0.406
0.5	0.25	0.15	0.10	π	0.25-0.75π	0.065
0.5	0.25	0.10	0.15	π	0.25-0.75π	0.236

Note that in the table above, changing the values of r and s has no effect on whether or not trade occurs or the expected gains from trade to either player. However, the correlation varies dramatically from case to case. Therefore, this model suggests that the correlation between v_b and v_s has no impact on the effectiveness of the fixed price mechanism. Rather, it is the relative likelihood of the buyer's value being low that is important, and this may correspond with either positive or negative correlations.

2.3 Buyer Sets Price

Now consider a mechanism in which the buyer can set the price. Let the utility functions be the same as in section 2.1. The seller's strategy is still a function $s_s : v_s \to \{Y, N\}$ but the buyer's strategy is now of the form $s_b = \pi \in \mathbb{R}$. As before, there are infinitely many nash equilibria (if the seller's strategy is to say no for all v_s, the buyer will have the same expected utility regardless of his strategy). Thus, I restrict attention to equilibria in which trade occurs.

Consider the following strategies:

$$s_s* = \begin{cases} Y & \text{if} \quad \pi \geq v_s \\ N & \text{otherwise} \end{cases}$$

and

$$s_b* = \pi = v_s^H$$

Theorem 2. *Assume $v_s^L \leq v_b^L p + v_b^H q$. (s_b^*, s_s^*) form a Bayesian Nash Equilibrium if and only if* $v_b^L(p+r) + v_b^H(q+s) - v_s^H \geq (p+q)\{v_b^L p + v_b^H q - v_s^L\}$

Proof. First assume (s_b^*, s_s^*) form a Bayesian Nash Equilibrium. Then we know the expected utility from setting a price $\pi = v_s^H$ must be greater than or equal to that from setting any other price. The expected utility to the buyer conditional on the seller agreeing to trade is given by:

$$E[U_b] = P(v_s \leq \pi)\{E[v_b | v_s \leq \pi] - \pi\} \tag{1}$$

If $\pi \geq v_s^H$, then (1) becomes:

$$E[U_b] = v_b^L(p+r) + v_b^H(q+s) - \pi$$

since $P(v_s \leq \pi) = 1$ and $E[v_b|v_s \leq \pi] = E[v_b]$. This expression is maximized when π takes its lowest possible value, i.e. $\pi = v_s^H$. Now assume $v_s^L \leq \pi < v_s^H$. (1) becomes:

$$E[U_b] = (p+q)\{v_b^L p + v_b^H q - \pi\}$$

This is maximized when $\pi = v_s^L$. Finally, assume $\pi < v_s^L$. In this case, it is clear the buyer's expected utility is zero because the seller will never agree to trade (since v_s is guaranteed to be greater than π). Therefore, if (s_b^*, s_s^*) form a Bayesian Nash Equilibrium, it must be true that:

$$v_b^L(p+r) + v_b^H(q+s) - v_s^H \geq (p+q)\{v_b^L p + v_b^H q - v_s^L\}$$

Now let $v_b^L(p+r) + v_b^H(q+s) - v_s^H \geq (p+q)\{v_b^L p + v_b^H q - v_s^L\}$. Then, from above, we know that the expected utility from setting a price $\pi \geq v_s^H$ is greater than or equal to setting a price $\pi < v_s^H$. Since the buyer's expected utility is maximized when setting the lowest possible price $\pi \geq v_s^H$, he sets a price $\pi = v_s^H$. Finally, the condition at the beginning of theorem 2 guarantees that the expected utility from trading will always be non-negative. □

There is also an equilibrium in which the buyer sets a price $\pi = v_s^L$. Consider the following strategies: Consider the following strategies:

$$s_s* = \begin{cases} Y & \text{if} \quad \pi \geq v_s \\ N & \text{otherwise} \end{cases}$$

and

$$s_b* = \pi = v_s^L$$

Theorem 3. *Assume $v_s^H \leq v_b^L(p+r) + v_b^H(q+s)$. (s_b^*, s_s^*) form a Bayesian Nash Equilibrium if and only if $v_b^L(p+r) + v_b^H(q+s) - v_s^H \leq (p+q)\{v_b^L p + v_b^H q - v_s^L\}$*

Proof. The proof is identical to the proof of theorem 2 except for the fact that the inequality sign in theorem 2 is flipped. Therefore, it is left to the reader to verify this as an equilibrium. □

2.4 Discussion

Although the conditions in theorems 2 and 3 appear complicated, they are quite intuitive. Note that the condition in theorem 2 can be rewritten as:

$$E[v_b] - v_s^H \geq E[U_b|\pi = v_s^L]$$

That is, the uninformed buyer will want to set a price $\pi = v_s^H$ if either his expected valuation for the good is very high or if v_s^H is not that large (i.e. if it does not cost the buyer a lot to guarantee the seller agrees to trade). But if v_s^H is very large or if the buyer's expected value for the good is not that high, he will want to set a price $\pi = v_s^L$ (note that the conditions at the beginning of theorems 2 and 3 ensure the expected utility from doing so will always be nonnegative). Note too that in this model, it is quite easy for inefficient trade to arise. For example, if p and q are zero, theorem 2 tells us that the buyer will set a price $\pi = v_s^H$ if:

$$v_b^L r + v_b^H s \geq v_s^H$$

If $v_b^L < v_s^H$, then we know with probability r inefficient trade will arise.

Finally, we again want to analyze what effect different probability distributions (and therefore correlations) have on expected gains from trade and whether or not trade occurs. For simplicity, assume that $(v_b^L, v_b^H) = (0.5, 2)$ and $(v_s^L, v_s^H) = (0.5, 0.75)$. For all of the following distributions, the buyer prices the good at $v_s^L = 0.50$, i.e. we are in the setting of theorem 3. The following table presents the expected gains from trade for both players along with the corresponding correlations between v_b and v_s.

Correlation, Distributions, and Gains from Trade						
p	q	r	s	E[U$_s$]	E[U$_b$]	ρ
0.5	0.25	0.25	0	0	0.0625	-0.333
0.5	0.25	0	0.25	0	0.0625	0.577
0.5	0.25	0.2	0.05	0	0.0625	-0.126
0.5	0.25	0.05	0.2	0	0.0625	0.406
0.5	0.25	0.15	0.10	0	0.0625	0.065
0.5	0.25	0.10	0.15	0	0.0625	0.236

Once again, notice that the expected gains from trade are independent of the correlation between v_b and v_s. In this model, it is not the correlation that is important in determining whether trade occurs or what the expected gains to trade may be. Rather, it is the seller's expected value for the good and the likelihood that the seller agrees to trade that are important here, and the correlation gives us little insight into either.

3 A Model with an Uninformed Buyer and Continuous Types

In this section, I will build upon the model to allow for a continuous range of asset values for both players. Assume that both the buyer's and the seller's asset values are given by

$$v_b \sim N(\mu_b, \sigma_b), \quad v_s \sim N(\mu_s, \sigma_s)$$

and that the buyer's and seller's values are correlated. Further assume that v_b and v_s are jointly normally distributed and that this distribution is common knowledge. Denote by f_{v_i} and F_{v_i} the probability density function and conditional distribution function of player i's value respectively, and denote by ρ the correlation between v_b and v_s.

3.1 Fixed Price Mechanism

As before, this game has infinitely many Bayesian Nash Equilibria. In particular, $(s_b = N, s_s)$ constitutes a Bayesian Nash Equilibrium for all strategies s_s. Therefore, I will again restrict attention only to equilibria in which trade occurs. Assuming the same equilibrium strategies from section 2.1, we get the following result:

Theorem 4. *(s_s^*, s_b^*) form a Bayesian Nash Equilibrium if and only if $\mu_b - \rho\sigma_b \frac{f_{v_s}(\pi)}{F_{v_s}(\pi)} \geq \pi$. Furthermore, this is the unique equilibrium in which trade occurs.*

Proof. First assume that (s_s^*, s_b^*) form a Bayesian Nash Equilibrium. This requires that the buyer's expected utility be nonnegative:

$$E[U_b] = E[U_b|a_s = Y]P(a_s = Y) = [E[v_b - \pi|a_s = Y]]P(a_s = Y) \geq 0$$

$$\implies P(v_s \leq \pi)[E[v_b|v_s \leq \pi] - \pi] \geq 0$$

$$\implies E[v_b|v_s \leq \pi] - \pi \geq 0$$

$$\implies E[v_b|v_s \leq \pi] \geq \pi$$

$$\implies \frac{1}{P(v_s \leq \pi)} \int_{-\infty}^{\pi} E[v_b|v_s] f_{v_s} dv_s \geq \pi$$

We can find the conditional density $f(v_b|v_s)$ by dividing the joint distribution by the marginal distribution of v_s. We then get that $E[v_b|v_s]$ is given by $\mu_b + \rho\sigma_b(\frac{v_s-\mu_s}{\sigma_s})$

Thus,

$$\implies \int_{-\infty}^{\pi} E[v_b|v_s]f(v_s)dv_s \geq \pi P(v_s \leq \pi)$$

$$\implies \int_{-\infty}^{\pi} (\mu_b + \rho\sigma_b(\frac{v_s - \mu_s}{\sigma_s}))f_{v_s}dv_s \geq \pi F_{v_s}(\pi)$$

$$\implies \int_{-\infty}^{\pi} \mu_b f_{v_s}dv_s + \int_{-\infty}^{\pi} \frac{\rho\sigma_b}{\sigma_s}(v_s - \mu_s)f_{v_s}dv_s \geq \pi F_{v_s}(\pi)$$

$$\implies \mu_b \int_{-\infty}^{\pi} f_{v_s}dv_s + \frac{\rho\sigma_b}{\sigma_s}\int_{-\infty}^{\pi} (v_s - \mu_s)f_{v_s}dv_s \geq \pi F_{v_s}(\pi)$$

$$\implies \mu_b F_{v_s}(\pi) + \frac{\rho\sigma_b}{\sigma_s}[\int_{-\infty}^{\pi} v_s f_{v_s}dv_s - \mu_s\int_{-\infty}^{\pi} f_{v_s}dv_s] \geq \pi F_{v_s}(\pi)$$

$$\implies \mu_b F_{v_s}(\pi) + \frac{\rho\sigma_b}{\sigma_s}\int_{-\infty}^{\pi} v_s f_{v_s}dv_s - \frac{\rho\sigma_b\mu_s}{\sigma_s}F_{v_s}(\pi) \geq \pi F_{v_s}(\pi)$$

$$\implies \mu_b + \frac{\rho\sigma_b}{\sigma_s}\cdot\frac{1}{F_{v_s}(\pi)}\int_{-\infty}^{\pi} v_s f_{v_s}dv_s - \frac{\rho\sigma_b\mu_s}{\sigma_s} \geq \pi$$

Now note that $\frac{1}{F_{v_s}(\pi)}\int_{-\infty}^{\pi} v_s f_{v_s}dv_s = E[v_s|v_s \leq \pi]$, which in the case of a normal distribution can be expressed as $E[v_s|v_s \leq \pi] = \mu_s - \sigma_s\frac{f_{v_s}(\pi)}{F_{v_s}(\pi)}$. Thus, we have

$$\mu_b + \frac{\rho\sigma_b}{\sigma_s}(\mu_s - \sigma_s\frac{f_{v_s}(\pi)}{F_{v_s}(\pi)}) - \frac{\rho\sigma_b\mu_s}{\sigma_s} \geq \pi$$

$$\mu_b + \frac{\rho\sigma_b\mu_s}{\sigma_s} - \rho\sigma_b\frac{f_{v_s}(\pi)}{F_{v_s}(\pi)} - \frac{\rho\sigma_b\mu_s}{\sigma_s} \geq \pi$$

$$\implies \mu_b - \rho\sigma_b\frac{f_{v_s}(\pi)}{F_{v_s}(\pi)} \geq \pi \tag{2}$$

Now consider the reverse direction. If (1) holds, the buyer's expected utility from always agreeing to trade is nonnegative, and therefore saying yes is a weakly dominant strategy. We need to show that $E[U_s(s_b^*, s_s^*, v_b, v_s))] \geq E[U_s(s_b^*, s_s, v_b, v_s)]$ for all strategies s_s. As before, the seller's strategy can be thought of as a threshold strategy in which the seller agrees to trade if $v_s \leq c$ for some c. The expected utility to the seller for threshold c is $F_{v_s}(c)(\pi - v_s)$. If $c < \pi$, $E[U_s(s_b^*, s_s, v_b, v_s)] < E[U_s(s_b^*, s_s^*, v_b, v_s)]$ because F_{v_s} is an increasing function. If $c > \pi$, we have

$$E[U_s(s_b^*, s_s, v_b, v_s)] = F_{v_s}(c)(\pi - v_s)$$

$$= F_{v_s}(\pi)(\pi - v_s) + (\pi - v_s)(F_{v_s}(c) - F_{v_s}(\pi))$$

$$< F_{v_s}(\pi)(\pi - v_s) = E[U_s(s_b^*, s_s^*, v_b, v_s)]$$

since $\pi - v_s < 0$ when $\pi < v_s < c$. Thus, $E[U_s(s_b^*, s_s^*, v_b, v_s)] \geq E[U_s(s_b^*, s_s, v_b, v_s)]$ for all strategies s_s, and (s_b^*, s_s^*) is a Bayesian Nash Equilibrium. Furthermore, this equilibrium is unique, since the seller's expected utility is maximized when he chooses a threshold $c = \pi$. $\square$

3.2 Analysis

Condition (2) can be thought of as a threshold for when the buyer will say yes to trade. Specifically, he says yes to trade as long as $\mu_b \geq \pi + \rho\sigma_b \frac{f_{v_s}(\pi)}{F_{v_s}(\pi)}$. When $\rho = 0$, the buyer agrees to trade if $\mu_b \geq \pi$. This makes intuitive sense - the buyer agrees to trade as long as on average, his value is above the price. When the correlation increases, this strategy may no longer be optimal. If $\rho\sigma_b \frac{f_{v_s}(\pi)}{F_{v_s}(\pi)}$ is sufficiently large that the inequality no longer holds, it might become optimal for the buyer to say no to trade. First note that this term is increasing in ρ. The buyer can use the seller's willingness to trade as a signal of the seller's value, and because values are correlated, this tells him something about his own value. Specifically, when the correlation is positive, the buyer infers from a trade agreement that the seller's value is relatively low, and therefore that his own value is low as well. Because there is a higher probability that his own value is low, the threshold on μ_b for which the buyer agrees to trade goes up. Similarly, for negative correlations, the buyer infers from a trade agreement that the seller's value is relatively low, and therefore his own value is relatively *high*. Because there is a higher probability that his own value is high, the threshold on μ_b for which it is optimal for the buyer to trade decreases.

Consider now the impact of a change in the variance of v_b on the buyer's willingness to trade. When $\rho = 0$, changing σ_b has no impact on the buyer's willingness to trade. This is because the investor is risk-neutral - the buyer will always agree to trade if on average, his value is above the price. When $\rho < 0$, increasing σ_b causes the threshold to decrease (i.e. increases the buyer's willingness to trade). Since both agents are assumed to be risk-neutral, this cannot be because the buyer has an intolerance to risk. To understand the intuition, remember that the buyer's expected value for the good conditional on v_s is given by:

$$E[v_b|v_s] = \mu_b + \rho\sigma_b(\frac{v_s - \mu_s}{\sigma_s})$$

and that

$$E[v_s|v_s \leq \pi] = \mu_s - \sigma_s \frac{f_{v_s}(\pi)}{F_{v_s}(\pi)} < \mu_s \quad \text{for all} \quad \pi$$

Thus, the seller infers from trade that the seller's expected value is less than μ_s. Then we know that for $\rho < 0$, increasing σ_b increases the expected value of v_b conditional on the seller agreeing to trade and thus makes him more willing to trade. Similarly, when $\rho > 0$, increasing σ_b will decrease the expected value of v_b conditional on trade, and therefore make the buyer less willing to trade.

This section has two important conclusions: First, the correlation *is* relevant in this model. In particular, the correlation determines whether or not the buyer will say yes to trade. Second, even if agents are risk-neutral, changing the variance of v_b can still impact whether or not the buyer agrees to trade. This is because the buyer's expected value conditional on the seller agreeing to trade is a function of σ_b, and the effect of σ_b on his expected value depends on the correlation. When values are negatively correlated, increasing σ_b makes the buyer more willing to trade because it raises his conditional expected value for the good. On the other hand, when $\rho > 0$, increasing σ_b makes the buyer less willing to trade because it decreases his conditional expected value for the good.

3.3 Buyer Sets Price

Now consider the exact model as in section 3.1 but assume that π is no longer fixed - rather, the buyer may set the price. We now have that the seller's strategy is a function $s_s : (v_s, \pi) \to \{Y, N\}$ as before, but the buyer's strategy is to set a price $\pi \in \mathbb{R}$.

Theorem 5. *The strategies* $s_s^* = \begin{cases} Y & \text{if} \quad \pi \geq v_s \\ N & \text{otherwise} \end{cases}$ *and* $s_b^* = \pi$ *such that* π *solves*

$$\mu_b - \pi - \rho\sigma_b(\frac{\mu_s - \pi}{\sigma_s^2}) = \frac{F_{v_s}(\pi)}{f_{v_s}(\pi)}$$

form a Bayesian Nash Equilibrium.

Proof. As before, as long as $\pi \geq v_s$, the seller's strategy to say yes to trade is weakly dominant. Now consider the buyer. The buyer faces the following problem:

$$\max_\pi E[U_b]$$

$$= \max_\pi E[U_b|v_s \leq \pi]P(v_s \leq \pi)$$

$$= \max_\pi F_{v_s}(\pi)(E[v_b|v_s \leq \pi] - \pi)$$

$$= \max_\pi F_{v_s}(\pi)(\frac{1}{F_{v_s}(\pi)} \int_{-\infty}^{\pi} E[v_b|v_s] f_{v_s} d_{v_s} - \pi)$$

$$= \max_\pi \int_{-\infty}^{\pi} E[v_b|v_s] f_{v_s} d_{v_s} - F_{v_s}(\pi)\pi$$

$$= \max_\pi \int_{-\infty}^{\pi} (\mu_b + \rho\sigma_b(\frac{v_s - \mu_s}{\sigma_s})) f_{v_s} d_{v_s} - F_{v_s}(\pi)\pi$$

This integral was evaluated in section 3, so we have

$$\max_\pi \mu_b F_{v_s}(\pi) + \frac{\rho\sigma_b}{\sigma_s} \int_{-\infty}^{\pi} v_s f_{v_s} dv_s - \frac{\rho\sigma_b\mu_s}{\sigma_s} F_{v_s}(\pi) - F_{v_s}(\pi)\pi$$

$$= \max_\pi F_{v_s}(\pi)[\mu_b + \frac{\rho\sigma_b}{\sigma_s} \cdot \frac{1}{F_{v_s}(\pi)} \int_{-\infty}^{\pi} v_s f_{v_s} dv_s - \frac{\rho\sigma_b\mu_s}{\sigma_s} - \pi]$$

$$= \max_\pi F_{v_s}(\pi)[\mu_b - \rho\sigma_b \frac{f_{v_s}(\pi)}{F_{v_s}(\pi)} - \pi]$$

$$= \max_\pi F_{v_s}(\pi)(\mu_b - \pi) - \rho\sigma_b f_{v_s}(\pi)$$

A quick look at the second-order conditions of this function show that the second derivative is of ambiguous sign. Therefore, I plotted the objective function in Mathematica for various parameter values to ensure we are solving for a maximum. In every case, the objective function was concave on the necessary region. One example is presented below.

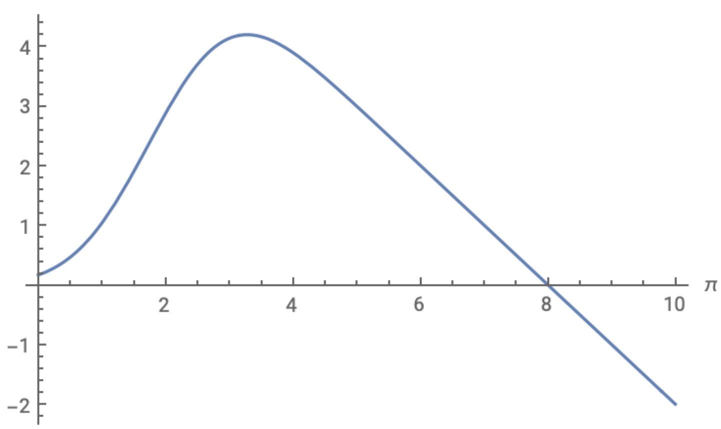

Figure 1: Graph of $F_{v_s}(\pi)(\mu_b - \pi) - \rho\sigma_b f_{v_s}(\pi)$

Thus, we can be reasonably confident that setting the first order condition to zero will solve for a maximum.

We now have that π must satisfy:

$$\frac{\partial}{\partial\pi}[F_{v_s}(\pi)\mu_b - \rho\sigma_b f_{v_s}(\pi) - F_{v_s}(\pi)\pi] = 0$$

$$\implies f_{v_s}(\pi)(\mu_b - \pi) - \rho\sigma_b f'_{v_s}(\pi) - F_{v_s}(\pi) = 0$$
$$\implies f_{v_s}(\pi)(\mu_b - \pi) - \rho\sigma_b f'_{v_s}(\pi) = F_{v_s}(\pi)$$

In the case of a normal distribution,

$$f'_{v_s}(\pi) = f_{v_s}(\pi)(\frac{\mu_s - \pi}{\sigma_s^2})$$

Plugging in and simplifying, we get that π must satisfy

$$\mu_b - \pi - \rho\sigma_b(\frac{\mu_s - \pi}{\sigma_s^2}) = \frac{F_{v_s}(\pi)}{f_{v_s}(\pi)} \tag{3}$$

Although there exists no simple, closed-form solution to (2), I use numerical methods in Mathematica to understand the economic properties behind the price formation.

Lemma 1. *Assume $\mu_b = \mu_s = \mu$. Then if $\rho < \frac{\sigma_s^2}{\sigma_b}$, there exists a unique, positive π that solves (2).*

Proof. Let $g(\pi) := (\mu - \pi)(1 - \frac{\rho\sigma_b}{\sigma_s^2}) - \frac{F_{v_s}(\pi)}{f_{v_s}(\pi)}$. Then equation (2) can be rewritten as $g(\pi) = 0$. To characterize existence of solutions, note first that $g(0) = \mu(1 - \frac{\rho\sigma_b}{\sigma_s^2}) > 0$ and that

$$\lim_{\pi \to \infty} g(\pi) = -\infty$$

By the Intermediate Value Theorem, we know that there exists a $\pi > 0$ where $g(\pi) = 0$. To show this solution is unique, it suffices to show that $g(\pi)$ is one to one. We know

$$\frac{\partial g(\pi)}{\partial \pi} = -(1 - \frac{\rho\sigma_b}{\sigma_s^2}) - \frac{\partial}{\partial \pi}\frac{F_{v_s}(\pi)}{f_{v_s}(\pi)} < 0 \quad \forall\, \pi$$

because the Mill's ratio $\frac{F_{v_s}(\pi)}{f_{v_s}(\pi)}$ is increasing in π. Therefore, $g(\pi)$ is one to one and the price π that solves (2) is unique. □

3.4 Analysis

The price function that satisfies (3) has several noteworthy economic properties. First consider the impact that an increase in the variance of the buyer's value would have on the price. Using Mathematica to numerically estimate the price for varying values of σ_b resulted in the following graphs:

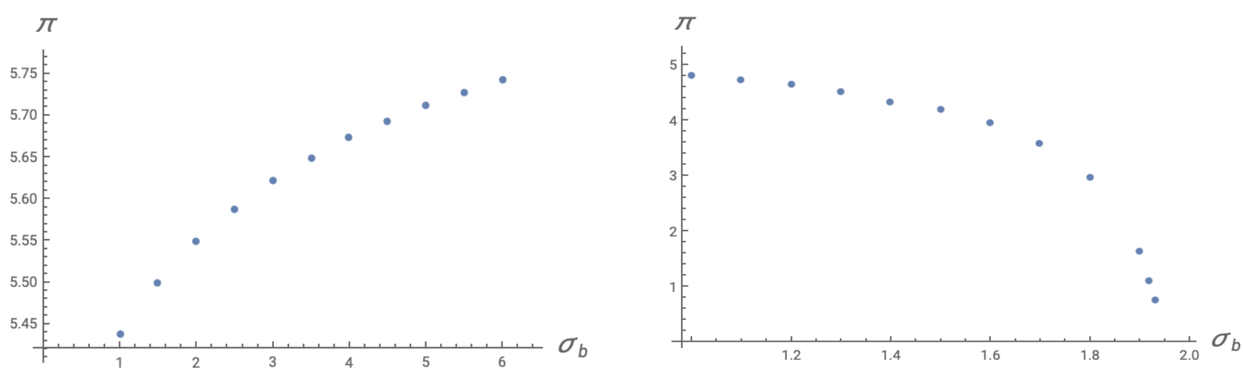

$\rho = -0.5$ $\rho = 0.5$

Figure 2: Price vs σ_b

(a) $v_s \sim N(6, 1)$, $v_b \sim N(6, 1)$

As we saw in section 3.2, the effect of changing σ_b depends on the correlation. From figure 3, we see that $\frac{\partial \pi}{\partial \sigma_b} > 0$ if $\rho < 0$ and $\frac{\partial \pi}{\partial \sigma_b} < 0$ if $\rho > 0$. The intuition is much the same as explained in section 3.2. When $\rho < 0$, raising σ_b increases the conditional expected value of v_b and makes the buyer more willing to trade (and therefore set a higher price). When $\rho > 0$, increasing σ_b decreases the conditional expected value of v_b, making him less willing to trade and therefore causing him to lower the price. Finally, note that when $\rho = 0$, the price function (3) which π solves simplifies to

$$\pi = \mu_b - \frac{F_{v_s}(\pi)}{f_{v_s}(\pi)}$$

There are two things to note here. First, the buyer sets a price by calculating his expected value for the good and subtracting $\frac{F_{v_s}(\pi)}{f_{v_s}(\pi)}$. This term that he subtracts accounts for the probability that the seller will say yes to trade, and it is clear that when the values are uncorrelated, the buyer will never set a price higher than his expected value. Second, this price function does not depend on σ_b. That is, an increase in σ_b now has no impact on the price he will set. This arises because the buyer is primarily concerned with the value $E[v_b|v_s \leq \pi]$. When the values are uncorrelated, this becomes simply $E[v_b] = \mu_2$. The buyer observes whether the seller says yes to trade, and then uses σ_b to infer from this decision something about his own value. When the values are uncorrelated, he can no longer do this.

What can we say about the price the buyer sets as a function of the correlation ρ? That is, what can be said about the sign of $\frac{\partial \pi}{\partial \rho}$? The answer is it depends where along the seller's distribution the buyer is pricing the good. Consider first an example where $\pi < \mu_s$ presented below.

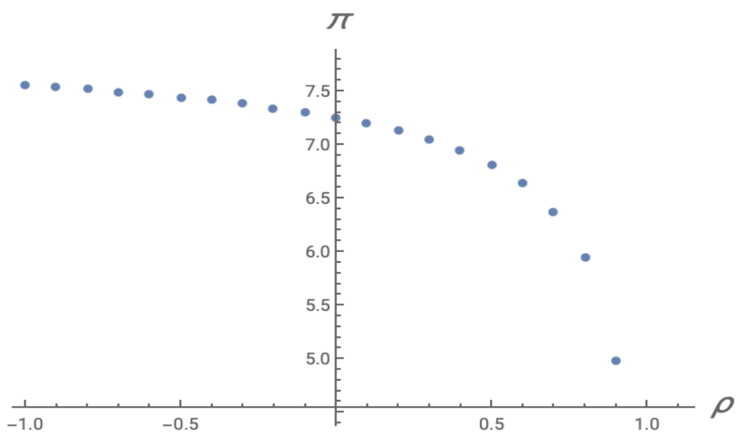

Figure 4: Price vs ρ

(a) $v_s \sim N(8, 1)$, $v_b \sim N(8, 1)$

As the correlation increases, the price the buyer sets decreases, and this negative relationship becomes stronger as v_s and v_b become more positively correlated. The economic intuition is the following: assume v_s and v_b are negatively correlated. Upon seeing that the seller agrees to trade, the buyer infers that the seller's value must be relatively low (lower than the price), and therefore his own value is relatively high. But as the correlation increases, the probability that his own value is low given that the seller's is low increases. To compensate himself for this added risk, the buyer must lower the price as the correlation increases. As the correlation approaches 1, the probability that the buyer's own value is low becomes so large that his willingness to pay for the good drops significantly. Now consider that we lower the seller's expected value for the good such that $\pi > \mu_s$.

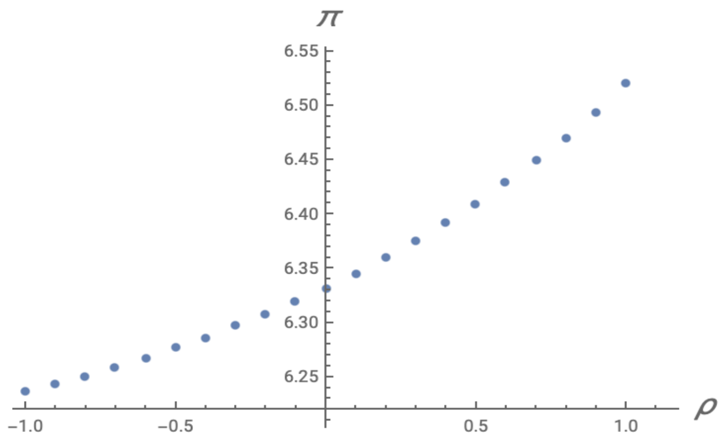

Figure 6: Price vs ρ

(a) $v_s \sim N(6,1)$, $v_b \sim N(8,1)$

In order to understand why the two cases differ, recall that the buyer's expected utility can be written as follows:

$$E[U_b] = F_{v_s}(\pi)[\mu_b - \pi + \rho\frac{\sigma_b}{\sigma_s}(E[v_s|v_s \leq \pi] - \mu_s)]$$

Note that $E[v_s|v_s \leq \pi] = \mu_s - \sigma_s \frac{f_{v_s}(\pi)}{F_{v_s}(\pi)} < \mu_s$ for all π. Therefore, when the correlation increases, we maximize expected utility by making $E[v_s|v_s \leq \pi] - \mu_s$ smaller. To understand this intuitively, consider the case where v_s and v_b are perfectly negatively correlated. The buyer would want to price the good in a range of low values of v_s, because these correlate with high values of v_b and he can benefit from both a higher v_b and a lower π. But if v_s and v_b were perfectly positively correlated, pricing in the range of low v_s would correlate with low values of v_b. Therefore as the correlation increases, the buyer wants to raise the price in order to increase the probability that the seller's value is high given that he agrees to trade, and therefore that his own value is high. We know that $F_{v_s}(\pi)$ is a strictly increasing function, and $f_{v_s}(\pi)$ is a decreasing function when $\pi > \mu_s$. Therefore, when pricing above μ_s, the buyer can increase $E[v_s|v_s \leq \pi]$ by increasing π and therefore ensure that he is pricing the good in a range of higher values of v_s.

3.5 Probability of Trade

Another important question to ask in addition to what price the buyer will set is whether or not the seller will agree to trade. That is, what is the probability that a transaction will occur at this price? We are interested in the following probability:

$$P(v_s \leq \pi)$$

It is clear that whenever the price increases, the probability of trade occurring increases as well (however this is not necessarily the case for changes in μ_s and σ_s, when the distribution of v_s changes as well). For example, consider the effect of an increase in the correlation. The buyer will set a lower price, and therefore the probability of trade goes down. This makes intuitive sense - two agents with positively correlated values for the asset will be less likely to agree to a trade at this price if they have similar values for the asset. The opposite is true of a decrease in ρ. As ρ decreases, the buyer will set a higher price and the probability of trade occurring will increase. If the seller values a good relatively low and the buyer relatively high, then a trade would benefit both parties, which is why we see an increase in the probability of trade.

Once again, the effect of an increase in σ_b is ambiguous. If $\rho < 0$, then the buyer will respond to a higher σ_b by increasing the price, and therefore increasing the probability of trade occurring. Similarly, if $\rho > 0$, the buyer will respond to a higher σ_b by decreasing the price and therefore reducing the probability of trade occurring.

The following table summarizes all the effects described in this section.

Comparative Statics			
	$\boldsymbol{\rho < 0}$	$\boldsymbol{\rho = 0}$	$\boldsymbol{\rho > 0}$
$\frac{\partial \pi}{\partial \mu_s}$	> 0	> 0	> 0
$\frac{\partial \pi}{\partial \mu_b}$	> 0	> 0	> 0
$\frac{\partial \pi}{\partial \sigma_b}$	> 0	$= 0$	< 0
$\frac{\partial \pi}{\partial \rho}$	$\begin{cases} > 0 & \text{if} \quad \pi > \mu_s \\ < 0 & \text{if} \quad \pi < \mu_s \end{cases}$	$\begin{cases} > 0 & \text{if} \quad \pi > \mu_s \\ < 0 & \text{if} \quad \pi < \mu_s \end{cases}$	$\begin{cases} > 0 & \text{if} \quad \pi > \mu_s \\ < 0 & \text{if} \quad \pi < \mu_s \end{cases}$

4 Two Sided Information Asymmetries

Consider the following framework. A seller possesses a single indivisible good which he is looking to trade to a buyer. The seller observes a private signal s_1 and the buyer observes a private signal s_2. (s_1, s_2) have joint distribution F and joint density f, and assume s_1 and s_2 are independent. Denote the seller's value $v_s(s_1, s_2)$ and the buyer's value $v_b(s_1, s_2)$. Finally, assume utility functions of the form:

$$U_s(s_1, s_2, \pi) = \begin{cases} \pi - v_s(s_1, s_2) & \text{if trade occurs} \\ 0 & \text{if trade does not occur} \end{cases}$$

and

$$U_b(s_1, s_2, \pi) = \begin{cases} v_b(s_1, s_2) - \pi & \text{if trade occurs} \\ 0 & \text{if trade does not occur} \end{cases}$$

where π denotes the price at which trade occurs.

4.1 Fixed Price Mechanisms

Let π denote an exogenously fixed price at which players can either agree or not agree to trade. Thus, the set of actions for each player is simply $\Omega = \{Y, N\}$. That is, the players can either say yes or no to trade.

Definition 2. *Let A_1 denote the subset of the domain of s_1 on which $E[U_s(s_1, s_2, \pi)|s_2 \in A_2]$ is positive. Similarly, let A_2 denote the subset of the domain of s_2 on which $E[U_b(s_1, s_2, \pi)|s_1 \in A_1]$ is positive. Consider the following strategies:*

$$s_1^* : \begin{cases} A_1 \to Y \\ A_1^C \to N \end{cases}$$

$$s_2^* : \begin{cases} A_2 \to Y \\ A_2^C \to N \end{cases}$$

Then if

$$\frac{1}{P(A_2)} \int_{A_2} v_s f(s_2|s_1) ds_2 \leq \pi \leq \frac{1}{P(A_1)} \int_{A_1} v_b f(s_1|s_2) ds_1$$

$\{s_1^*, s_2^*\}$ *constitute a Bayesian Nash Equilibrium.*

Proof. Consider first the point of view of the seller. Saying yes is a weakly dominant strategy for the seller if $E[U_s] \geq 0$. We have

$$\begin{aligned} & E[U_s] \geq 0 \\ \implies & E[U_s|\omega_b = Y]P(\omega_b = Y) + E[U_s|\omega_b = N]P(\omega_b = N) \geq 0 \\ \implies & E[U_s|\omega_b = Y] \geq 0 \\ \implies & E[\pi - v_s|s_2 \in A_2] \geq 0 \\ \implies & E[v_s|s_2 \in A_2] \leq \pi \\ \implies & \frac{1}{P(s_2 \in A_2)} E[v_s \cdot \mathbb{1}_{A_2}] \leq \pi \\ \implies & \frac{1}{P(A_2)} \int_{A_2} v_s(s_1, s_2) f(s_2|s_1) ds_2 \leq \pi \end{aligned} \tag{4}$$

An identical argument for the buyer gives the reverse inequality:

$$\frac{1}{P(A_1)} \int_{A_1} v_b(s_1, s_2) f(s_1|s_2) ds_1 \geq \pi \tag{5}$$

Thus, if inequalities (5) and (6) hold, $\{s_1^*, s_2^*\}$ constitute a Bayesian Nash Equilibrium.

□

Proposition 1. *Let $f(\cdot)$ be a bounded and atomless density over s_1 and s_2. Then there exists a pure strategy nash equilibrium in the fixed price mechanism where both the buyer's and seller's strategy are nondecreasing*

Proof. The proof uses theorem 1 in Athey (2001). Athey's result shows that in finite action games, there exists a pure strategy nash equilibrium if whenever one player uses a nondecreasing strategy, the other player's objective function, given by

$$u_i(a_i, s_i; \alpha_{-i}(\cdot)) \equiv \int_{s_{-i}} U_i((a_i, \alpha_{-i}(s_{-i})), s) f(s_{-i}|s_i) ds_{-i}$$

satisfies the Single Crossing Property of Incremental Returns (SCP-IR), which for differentiable objective functions is given by:

$$\frac{\partial^2}{\partial a_i \partial s_i} u_i(a_i, s_i; \alpha_{-i}(\cdot)) \geq 0$$

To clarify notation above, a_i is an element taken from the set of actions, α_i represents player $i's$ strategy, s_i represents player $i's$ type (in this game, his signal), and the subscript $-i$ will refer the corresponding notation for the other player.

Let the support of $f(\cdot)$ be given by $[\underline{s_i}, \overline{s_i}]$ Consider first the buyer. Since U_b is increasing in s_2, we know his strategy will take the following form:

$$\alpha_b = \begin{cases} Y & \text{if} \quad s_2 \geq c \\ N & \text{otherwise} \end{cases}$$

for some $c \in \mathbb{R}$ We can then write the seller's corresponding objective function as:

$$u_s(a_s, s_1; \alpha_b) \equiv \int_c^{\overline{s_2}} (\pi - v_s(s_1, s_2)) f(s_2) ds_2$$

It is clear that $\frac{\partial^2}{\partial a_i \partial s_i} u_i(a_i, s_i; \alpha_{-i}(\cdot)) = 0$. Thus, the SCP-IR is satisfied.

Now consider the seller. Since U_s is decreasing in s_1, we expect his strategy to take the following form:

$$\alpha_s = \begin{cases} Y & \text{if} \quad s_1 \leq d \\ N & \text{otherwise} \end{cases}$$

for some $d \in \mathbb{R}$. We can write the buyer's objective function as:

$$u_b(a_b, s_2; \alpha_s) \equiv \int_{\underline{s_1}}^{d} (v_b(s_1, s_2) - \pi) f(s_1) ds_1$$

Once again, we have that $\frac{\partial^2}{\partial a_i \partial s_i} u_i(a_i, s_i; \alpha_{-i}(\cdot)) = 0$. Thus, the SCP-IR is satisfied. This completes the proof □

4.2 Application

Consider an application of the model presented above. Let s_1 and s_2 be independent random variables given by $s_1 \sim U[0,1]$ and $s_2 \sim U[0,1]$. Let the value functions take the following form:

$$v_s(s_1, s_2) = s_1 + \alpha s_2 \qquad v_b(s_1, s_2) = s_2 + \beta s_1$$

for some $-1 \leq \alpha \leq 1$ and $-1 \leq \beta \leq -1$. That is, both the buyer and seller care not only about their own signal, but about the signal observed by the other agent. Thus in this model there is a double-sided information asymmetry where neither the buyer nor the seller observes the other's signal. Note that this model has the nice property that, by setting either α or β equal to zero, you get the corresponding model with a one-sided information asymmetry. A quick computation [3] shows that the correlation between v_s and v_b is given by

$$\rho = \frac{\alpha + \beta}{\sqrt{(\alpha^2+1)(\beta^2+1)}} \tag{6}$$

Once again, we only care about equilibria in which trade actually occurs. The following theorem characterizes such a Nash equilibrium.

Theorem 6. *Consider the following strategies:*

$$s_s^* = \begin{cases} Y & \text{if} \quad s_1 \le \frac{4\pi - 2\pi\alpha - 2\alpha}{4-\alpha\beta} \\ N & \text{otherwise} \end{cases}$$

$$s_b^* = \begin{cases} Y & \text{if} \quad s_2 \ge \pi - \frac{\beta}{2}(\frac{4\pi - 2\pi\alpha - 2\alpha}{4-\alpha\beta}) \\ N & \text{otherwise} \end{cases}$$

s_b^* and s_s^* form a Bayesian Nash Equilibrium.

Proof. We can apply theorem 4 to show that this is a Bayesian Nash Equilibrium. First note that the sets A_1 and A_2 must take the forms

$$A_1 := \{s_1 | s_1 \le c\} \qquad A_2 := \{s_2 | s_2 \ge d\}$$

for some constants c, d since $U_s(s_1, s_2)$ is decreasing in s_1 and $U_b(s_1, s_2)$ is increasing in s_2. Applying theorem 4 gives us:

$$\begin{aligned} &\frac{1}{P(A_2)} \int_{A_2} v_s f(s_2|s_1) ds_2 \le \pi \\ \implies &\frac{1}{P(s_2 \ge d)} \int_d^1 (s_1 + \alpha s_2) f(s_2) ds_2 \le \pi \\ \implies &\frac{1}{1-d}\{s_1 \int_d^1 1 ds_2 + \alpha \int_d^1 s_2 ds_2\} \le \pi \\ \implies &\frac{1}{1-d}\{s_1(1-d) + \alpha\frac{1}{2}(1-d^2)\} \le \pi \\ \implies &s_1 \le \pi - \frac{\alpha}{2}(1+d) \end{aligned} \tag{7}$$

Similarly, we have:

$$\begin{aligned} &\frac{1}{P(A_1)} \int_{A_1} v_b f(s_1|s_2) ds_1 \ge \pi \\ \implies &\frac{1}{P(s_1 \le c)} \int_0^c (s_2 + \beta s_1) f(s_1) ds_1 \ge \pi \\ \implies &\frac{1}{F_1(c)}\{\int_0^c s_2 f(s_1) ds_1 + \int_0^c \beta s_1 f(s_1) ds_1\} \ge \pi \\ &\frac{1}{c}(s_2 c + \beta \frac{1}{2} c^2) \ge \pi \\ \implies &s_2 \ge \pi - \frac{\beta}{2} c \end{aligned} \tag{8}$$

Equations (10) and (11) give us a system of equations of the following form:

$$\begin{cases} c = \pi - \frac{\alpha}{2}(1+d) \\ d = \pi - \frac{\beta}{2} c \end{cases}$$

Solving this system gives $c = \frac{4\pi - 2\pi\alpha - 2\alpha}{4-\alpha\beta}$ and $d = \pi - \frac{\beta}{2}\frac{4\pi - 2\pi\alpha - 2\alpha}{4-\alpha\beta}$ $\square$

4.3 Buyer Sets Price

Now consider a mechanism where the buyer is free to set the price. The seller's strategy is of the form $s_s : (s_1, p) \to \{Y, N\}$, and the buyer's strategy is now of the form $s_b : s_2 \to p \in \mathbb{R}^+$ where $p = \pi(s_2)$ denotes the price observed by the seller.

Definition 3. *Let A_1 denote the set of (s_1, p) such that $E[U_s(s_1, s_2)|\pi^{-1}(p)]$ is positive. If*

$$\frac{\partial}{\partial \pi^2} \int_{A_1} v_b f(s_1) ds_1 < 2\frac{\partial}{\partial \pi} P(A_1) + \pi \frac{\partial}{\partial \pi^2} P(A_1)$$

then the strategies $s_s* : \begin{cases} A_1 \to Y \\ A_1^C \to N \end{cases}$ *and* $s_b^* = \pi$ *such that* π *solves*

$$\frac{\partial}{\partial \pi} \int_{A_1} v_b(s_1, s_2) f(s_1) ds_1 = P(A_1) + \pi \frac{\partial}{\partial \pi} P(A_1)$$

constitute a Bayesian Nash Equilibrium

Proof. Consider first the problem for the buyer:

$$\max_{\pi} E[U_b]$$
$$\implies \max_{\pi} E[U_b | s_1 \in A_1] P(s_1 \in A_1)$$
$$\implies \max_{\pi} E[v_b(s_1, s_2) - \pi | s_1 \in A_1] P(s_1 \in A_1)$$
$$\implies \max_{\pi} \{E[v_b(s_1, s_2)|A_1] - \pi\} P(A_1)$$
$$\implies \max_{\pi} \{\frac{1}{P(A_1)} \int_{A_1} v_b(s_1, s_2) f(s_1) ds_1 - \pi\} P(A_1)$$
$$\implies \max_{\pi} \int_{A_1} v_b(s_1, s_2) f(s_1) ds_1 - \pi P(A_1)$$

The first condition becomes:

$$\frac{\partial}{\partial \pi} \int_{A_1} v_b(s_1, s_2) f(s_1) ds_1 = P(A_1) + \pi P'(A_1)$$

Finally, the second order condition must be negative in order to ensure π solves for a maximum:

$$\frac{\partial}{\partial \pi} \{\frac{\partial}{\partial \pi} \int_{A_1} v_b(s_1, s_2) f(s_1) ds_1 - P(A_1) - \pi \frac{\partial}{\partial \pi} P(A_1)\} < 0$$

$$\frac{\partial}{\partial \pi^2} \int_{A_1} v_b f(s_1) ds_1 - \frac{\partial}{\partial \pi} P(A_1) - \frac{\partial}{\partial \pi} P(A_1) - \pi \frac{\partial}{\partial \pi^2} P(A_1) < 0$$

$$\frac{\partial}{\partial \pi^2} \int_{A_1} v_b f(s_1) ds_1 < 2\frac{\partial}{\partial \pi} P(A_1) + \pi \frac{\partial}{\partial \pi^2} P(A_1)$$

□

4.4 Application - Linear Strategies

Now consider an application. Let s_1,s_2, and the value functions be the same as in section 4.2.

We are still interested only in equilibria where trade occurs. However, it is much more ambiguous what form such equilibrium strategies must take. For the purposes of this section, I restrict attention to Bayesian Nash Equilibria in linear strategies. Consider the following strategies:

$$s_b^*(s_2) = \pi = s_2 \frac{\alpha\beta - 1}{\beta - 2}$$

and

$$s_s^*(s_1, \pi) = \begin{cases} Y & \text{if} \quad s_1 \leq \pi(1 - \frac{\alpha(\beta-2)}{\alpha\beta-1}) \\ N & \text{otherwise} \end{cases}$$

Theorem 7. *Let $\alpha < \frac{1}{2}$. (s_b^*, s_s^*) constitute a Bayesian Nash Equilibrium.*

Proof. If the seller observes a price p, then his best response is to say yes if his expected utility is positive. That is,

$$E[U_s(s_1, s_2)|\pi^{-1}(p)] \geq 0$$
$$\implies s_1 + \alpha\pi^{-1}(p) \leq p$$
$$\implies s_1 \leq p - \alpha\pi^{-1}(p)$$
$$\implies s_1 \leq p - \alpha(p\frac{\beta - 2}{\alpha\beta - 1}) = p(1 - \frac{\alpha(\beta - 2)}{\alpha\beta - 1})$$

For simplicity, let $c := 1 - \frac{\alpha(\beta-2)}{\alpha\beta-1}$ Now the buyer's best response is given by

$$\max_{\pi} E[U_b(s_1, s_2)]$$

$$\max_{\pi} E[U_b(s_1, s_2)|s_1 \leq \pi c]F_1(\pi c)$$

$$\max_{\pi} E[s_2 + \beta s_1 - \pi|s_1 \leq \pi c]F_1(\pi c)$$

$$\implies \max_{\pi}(s_2 - \pi)\pi c + \beta E[s_1|s_1 \leq \pi c]\pi c$$

$$\implies \max_{\pi}(s_2 - \pi)\pi c + \frac{\beta}{\pi c}\int_0^{\pi c} s_1 f(s_1) ds_1 \pi c$$

$$\implies \max_{\pi}(s_2 - \pi)\pi c + \frac{\beta}{2}(\pi c)^2$$

The first order condition gives us:

$$-\pi c + (s_2 - \pi)c + \beta\pi c^2 = 0$$
$$\implies \pi = \frac{s_2}{2 - \beta c}$$

Which upon substituting for c yields:

$$\pi(s_2) = s_2 \frac{\alpha\beta - 1}{\beta - 2} \tag{9}$$

Finally, we need to check that this is a maximum. The second order condition is given by:

$$-2c + \beta c^2$$

which upon substituting for c yields

$$-\frac{(2\alpha - 1)(\beta - 2)}{(\alpha\beta - 1)^2}$$

The second order condition is negative when $\alpha < \frac{1}{2}$ $\square$

Note that this is also just a straightforward application of definition 3. We have that in equilibrium, the price $p = \pi(s_2)$ must solve:

$$\frac{\partial}{\partial p}\int_{A_1} v_b(s_1, s_2) f(s_1) ds_1 = P(A_1) + p\frac{\partial}{\partial p}P(A_1)$$

For linear strategies, the seller's acceptance set A_1 can be written as $A_1 := \{s_1 | s_1 \leq pc\}$ for some constant c, as we did in the proof of theorem 5. Then the above can be written as

$$\frac{\partial}{\partial p}\int_0^{pc} s_2 + \beta s_1 ds_1 = pc + pc$$

$$\implies (s_2 + \beta pc)c = 2pc$$

$$\implies p = \frac{s_2}{2 - \beta c}$$

Upon seeing this price, the seller's best response is to say yes if:

$$s_1 \leq p - \alpha\pi^{-1}(p)$$

$$\implies s_1 \leq p(1 - \alpha(2 - \beta c))$$

We therefore know that

$$(1 - \alpha(2 - \beta c)) = c$$

$$\implies c = \frac{1 - 2\alpha}{1 - \alpha\beta}$$

To show that this is the same result, we can rewrite c as follows:

$$c = \frac{1 - 2\alpha}{1 - \alpha\beta} = \frac{2\alpha - 1}{\alpha\beta - 1} = \frac{\alpha\beta - 1 - \alpha\beta + 2\alpha}{\alpha\beta - 1} = 1 - \frac{\alpha(\beta - 2)}{\alpha\beta - 1}$$

All that remains is to show that the first condition (concavity) in definition 5 is satisfied:

$$\frac{\partial}{\partial \pi^2}\int_{A_1} v_b f(s_1) ds_1 < 2\frac{\partial}{\partial \pi}P(A_1) + \pi\frac{\partial}{\partial \pi^2}P(A_1)$$

$$\implies \frac{\partial}{\partial \pi}(s_2 + \beta\pi c)c < 2c + 0$$

$$\implies \beta c^2 - 2c < 0$$

It was shown above that this is satisfied whenever $\alpha < \frac{1}{2}$. This completes the proof.

Proposition 2. *All trade that occurs in this equilibrium is efficient*

Proof. From the equilibrium strategies in theorem 7, we know trade occurs if:

$$s_1 \leq \pi(1 - \frac{\alpha(\beta - 2)}{\alpha\beta - 1})$$

$$\implies s_1 \leq s_2(\frac{\alpha\beta - 1}{\beta - 2})(1 - \frac{\alpha(\beta - 2)}{\alpha\beta 1})$$

Upon simplifying, we get that trade occurs if and only if:

$$s_2 \geq s_1\frac{2 - \beta}{1 - 2\alpha}$$

We also know that $v_b \geq v_s$ implies:

$$s_2 \geq s_1\frac{1 - \beta}{1 - \alpha}$$

Thus, to show all trade will be efficient, we need to show that $s_2 \geq s_1 \frac{2-\beta}{1-2\alpha} \implies s_2 \geq s_1 \frac{1-\beta}{1-\alpha}$, i.e. that $\frac{2-\beta}{1-2\alpha} \geq \frac{1-\beta}{1-\alpha}$ We have:

$$\frac{2-\beta}{1-2\alpha} \geq \frac{1-\beta}{1-\alpha}$$
$$\implies 2-\beta-2\alpha+\alpha\beta \geq 1-2\alpha-\beta+2\alpha\beta$$
$$\implies 1 \geq \alpha\beta$$

This is satisfied for all α, β. Thus, trade under this equilibrium will always be efficient.

□

4.5 Application - Nonlinear Strategies

Of course, there may also exist Bayesian Nash Equilibria in nonlinear strategies. Although no such equilibrium is presented here, I do derive a differential equation that such a solution must satisfy. Future research can investigate what boundary conditions to impose and attempt to numerically approximate this solution, as it is difficult to solve analytically.

We will start with the seller's point of view. It was shown above that the seller will agree to trade if and only if:

$$s_1 \leq p - \alpha\pi^{-1}(p)$$

The buyer's best response to this strategy is to solve the following maximization problem:

$$\max_p (s_2 - p)(p - \alpha\pi^{-1}(p)) + \frac{\beta}{2}(p - \alpha\pi^{-1}(p))^2$$

Taking the first derivative with respect to p gives us:

$$-(p-\alpha\pi^{-1}(p)) + (s_2 - p)(1 - \alpha\frac{1}{\pi'(\pi^{-1}(p))}) + \beta(p - \alpha\pi^{-1}(p))(1 - \alpha\frac{1}{\pi'(\pi^{-1}(p)})$$

A necessary condition for an equilibrium is that this equals zero when $p = \pi(s_2) \iff s_2 = \pi^{-1}(p)$:

$$-(\pi(s_2) - \alpha s_2) + (s_2 - \pi(s_2))(1 - \alpha\frac{1}{\pi'(s_2)}) + \beta(\pi(s_2) - \alpha s_2)(1 - \alpha\frac{1}{\pi'(s_2)}) = 0 \tag{10}$$

Thus, any equilibrium pricing strategy must satisfy (9).

Lemma 2. *The price function presented in 6.2 is the solution to (9) for boundary condition* $\pi(0) = 0$

Proof. It is clear that the price function given by

$$\pi(s_2) = s_2\frac{\alpha\beta - 1}{\beta - 2}$$

satisfies the boundary condition $\pi(0) = 0$. Thus we only need to show it satisfies (9). Plugging in $\pi^{'}(s_2) = \frac{\alpha\beta-1}{\beta-2}$ to the left hand side of (9) gives

$$\frac{(\alpha-1)\{\pi(s_2)(\beta-2) + s_2 - \alpha\beta s_2\}}{\alpha\beta - 1}$$

Finally, plugging in $\pi(s_2) = s_2\frac{\alpha\beta-1}{\beta-2}$ yields

$$\frac{(\alpha-1)\{s_2(\alpha\beta-1) + s_2 - \alpha\beta s_2\}}{\alpha\beta - 1} = 0$$

□

4.6 Application - Price Floors

What happens if we impose a price floor ($p^* > 0$) on the market? That is, we impose the condition that the buyer can set any price as long as $p \geq p^*$? The equilibrium discussion becomes more complicated because equilibria may now be part pooling and part revealing.

We can imagine that the buyer would want to leave the price at the price floor for all signals $s_2 \leq s_2^*$ where s_2^* is the signal that is mapped to p^*, when he would decide to raise it. Now consider the seller. If the seller observes a price $p > p^*$, his strategy will be the same as before. But if he observes a price p^*, he now can no longer infer the buyer's signal - he must take an expectation over s_2. Upon observing p^*, the seller infers:

$$E[s_2|s_2 \leq \pi^{-1}(s_2^*)] = \frac{1}{s_2^*}\int_0^{s_2^*} s_2 f(s_2)ds_2 = \frac{s_2^*}{2}$$

Therefore, the buyer's and seller's strategies may take the following form:

$$s_b = \begin{cases} p & \text{if} \quad s_2 > s_2^* \\ p^* & \text{if} \quad s_2 \leq s_2^* \end{cases}$$

$$s_s = \begin{cases} Y & \text{if} \quad p > p^* \quad \text{and} \quad s_1 \leq p - \alpha\pi^{-1}(p) \\ Y & \text{if} \quad p = p^* \quad \text{and} \quad s_1 \leq p - \frac{\alpha}{2}\pi^{-1}(p) \\ N & \text{otherwise} \end{cases}$$

The trouble, however, is that this introduces a discontinuity in the seller's strategy. This problem is best understood graphically:

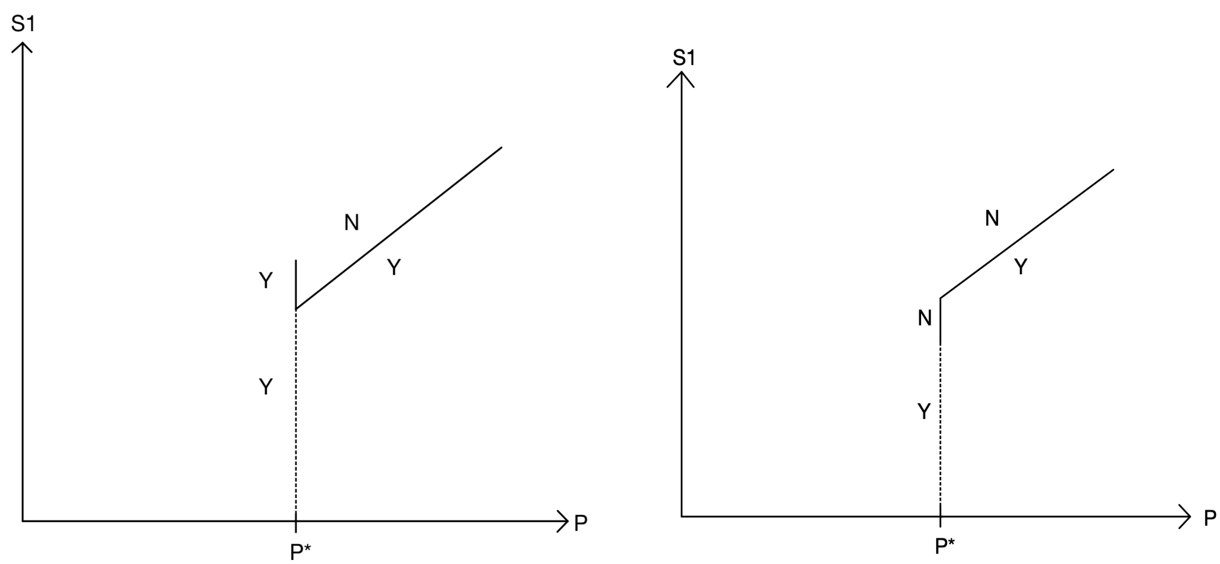

$\alpha > 0$ $\alpha < 0$

Figure 8: Seller's Strategy

(a) $\alpha \neq 0$

Figure 10 shows the region of the (p, s_1) plane where the seller will agree to trade - notice there exists a jump in his strategy at $p = p^*$. Consider first the case where $\alpha > 0$. If the buyer had signal $s_2^* + \epsilon$ for arbitrarily small ϵ, he would have an incentive to deviate from his strategy and keep the price at p^* to increase the probability of trade. Now consider $\alpha < 0$. If the seller had a signal s_2^*, he would have an incentive to deviate and set a price $p = p^* + \epsilon$, once again to increase the probability of trade.

To avoid the complications posed by pooling, we could instead consider only fully revealing equilibria, such as those given by strategies of the following form:

$$s_b = \begin{cases} p \quad \text{if} \quad s_2 \geq s_2^* \\ \text{Drop Out} \quad \text{if} \quad s_2 < s_2^* \end{cases}$$

$$s_s = \begin{cases} Y \quad \text{if} \quad s_1 \leq p - \alpha\pi^{-1}(p) \\ N \quad \text{otherwise} \end{cases}$$

In order for such strategies to constitute a Bayesian Nash Equilibrium, we must impose one more requirement: the buyer's expected utility at the price floor must be zero, i.e. he must be indifferent between trading and not trading. To understand why this is the case, imagine that the buyer's expected utility at the price floor is positive. Then a buyer observing signal $s_2^* - \epsilon$ for arbitrarily small ϵ will have an incentive to set a price p^* and signal to the seller that he actually has observed s_2^*. If his expected utility at the price floor were negative, he would clearly do better by dropping out. Thus, it must be the case that in equilibrium, the buyer's expected utility at the price floor is zero.

Proposition 3. *The price floor game contains no fully revealing Bayesian Nash Equilibria in linear strategies*

To prove this I first show that the price formula presented in section 6.2 is the only linear solution to (9). I then show that the expected utility from a positive price floor will always be positive.

Assume that there exists a price function of the form $\pi(s_2) = s_2 m + b$ which forms an equilibrium strategy in the signaling game. Then the seller's best response function is to say yes if and only if:

$$s_1 \leq p - \alpha\pi^{-1}(p) = p - \alpha(\frac{p-b}{m})$$

The buyer's best response is then given by

$$\max_p (s_2 - p)(p - \alpha(\frac{p-b}{m})) + \frac{\beta}{2}(p - \alpha(\frac{p-b}{m}))^2$$

Setting the first derivative with respect to p equal to zero gives us:

$$-(p - \alpha(\frac{p-b}{m})) + (s_2 - p)(1 - \frac{\alpha}{m}) + \beta(p - \alpha(\frac{p-b}{m}))(1 - \frac{\alpha}{m}) = 0$$

Solving for p gives

$$p = s_2(\frac{\alpha m - m^2}{(\alpha - m)(\alpha\beta + 2m - \beta m)}) + (\frac{\alpha^2\beta b + \alpha b m - \alpha\beta b m}{(\alpha - m)(\alpha\beta + 2m - \beta m)})$$

We then get the following system for m and b:

$$\begin{cases} m = \frac{\alpha m - m^2}{(\alpha - m)(\alpha\beta + 2m - \beta m)} \\ b = \frac{\alpha^2\beta b + \alpha b m - \alpha\beta b m}{(\alpha - m)(\alpha\beta + 2m - \beta m)} \end{cases}$$

Solving yields $m = \frac{\alpha\beta - 1}{\beta - 2}$ and $b = 0$. Thus, this must be the only linear price function in the signaling game.

The expected utility from setting a price $\pi(s_2) = s_2(\frac{\alpha\beta - 1}{\beta - 2})$ is given by

$$E[U_b | s_1 \leq p - \alpha\pi^{-1}(p)](p - \alpha\pi^{-1}(p))$$

which upon plugging in simplifies to:

$$E[U_b] = \frac{s_2^2(2\alpha - 1)}{2(\beta - 2)}$$

Remember that the price only maximizes utility if $\alpha < \frac{1}{2}$. Thus, the expected utility can only be zero if $s_2 = 0$. For positive price floors, this is never the case. Thus, there exists no fully revealing Bayesian Nash Equilibrium in linear strategies.

5 Efficiency

5.1 Independent Private Values in the Fixed Price Mechanism

Imagine a buyer and a seller are looking to engage in trade, and neither agent cares about the other's valuation for the good - they care only about their own value. This is the well-studied independent private value case, and in the model presented in section 5, corresponds to the case in which $\alpha = \beta = 0$. In the fixed price mechanism, the equilibrium strategies (using theorem 6) are for the seller to say yes to trade if $s_1 \leq \pi$ and for the buyer to say yes to trade if $s_2 \geq \pi$. Therefore, trade occurs only if $s_2 \geq s_1 \implies v_b \geq v_s$. Thus, in the fixed price mechanism with independent private values, trade will always be efficient. Indeed, interactive modeling in Mathematica confirms this:

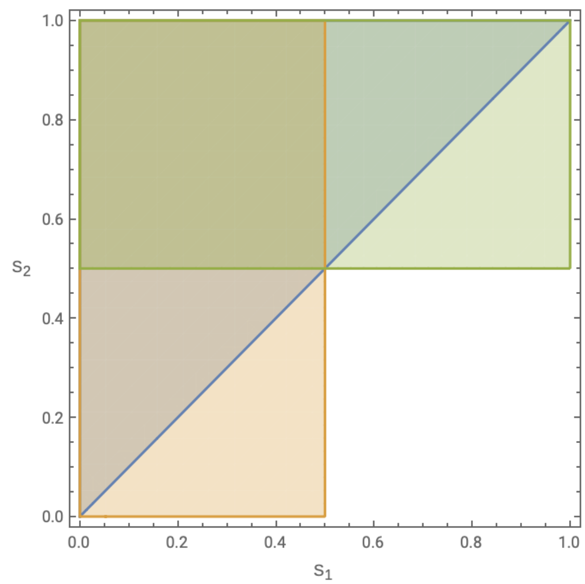

Figure 10: $s_1, s_2 \sim U[0, 1]$, $\alpha = \beta = 0$, $\pi = 0.5$

In figure (10), the green region is the set of s_2 for which the buyer agrees to trade, the orange region is the set of s_1 for which the seller will agree to trade, and their intersection represents the region where trade actually occurs. The blue line dictates the efficient region, with any trade occurring above the blue line being efficient. Varying the price π confirms that the region of trade always lies above the blue line, and is therefore always efficient. However, the well-known problem with the fixed price mechanism is that too little trade occurs. Only 50% of the efficient region in figure 10 is captured by trade. A natural question, then, is does the fixed price mechanism perform differently when we vary α and β (and therefore the correlation between v_b and v_s)?

5.2 Fixed Price Mechanism - Correlated Values

To answer this, consider first changes in α:

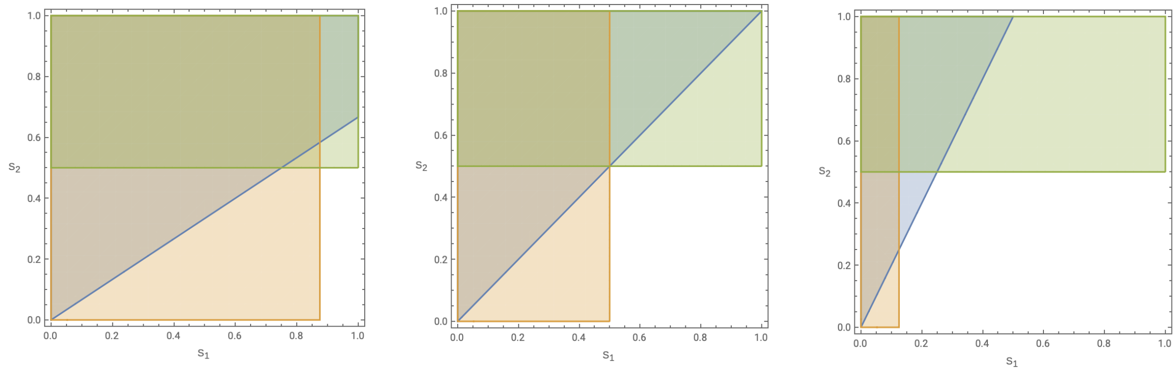

$\alpha = -0.5$ $\alpha = 0$ $\alpha = 0.5$

Figure 11: Effect of Changes in α

(a) $s_1, s_2 \sim U[0,1], \quad \pi = 0.50,\, \beta = 0$

There are several things to note from figure 11. First, changing α has no impact on the set of s_2 for which the buyer agrees to trade. This makes sense, the buyer's value for the good does not depend on α. Second, note that decreasing α increases the region for which the seller agrees to trade, and increasing α will decrease this region. To understand why this occurs, consider a situation in which a buyer and seller are looking to transact a good, but the seller puts a negative weight on the buyer's signal ($\alpha < 0$). When the seller sees that the buyer wants to engage in trade, he infers that the buyer's signal is high. Because α is negative, this reduces the seller's perception of his own value, and *increases* his willingness to trade for a given s_1. Similarly, if the seller puts a positive weight on the buyer's signal, the buyer's willingness to trade would increase the seller's perception of his valuation, and *decrease* his willingness to trade. Finally, note that for sufficiently large decrease in α, inefficient trade may arise under the fixed price mechanism.

Now consider changes in β:

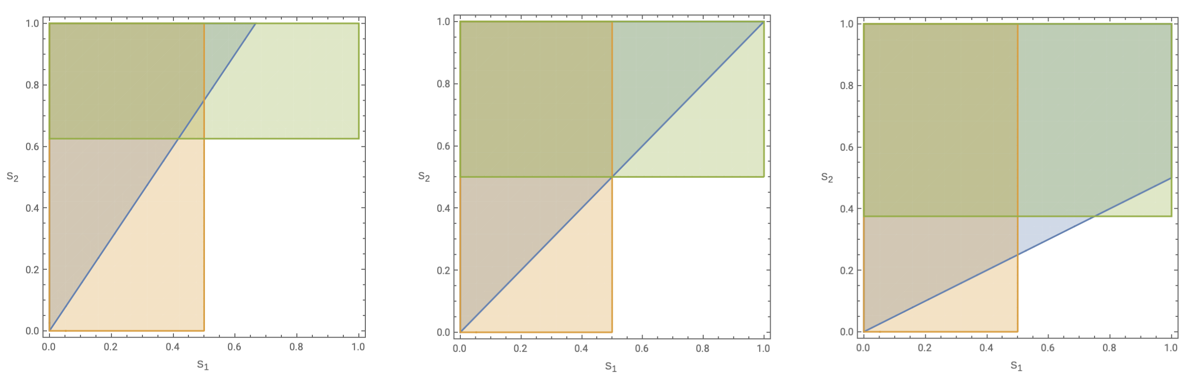

$\beta = -0.5$ $\beta = 0$ $\beta = 0.5$

Figure 13: Effect of Changes in β

(a) $s_1, s_2 \sim U[0,1], \quad \pi = 0.50,\, \alpha = 0$

As before, changes in β have no effect on the seller's willingness to trade - since the seller's valuation does not depend on β, the set of s_1 for which the seller agrees to trade remains constant. However, changing β has the opposite effect on the buyer's willingness to trade than changing α has on the seller's willingness to trade. Decreasing β reduces the set of s_2 for which the buyer will engage in trade, and increasing β increases this set. As before, imagine a buyer and seller looking to transact a good, and imagine that the buyer puts a negative weight on the seller's signal. Upon seeing the seller willing to trade, he will infer the seller's signal and, because of the negative weight, reduce his own value for the good. This reduces the buyer's willingness to trade. Similarly, a positive value of β will increase the buyer's perception of his own valuation for any given (s_1, s_2) and therefore make him more willing to engage in trade. Finally, as before, a negative value of β may lead to inefficient trade occurring.

One interesting implication of the model above is that for a given correlation, the fixed price mechanism can perform quite differently. Consider the following two cases:

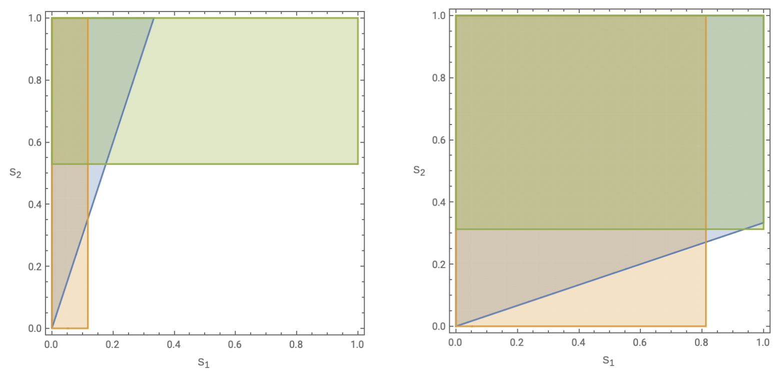

$(\alpha, \beta) = (0.5, -0.5)$ $\quad$ $(\alpha, \beta) = (-0.5, 0.5)$

Figure 15: $\rho = 0$

(a) $s_1, s_2 \sim U[0, 1], \quad \pi = 0.50$

In figure 15, both cases correspond to zero correlation between the buyer's and seller's values, but the fixed price mechanism performs quite differently in the two cases. The graphs suggest that the fixed price mechanism works best when α is negative and β is positive. Intuitively, it works best when the seller negatively weighs the buyer's signal (thereby reducing his value for any given signals and increasing his willingness to trade) and when the buyer positively weighs the seller's signal (thereby increasing his value for any given signals and increasing his willingness to trade). Another implication is that in addition to how the values are correlated, it is important for the mechanism designer to understand *why* they are correlated thus. Figure 15 shows that the fixed price mechanism may perform differently for the same correlation.

5.3 Independent Private Values - Buyer Sets Price

Once again, we are interested for what values of s_1 and s_2 trade occurs in this game, and how does this change as we change α and β. In the independent private value case, where each agent cares only about his own signal, the equilibrium strategies are (according to section 6.2) as follows:

$$s_b^* = \pi(s_2) = \frac{s_2}{2} \quad s_s^* = \begin{cases} Y & \text{if} \quad s_1 \leq \pi \\ N & \text{otherwise} \end{cases}$$

Since $s_1 \leq \pi \implies 2s_1 \leq s_2 \implies v_b \geq v_s$, we know all trade that occurs will be efficient. Indeed, this is what we see when plotting the regions in Mathematica:

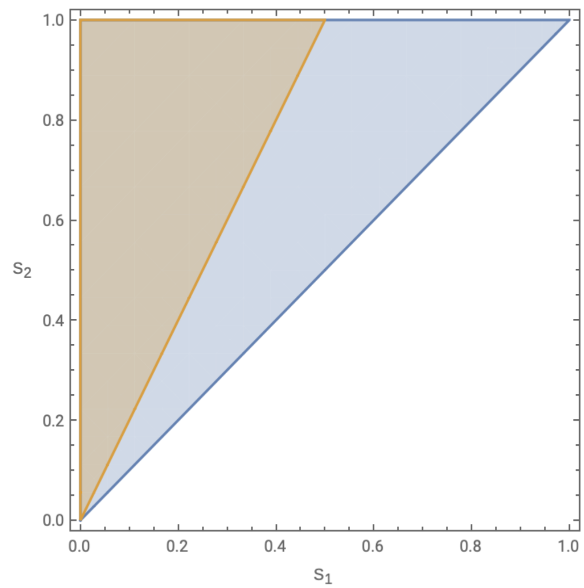

Figure 17: $s_1, s_2 \sim U[0, 1]$, $\alpha = \beta = 0$

In figure 19, the blue region is the region where trade would be efficient (i.e. the buyer values the good more than the seller), and the orange region is the region where trade actually occurs. Clearly, all trade is efficient in this game, but once again we want to know if changing α and β can improve the efficiency of this mechanism.

5.4 Correlated Values- Buyer Sets Price

To understand how changing α and β impacts the efficiency of the mechanism, consider first changes in α:

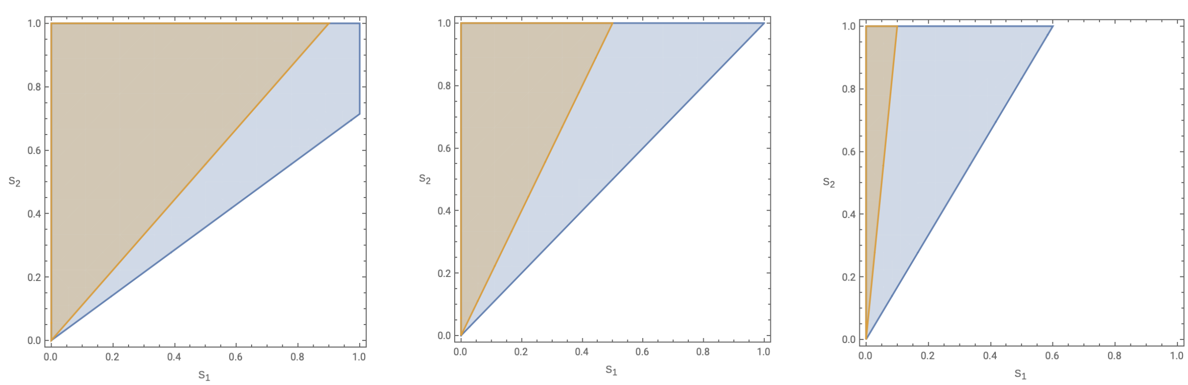

$\alpha = -0.4$ $\qquad$ $\alpha = 0$ $\qquad$ $\alpha = 0.4$

Figure 18: Effect of Changes in α

(a) $s_1, s_2 \sim U[0, 1], \quad \beta = 0$

We see similar results as we did in section 7.2: namely, that lower values of α promote trade while higher values of α discourage trade. When α decreases, the seller lowers his valuation of the good for any given (s_1, s_2), and is therefore more willing to trade. On the other hand, the buyer responds to a lower α by decreasing the price he charges (since $\frac{\partial \pi}{\partial \alpha} > 0$), which would lower the seller's willingness to trade. The results in figure 20 indicate that the first effect is more significant; that is, the seller reduces his value for the good by more than the buyer will decrease the price, so that the overall effect is an overall increase in his willingness to trade.

Therefore, we can see that a mechanism in which the buyer sets the price is much more efficient in markets with low α's, whereas they are quite inefficient in markets with high α's. Now consider the effect of changes in β:

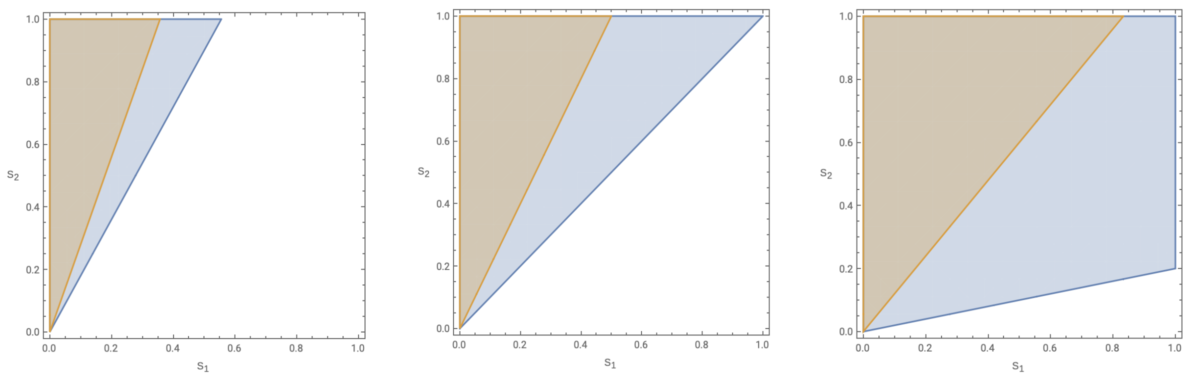

$\beta = -0.8$ $\qquad$ $\beta = 0$ $\qquad$ $\beta = 0.8$

Figure 20: Effect of Changes in β

(a) $s_1, s_2 \sim U[0,1], \quad \alpha = 0$

As expected, we see the opposite effects. Increases in β promote trade while decreases in β discourage trade. When β increases, the buyer's valuation for any given (s_1, s_2) increases, and he is therefore more willing to engage in trade. Thus, mechanisms in which the buyer can set the price are more efficient in markets with high β's and less efficient in markets with high β's.

There is one more thing to note from figures 20 and 22: changing β has a much smaller impact on the amount of trade that occurs than changing α. The likely intuitive explanation for this is that in this mechanism, the buyer has control over the price. He can therefore respond to a higher β by adjusting the price he sets, and thus his expected utility need not fluctuate widely from changes in β. But the seller has much less scope to respond to changes in α - he must always choose one of two actions (to say yes or no to trade). If the seller's value of α is high, and he infers from the price that the buyer's signal is high, he does not have the freedom to set or negotiate a higher price. Thus, he may simply respond by refusing to trade. We can therefore conjecture that, in the reverse case where the seller is able to set the price and the buyer must either choose to say yes or no to trade, the region of trade will respond much more drastically to changes in β than to changes in α.

Finally, to once again illustrate how differently this mechanism may play out for the same correlation, consider the following to cases:

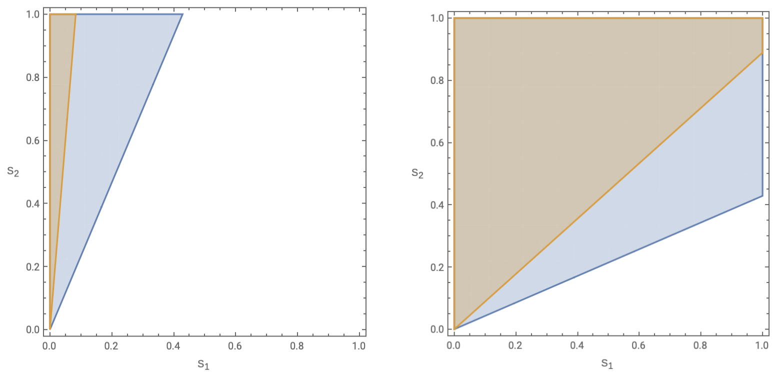

$(\alpha, \beta) = (0.4, -0.4)$ $(\alpha, \beta) = (-0.4, 0.4)$

Figure 22: $\rho = 0$

(a) $s_1, s_2 \sim U[0, 1]$

Figure 24 demonstrates two different markets in which v_b and v_s both have a correlation of 0, yet the mechanism is much more efficient in the market on the right than the market on the left. This suggests that, just as with the fixed price mechanism, the mechanism where the buyer sets the price is more efficient when α is low and β is high but less efficient in the opposite case.

6 Discussion

6.1 Implications

There are several noteworthy implications of the models studied above. The first is that, in the models presented in sections 2 and 4, the correlation between v_b and v_s did not seem to be directly responsible for market outcomes. In particular, recall that in the model presented in section 4, one could construct two markets with the same correlation but vastly different market outcomes and efficiency properties. Rather, it was important to understand *why* the buyer's and seller's values exhibited a certain correlation. Note also that in the model presented in section 3, although the price the buyer set did depend on ρ, the effect of changing ρ itself depended on σ_b. What's more, in that model a given correlation corresponded to more than one price, just as was seen in the model with the two sided information asymmetry. This again suggests that what's most important is to understand why the values exhibit a certain correlation. For example, are values positively correlated because the buyer puts a positive weight on the seller's signal, or because the seller puts a positive weight on the buyer's signal? The two cases have opposite economic properties, and it is important to differentiate between them.

Second, the fixed price mechanism is sometimes more efficient than the mechanism in which the buyer can set the price. Although a true analysis of efficiency gains requires that we compare the fixed price mechanism not with a freely floating price but with a price floor mechanism where the price floor is set at the fixed price, it is still interesting to discuss why we see the fixed price mechanism outperform in some markets. Computer simulations suggest that in the model presented in this paper, the fixed price mechanism outperforms in markets with high α's and low β's. For example, in a market with $(\alpha, \beta) = (0.4, 0)$, the probability of trade occurring is 10% in the fixed price mechanism but only 5% when the buyer can set the price. Why? The likely explanation is because the game in which the buyer can set the price is a signaling game. In the real world, we can only use private information to the extent that we are willing to reveal it (or something about it). In markets with low α's, the marginal cost to revealing private information is low, because the seller does not put a significant weight on the buyer's signal. But in markets with high α's, the marginal cost to revealing private information is much higher, and this reduces the efficiency of the mechanism. This helps explain why the fixed price mechanism performs better in high α contexts: the seller can longer infer anything about the buyer's signal if the price is fixed by a third party.

To illustrate a more concrete example, consider the insurance market. Insurance markets are plagued by adverse selection (people who are sicker tend to value insurance more). We can therefore expect such markets to have high α's and low β's, since the insurance underwriter likely cares a lot about the buyer's private signal regarding his health, but the buyer probably does not care too much about the insurer's value for the contract. In other words, this is a context where the marginal cost to the buyer of revealing his private signal is potentially very large. In such a market, this model suggests it would be preferable to let a third party such as the government fix the price of the contract rather than let the buyer propose a price, because the insurer will try to use that price to infer the true health status of the buyer.

Third, with information asymmetries, trade under the fixed price mechanism is *not* always efficient. The first panels in figures 11 and 13 demonstrate that trade may occur even if $v_s > v_b$, something that would not happen with complete information. However, proposition 2 proves that trade under the equilibrium proposed in section 6.4 is always efficient. This demonstrates a trade-off between the two mechanisms. The fixed price mechanism may be preferable in high α contexts because it removes the possibility of the buyer losing from revealing his information, but the fixed price leaves the buyer and seller with little flexibility. If the buyer can set the price, he can adjust the price according to his signal and the values of α and β, and therefore he has much more control over his (and the seller's) expected gains from trade. But in the fixed price mechanism, agents can only decide between saying yes or no to trade. In other words, when rational agents have control over the price at which they trade, they may adjust this price to prevent inefficiencies from arising. But when the price is fixed by a third party, they do not have this freedom, and must either agree to trade in a mechanism which may result in inefficient trade, or say no to trade altogether.

6.2 Limitations and Future Research

There are several limitations of the results presented in this paper. The largest and most significant is that, with the exception of proposition 1, all the results discussed in this paper are based on the assumption of specific probability distributions. Future research should investigate the model under other probability distributions. Even more effective would be an understanding of how these mechanisms operate under the assumption of an arbitrary density function. Second, the analysis of the game in which the buyer sets the price is incomplete. There exist more solutions to the differential equation proposed in section 4.5, and future research should investigate other boundary conditions and attempt to approximate other solutions, either numerically or analytically. Third, a true discussion of efficiency gains requires us to compare the fixed price mechanism with the price floor mechanism in which the buyer may raise the price, rather than a game in which the buyer can set any price. The price floor game was briefly investigated in 4.6, and some of the challenges were discussed. It was shown that in this game, there exists no fully revealing Bayesian Nash Equilibrium in linear strategies. However, there may exist fully revealing equilibria in nonlinear strategies (of a form that satisfies equation (9)), and future research should investigate this possibility. Furthermore, there may exist equilibria that are part pooling and part revealing. Thus, a more general analysis of equilibria in the price floor game under correlated values is still needed. Finally, a limitation for the mechanism designer is that it may be quite hard to estimate the values of α and β in practice. It is likely that the buyer and seller do not even know there own implicit values of α and β, and so implementing the theoretical conclusions in this paper may be somewhat challenging. Future empirical work could investigate if there is a way to extrapolate implicit values of α and β from behaviors observed in real life markets.

7 Conclusion

This paper studied the effect of information asymmetries on trade and price formation in markets with correlated values. The models presented in sections 2 and 4 suggest that it is not the correlation itself that dictates market outcomes, but why we observe that correlation in the first place. In particular, we can construct theoretical markets with the same correlation but widely varying market outcomes. The correlation does play a significant role in the model presented in section 3, but it is still true that the same correlation may correspond to more than one price being set. Section 4 proposed a model with a double-sided information asymmetry which has the convenient property that by setting either α or β equal to zero, you achieve the corresponding model with a one-sided asymmetry. Throughout the paper, I focused on two mechanisms: a fixed price and a mechanism in which the buyer can set the price. It was shown that fixed price mechanisms outperform in markets where the marginal cost to revealing private information is high, such as insurance markets. However, due to their having less flexibility, fixed price mechanisms may also result in some inefficient trade. Future research should investigate the efficiency properties of these mechanisms further, and extend these results to other plausible trading mechanisms.

8 Appendix

1. We have $E[v_b] = v_b^L(p+r) + v_b^H(q+s)$ and $E[v_s] = v_s^L(p+q) + v_s^H(r+s)$. We also know that $E[v_b v_s] = v_b^L v_s^L p + v_b^L v_s^H r + v_b^H v_s^L q + v_b^H v_s^H s$. Finally, we have that $Var(v_b) = E[v_b^2] - E[v_b]^2 = (v_b^L)^2(p+r) + (v_b^H)^2(q+s) - (v_b^L(p+r) + v_b^H(q+s))^2$ and $Var(v_s) = E[v_s^2] - E[v_s]^2 = (v_s^L)^2(p+q) + (v_s^H)^2(r+s) - (v_s^L(p+q) + v_s^H(r+s))^2$. Plugging in $v_b^L = v_s^L = 0$, $v_b^H = v_s^H = 1$, and using the fact that $p+q+r+s = 1$, we have that:

$$\rho(q,r,s) = \frac{-s(s+r-1) - q(r+s)}{\sqrt{(q^2 + s(s-1) + q(2s-1))(r^2 + s(s-1) + r(2s-1))}}$$

2. We have that $E[v_s] = \mu_{1+}\alpha\mu_2$ and $E[v_b] = \mu_2 + \beta\mu_1$. Furthermore, $v_b v_s = \beta s_1^2 + \alpha s_2^2 + s_1 s_2(1+\alpha\beta)$, so $E[v_b v_s] = [s_1^2] + \alpha E[s_1^2] + (1+\alpha\beta)E[s_1 s_2] = \beta(\mu_1^2 + \sigma_1^2) + \alpha(\mu_2^2 + \sigma_2^2) + (1+\alpha\beta)\mu_1\mu_2$. This now gives us the covariance: $cov(v_b, v_s) = E[v_b v_s] - E[v_b]E[v_s] = \beta\sigma_1^2 + \alpha\sigma_2^2$. Finally, dividing by the standard deviations gives us the correlation: $\rho = \frac{\beta\sigma_1^2 + \alpha\sigma_2^2}{\sqrt{\sigma_1^2 + \alpha^2\sigma_2^2}\sqrt{\sigma_2^2 + \beta^2\sigma_1^2}}$.

3. We have that $E[v_s] = E[s_1 + \alpha s_2] = \frac{1}{2} + \frac{\alpha}{2}$. Similarly, we have $E[v_b] = \frac{1}{2} + \frac{\beta}{2}$. Now, $v_s v_b = (s_1 + \alpha s_2)(s_2 + \beta s_1) = \alpha s_2^2 + \beta s_1^2 + s_1 s_2(1+\alpha\beta)$ and $E[v_s v_b] = \frac{\alpha}{3} + \frac{\beta}{3} + \frac{1+\alpha\beta}{4}$. To calculate covariance, we have $cov(v_s, v_b) = E[v_s v_b] - E[v_s]E[v_b] = \frac{\alpha+\beta}{12}$. We also know that $Var(v_s) = Var(s_1 + \alpha s_2) = Var(s_1) + \alpha^2 Var(s_2) = \frac{\alpha^2+1}{12}$ and $Var(v_b) = \frac{\beta^2+1}{12}$. Dividing by the standard deviations gives us the correlation: $\rho = \frac{\alpha+\beta}{\sqrt{(\alpha^2+1)(\beta^2+1)}}$.

References

[1] Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. The Quarterly Journal of Economics, 1970, 84(3), 488.

[2] Athey, S. (2001). Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. Econometrica, 69(4), 861-889.

[3] Börgers, T.; Li, J. (2019). Strategically simple mechanisms. Econometrica, 87(6), 2003-2035.

[4] Cremer, J.; McLean, R. P. (1985). Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. Econometrica, 53(2), 345.

[5] Cremer, J.; McLean, R. P. (1988). Full extraction of the surplus in Bayesian and dominant strategy auctions. Econometrica, 56(6), 1247.

[6] Hagerty, K. M.; Rogerson, W. P. (1987). Robust trading mechanisms. Journal of Economic Theory, 42(1), 94-107.

[7] Hayek, F. (1945). The Use of Knowledge in Society. American Economic Review, 519-530.

[8] Milgrom, P.; Stokey, N. (1982). Information, trade and common knowledge. Journal of Economic Theory, 26(1), 17-27.

[9] Myerson, R. B.; Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. Journal of Economic Theory, 29(2), 265-281.