

Blouin, Max; Serrano, Roberto

Working Paper

A Decentralized Market with Common Values Uncertainty: Non-Steady States

Working Paper, No. 1998-5

Provided in Cooperation with:

Department of Economics, Brown University

Suggested Citation: Blouin, Max; Serrano, Roberto (1998) : A Decentralized Market with Common Values Uncertainty: Non-Steady States, Working Paper, No. 1998-5, Brown University, Department of Economics, Providence, RI

This Version is available at:

<https://hdl.handle.net/10419/102608>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

A Decentralized Market with Common Values Uncertainty: Non-Steady States*

Max R. Blouin
Brown University

Roberto Serrano
Brown University

August 10, 1998

Abstract

We analyze a market where (i) trade proceeds by random and anonymous pairwise meetings with bargaining; (ii) agents are asymmetrically informed about the value of the traded good; and (iii) no new entrants are allowed once the market is open. We depict this infinite-horizon market as a strategic-form game and, adapting existence theorems from the literature on anonymous games, we prove the existence of a Nash equilibrium for every value of the discount factor. The main questions are whether information revelation and efficiency obtain as the economy becomes approximately frictionless (i.e. as discounting is gradually removed). We show that this is not the case. This negative result holds whether the asymmetry is two-sided or restricted to one side of the market. This contrasts with the earlier literature, which was based on steady-state equilibria.

KEYWORDS: information revelation, interim incentive efficiency, decentralized markets, non-steady states.

JEL CLASSIFICATION NUMBERS: C72, C78, D82, D83.

1 Introduction

In the Walrasian model of competitive market equilibrium, trade is *centralized*: all sellers and buyers of a particular commodity meet in the same location, and all trade takes place simultaneously at the same price. This price is chosen by an auctioneer to equate demand with supply, and no trade takes place until it has been made public. The model is frictionless, in that there are no impediments to trade.

*Blouin thanks the Morgan Stanley Corporation for financial support and Serrano acknowledges a Salomon research award from Brown University.

A variation of this setup allows for asymmetric information about the value of the good being traded, an example of a problem with common values uncertainty. In such a market, the auctioneer plays a key informational role. Suppose the value of the good being traded depends on the state of the world. If the equilibrium price is known to vary with the state, then inference about the state (and hence about the value of the good) can be drawn as soon as the price is announced. The *rational expectations equilibrium* (REE) is an equilibrium concept used in this framework. In a centralized analog of the case we study, it exhibits full information revelation and is an interim incentive efficient mechanism, a sort of first welfare theorem.¹

Our main goal in this paper is to study the information revelation and efficiency properties of a *decentralized* market, where there is no auctioneer and transactions take place in pairwise meetings of agents. In this context, information about the state of the world cannot be aggregated in a centrally called price. Any information revelation must be carried out by the trading mechanism itself.

A recent line of research, begun by Wolinsky (1990), has looked at the extent to which this is possible.² Wolinsky (1990), and Serrano and Yosha (1993) after him, perform the analysis under a very strong steady-state restriction: each period a constant population of agents enters the market, and an equal population must trade and exit the market. We depart in this paper from the steady state analysis by studying a one-time entry market. A different approach is required both to establish existence of equilibrium and to answer our basic questions. In particular, we employ machinery for infinite-dimensional sequences. We describe our model presently.

The market consists of a continuum of buyers and a continuum of sellers, all present from the outset. No new traders beyond these initial populations participate in the market. Trade is decentralized: buyers and sellers do not all meet simultaneously in a central market. Rather, each seller meets one buyer at a time, and vice versa. When a buyer and seller meet, they bargain. If they agree on a price, they transact and leave the market. If they disagree, they remain in the market: the following period, the seller meets another buyer, and the buyer meets another seller. These meetings are random and anonymous. The good for sale is indivisible. Each seller has one unit, and each buyer wants at most one unit. Its value is either high (H) or low (L), and it is the same for all units of the good in all periods. That is, only one state (value) prevails throughout the game. However, some agents do not know whether it is H or L .

¹See Holmström and Myerson (1983) for this efficiency concept, and Laffont (1985) and Serrano and Yosha (1996) for two uses of it.

²His paper continues a literature of decentralized strategic markets that started with Rubinstein and Wolinsky (1985) and Gale (1986a, 1986b). These papers, though, deal with economies without asymmetric information. See also Osborne and Rubinstein (1990, chapters 7-8) and the references therein.

The bargaining process, taken from Wolinsky (1990), is simple. When a buyer and seller meet, they simultaneously announce bargaining positions: either *tough* or *soft*. For a seller, playing tough means insisting that the value is high and asking for an accordingly high price; playing soft means agreeing to a low price. For a buyer, it is the reverse. If both play tough, the meeting ends in disagreement. Agreement results in all other cases, and the price of the transaction depends on who played soft. There is a cost to delay (through disagreement), represented by a discount factor $\delta < 1$ common to all agents.

By playing tough for a number of periods, an agent effectively searches for a trading partner who will agree to his conditions. Search is costly, however, and the agent may eventually give in and agree to his next partner's conditions. In this sense, the market is a vast war of attrition.³ We can imagine: (i) a buyer shopping for an item, going from store to store until either he finds a low enough price or he gets tired of shopping and buys at a high price; and (ii) a merchant seeing a succession of customers walk in, look at the price tag, and leave, until either one of them buys the item or he (the merchant) gets tired of waiting and lowers the price.

If any information is conveyed to uninformed agents about the value of the good, it is through this process of repeated bargaining. For example, an uninformed buyer might expect all sellers to play soft if the value is L . If he meets a seller playing tough, he deduces that the value is actually H . Typically, though, the buyer's problem will not be that easy. He will expect to meet tough or soft sellers whatever the value of the good, albeit with different probabilities. His inference will be based on these probabilities and will be drawn over several periods. The question is whether discounting will wear him out before he can draw the correct inference.

Discounting is in fact the only friction in this model. Therefore, if we make the discount factor arbitrarily close to 1, we can make a legitimate comparison between our results and the centralized benchmark. We ask whether (as $\delta \rightarrow 1$) the economy moves towards a situation of (i) ex-post individually rational (EPIR) trade in both states of the world, and (ii) efficiency in both states.

We distinguish between two-sided and one-sided uncertainty. Two-sided uncertainty refers to the case where both sellers and buyers can be uninformed; one-sided uncertainty is the case where only some buyers (or only some sellers) can be. In the two-sided case, we find that (as $\delta \rightarrow 1$) a sizeable fraction of agents either does not trade or does so at non-EPIR prices. In the one-sided version, all agents trade, but (again as $\delta \rightarrow 1$) a non-negligible fraction does so at the wrong prices. Also, in both cases, there is asymptotic inefficiency. That is, the lack of information revelation

³To our knowledge, our paper is the first study of the war of attrition in a pure common values setting. See, for example, Krishna and Morgan (1997) for a model that combines the war of attrition and private values.

does not merely cause a pure transfer from the uninformed to the informed: society as a whole loses out, due to excessive delay in learning.

Wolinsky (1990) solved for equilibria with two-sided uncertainty. He found that in such a configuration a sizeable fraction of those agents who trade do so at prices which are not EPIR, even as the discount factor is made arbitrarily close to 1. Serrano and Yosha (1993) performed the same test for one-sided uncertainty and found a quite different result: typically all agents who trade do so at EPIR prices.

It is instructive to examine the forces at work in order to understand the results. We detect four economic forces. The first is *cost of learning* (CL): as the discount factor δ approaches 1, search becomes less costly for uninformed agents. The second is *misrepresentation of information* (MI): as $\delta \rightarrow 1$, it is cheaper for an informed agent to lie in his favor about the value of the good. These first two forces are present in both versions of the model. The next two appear only in the two-sided version. The third force is *noise* (N): as $\delta \rightarrow 1$, meetings between two uninformed agents may become more and more common, due to force CL. Meetings of this sort are not conducive to learning: in the extreme, if all agents were uninformed, the set of equilibria would not depend on the state the world at all. Finally, there is *fear* (F): for some parameter values, uninformed agents may play tough indefinitely in order to avoid a loss which they believe likely. Such pessimism feeds upon itself: uninformed sellers are afraid only if they know a good number of uninformed buyers to be afraid, and vice versa. Force CL works in favor of information revelation, while forces MI and N work against it. Force F, when present, works to prevent trade from taking place.

Our investigation shows that in the one-sided model, force MI overcomes force CL. This is in sharp contrast with the steady-state result of Serrano and Yosha (1993), where the opposite happens. Here informed agents who misrepresent their information are not competing with new learners who arrive fresh every period, but rather with learners who have been worn out by discounting since the beginning of the game. The informed agents end up capitalizing on their information.

In the two-sided model, noise is the dominant factor. Although force MI is present, our proof shows clearly that the problem reduces to solving the tension between forces CL and N. That is, what informed agents do is secondary to obtaining the result. And force N takes over.⁴ This echoes Wolinsky's (1990) steady-state result.

In Wolinsky (1990) and Serrano and Yosha (1993), much of the analysis focuses on limiting agents' behavior so that in each period as many leave the market as

⁴Because it prevents uninformed agents from trading, force F can overpower the other three. However, its presence signals rather extreme parameter values. We find weak sufficient conditions which guarantee that all agents trade, i.e. that force F vanishes. These are weaker than the conditions needed by Wolinsky (1990) to show existence of equilibrium.

enter it. In our one-time entry model, the market opens with a certain population of agents, no new entrants are allowed, and agents' behavior is not artificially restricted. Thus we can proceed directly from the definition of Nash equilibrium.

Another difficulty which we avoid by dropping the steady-state assumption is related to the feasibility constraints of the economy. When the steady state restriction is imposed, it is not clear what is the right notion of feasibility in the economy. It is often the case that, whatever this notion is, it will be satisfied on the equilibrium path, but violated as soon as we abandon it, due to the unbounded accumulation of agents in the market. For other criticisms of steady state models, see Gale (1987).

At a technical level, we show existence of equilibrium by adapting to our environment theorems from the literature on anonymous games [Mas-Colell (1984), Khan (1989)]. This turns out to be an exercise of interest in its own right. In those papers, payoffs depend on the total measure of players playing a given strategy, not distinguishing among types of players. We allow for a finite number of types in the payoff functions. In our treatment, the anonymity assumed in the procedure is the key, as it allows a simple strategic-form representation of an infinite-horizon game.

In a common values auction context, Pesendorfer and Swinkels (1997) show a remarkable positive result of information aggregation: the symmetric equilibrium price converges in probability to the true value of the object if and only if both supply and excess demand grow unbounded. In their result, an ingenious loser's curse argument replaces the strong conditions on the signal structure found in Wilson (1977) and Milgrom (1979, 1981). Investigating the connections between the auctions and matching and bargaining literatures seems to be an important unexplored area of research. On the other hand, the strategy adopted by the uninformed agents in our model can be interpreted as the experimentation of optimizing agents before a bandit problem [see Rothschild (1974)]. To play soft is to opt for the arm with an immediate payoff, while to play tough is to keep experimenting in order to increase learning. Bergemann and Välimäki (1996) show that, in a bandit problem between a single buyer and several sellers, the experimentation that goes on in every Markov perfect equilibrium of their model does not cause a loss in efficiency. As in their framework, the prizes associated with the two arms in our model are endogenous, since they are a function of the strategies employed by the other traders in the market. One way to interpret our results is based on the externalities caused by many buyers and many sellers trying to learn at the same time. Finally, in the literature on matching and bargaining with private values uncertainty, the positive result of connecting with the Walrasian outcome arises [see Gale (1986a, section 5)]. Our findings show that the same positive results in the common values framework are for now elusive.

The paper is organized as follows. Sections 2-5 deal specifically with the two-sided model. Section 2 gives the particulars of the model. Section 3 proves existence

of equilibrium and shows conditions under which all agents trade in finite time. Section 4 provides the information revelation result, and Section 5 the efficiency result. Section 6 contains the analysis for the one-sided version of the model. This requires separate study, as different forces are at work. Section 7 concludes with a brief discussion.

2 The Model

There are two populations of agents in the economy: sellers and buyers. There is a single indivisible good. Each seller has one unit of the good for sale, and each buyer is interested in buying one unit. Each population is an atomless continuum of measure 1.⁵

Time elapses discretely according to $t = 0, 1, 2, \dots$. All agents enter the market at the beginning of period 0. There is no entry of new agents in subsequent periods.

In period 0, each agent is randomly matched with one agent of the other population. The pair then tries to agree on a price at which to transact the good. If there is agreement, the transaction takes place, the two agents receive their payoffs and exit the market forever. If there is disagreement, no transaction takes place and the two agents remain in the market. In period 1, each of the *remaining* agents is again randomly matched with an agent of the other population. The pair tries to agree on a price, etc. This cycle is repeated infinitely many times, or until all agents have transacted and left the market. Note that there is always an equal measure of sellers and buyers remaining in the market.

The payoffs to two agents reaching agreement depend on the price at which they transact and on the state of the world. The state of the world is either H or L , and it never changes. If the state is H , all units of the good have high value in all periods; payoffs to seller and buyer are $p - c_H$ and $u_H - p$, respectively, where p is the price they agree upon. If the state is L , all units of the good have low value in all periods; payoffs to seller and buyer are $p - c_L$ and $u_L - p$, respectively.

A fraction x_S of the initial population of sellers and a fraction x_B of the initial population of buyers know the state of the world. The rest do not know the state but enter the market believing that with probability α the state is H . We assume $\alpha \in (0, 1)$, and $x_S, x_B \in [0, 1]$. That is, we first consider a two-sided information structure, where there are uninformed traders on both sides of the market. We shall consider the one-sided information case ($x_S = 1$) in section 6.

In each period, bargaining between a seller and buyer proceeds as follows. The two agents simultaneously take bargaining positions: each agent can play *soft* or

⁵We avoid the additional friction of having unmatched agents in every period. We discuss this assumption in our last section.

tough. If the seller plays soft and the buyer plays tough, they transact at price p_L . If the seller plays tough and the buyer plays soft, they transact at price p_H . If the seller and buyer both play soft, they transact at price p_M . Finally, if the seller and buyer both play tough, there is disagreement and no transaction takes place: both agents remain in the market until next period, at which point they will be matched with new opponents. The process is summarized in Figure 1.

		BUYER	
		soft	tough
SELLER	soft	$p_M - c_v, u_v - p_M$	$p_L - c_v, u_v - p_L$
	tough	$p_H - c_v, u_v - p_H$	disagreement

Figure 1. Bargaining process and outcomes; $v = H, L$.

The parameters are related as follows:

$$u_H > p_H > c_H > p_M > u_L > p_L > c_L \geq 0. \quad (1)$$

It can be seen that if the state is H , then p_H is the only price which is EPIR for both sellers and buyers; likewise, p_L is the only EPIR price for state L . For this reason we will sometimes call p_H the “right” price for state H and p_L the “right” price for state L . These are the prices which would prevail in a perfect-information setting.

The cost of delay (through disagreement) is embodied in the discount factor $\delta \in (0, 1)$, common to all agents. Perpetual disagreement entails a payoff of zero.

We turn now to individual agents’ strategies. There are three things to consider. The first is anonymity. An agent never knows the identity of his opponent and can never tell whether his opponent is informed (i.e. knows the state of the world) or uninformed. He must treat all opponents the same way. Second, we note that an agent leaves the market as soon as he (or an opponent) plays soft. Thus if he is still active in period t , this means he has played tough in periods $0, 1, \dots, t - 1$ (and met a tough-playing opponent each time). Third, it is assumed that during his stay in the market, an agent observes only the results of his own meetings. That is to say, he obtains no useful information along the way, other than the very fact that he is still in the market.

It follows from these considerations that an agent's strategy for the game is simply the number of periods he is prepared to play tough. His history at the beginning of any period t (if he is still active) can be summarized by the number t itself. He can therefore calculate right at the start of the game how long it will be optimal to play tough and when it will be optimal to play soft. In other words, he can decide from the outset how long to hold out for the most advantageous price, and when to give in. Thus the strategy space is $A \equiv \mathbb{N} \cup \{\infty\}$, where \mathbb{N} is the set of non-negative integers. An agent playing ∞ plays tough all the time.

Within this framework we seek a Nash equilibrium, a profile of strategies where each agent is maximizing his expected payoff, given the strategies of the other agents. All parameters $(\delta, x_S, x_B, \alpha, p_H, p_M, p_L, c_H, c_L, u_H, u_L)$ are common knowledge, as are all equilibrium strategies.

To summarize, a game G consists of $(\delta, x_S, x_B, \alpha, p_H, p_M, p_L, c_H, c_L, u_H, u_L)$ satisfying $\alpha, \delta \in (0, 1)$; $x_S, x_B \in [0, 1]$; $p_H, p_M, p_L, c_H, c_L, u_H, u_L \in \mathbb{R}$; inequality (1); and proceeding as described.

3 Equilibrium

An informed agent will in general act differently in the two states, whereas an uninformed agent cannot. There are thus six types of behavior to account for: informed sellers in state H (SH), informed sellers in state L (SL), uninformed sellers (S), informed buyers in state H (BH), informed buyers in state L (BL), and uninformed buyers (B). Let $K = \{\text{SH, SL, S, BH, BL, B}\}$ denote the set of possible types.

Agents belonging to the same type solve the same problem. However, they will not necessarily adopt the same strategy, as several strategies may be equally optimal. For any subset of possible strategies $X \subset A$, we denote by $\phi^k(X)$ the fraction of the initial population of type- k agents who in equilibrium play strategies contained in X . For simplicity we write $\phi^k(a) \equiv \phi^k(\{a\})$ and $\phi^k(a, b) \equiv \phi^k(\{a, \dots, b\})$, for all $a, b \in A$. In particular, $\phi^k(a, \infty)$ will signify $\phi^k(\{a, a+1, \dots\} \cup \{\infty\})$.

The set function ϕ^k has all the features of a probability measure, and we will treat it as such.⁶ The set of probability measures on A will be denoted Φ . Candidates for equilibrium will take the form $\phi = (\phi^{SH}, \phi^{SL}, \phi^S, \phi^{BH}, \phi^{BL}, \phi^B) \in \Phi^6$.

The function $\pi^k(a; \phi)$ will be the expected payoff to a type- k agent from playing strategy a (i.e. playing tough a times and playing soft in period a), given that other agents' strategies conform to ϕ . Often we will omit strategies from the argument and write simply $\pi^k(a)$.

⁶We will see in the proof of Theorem 1 that the Borel σ -algebra of A (call it \mathcal{B}) coincides with the set of subsets of A . Hence, for every $k \in K$, $\phi^k : \mathcal{B} \rightarrow \mathbb{R}$ is a Borel probability measure.

DEFINITION 1. A Nash equilibrium of a game G consists of $\phi = (\phi^{SH}, \phi^{SL}, \phi^S, \phi^{BH}, \phi^{BL}, \phi^B) \in \Phi^6$ such that, for all $a \in A$ and for all $k \in K$,

$$\phi^k(a) > 0 \quad \text{implies} \quad a \in \arg \max_{b \in A} \pi^k(b; \phi). \quad (2)$$

3.1 Expected Payoffs

We now calculate expected payoffs for each type of agent. Agents know the distributions of equilibrium strategies for the types they are likely to be paired with, and use these to calculate the probability of meeting a tough or soft opponent in any given period.

Define $\sigma^H(t) \equiv x_S \phi^{SH}(t) + (1 - x_S) \phi^S(t)$ and $\beta^H(t) \equiv x_B \phi^{BH}(t) + (1 - x_B) \phi^B(t)$. These are the measures of sellers and buyers, respectively, who play strategy t in state H . $\sigma^H(t, \infty)$ and the like are defined analogously. Consider a type-SH agent playing strategy a . In any period t , a fraction $\beta^H(t)/\beta^H(t, \infty)$ of the remaining buyers (if there are any) plays soft. Thus, for any $t \leq a$, the probability of meeting a soft buyer in period t is equal to this fraction multiplied by the probability of reaching t . The result is, quite simply, $\beta^H(t)$. With this in hand, we compute⁷

$$\begin{aligned} \pi^{SH}(a) = & (p_H - c_H) \sum_{t=0}^{a-1} \delta^t [x_B \phi^{BH}(t) + (1 - x_B) \phi^B(t)] \\ & + (p_M - c_H) \delta^a [x_B \phi^{BH}(a) + (1 - x_B) \phi^B(a)] \\ & + (p_L - c_H) \delta^a [x_B \phi^{BH}(a + 1, \infty) + (1 - x_B) \phi^B(a + 1, \infty)]. \end{aligned} \quad (3)$$

In a similar fashion we calculate, for all $a \in A$,

$$\begin{aligned} \pi^{SL}(a) = & (p_H - c_L) \sum_{t=0}^{a-1} \delta^t [x_B \phi^{BL}(t) + (1 - x_B) \phi^B(t)] \\ & + (p_M - c_L) \delta^a [x_B \phi^{BL}(a) + (1 - x_B) \phi^B(a)] \\ & + (p_L - c_L) \delta^a [x_B \phi^{BL}(a + 1, \infty) + (1 - x_B) \phi^B(a + 1, \infty)]; \end{aligned} \quad (4)$$

$$\begin{aligned} \pi^{BH}(a) = & (u_H - p_L) \sum_{t=0}^{a-1} \delta^t [x_S \phi^{SH}(t) + (1 - x_S) \phi^S(t)] \\ & + (u_H - p_M) \delta^a [x_S \phi^{SH}(a) + (1 - x_S) \phi^S(a)] \\ & + (u_H - p_H) \delta^a [x_S \phi^{SH}(a + 1, \infty) + (1 - x_S) \phi^S(a + 1, \infty)]; \end{aligned} \quad (5)$$

⁷Throughout the paper, we adopt the convention that $\sum_{t=0}^{-1}(\cdot) = 0$.

$$\begin{aligned}
\pi^{BL}(a) = & (u_L - p_L) \sum_{t=0}^{a-1} \delta^t [x_S \phi^{SL}(t) + (1 - x_S) \phi^S(t)] \\
& + (u_L - p_M) \delta^a [x_S \phi^{SL}(a) + (1 - x_S) \phi^S(a)] \\
& + (u_L - p_H) \delta^a [x_S \phi^{SL}(a + 1, \infty) + (1 - x_S) \phi^S(a + 1, \infty)];
\end{aligned} \tag{6}$$

$$\pi^S(a) = \alpha \pi^{SH}(a) + (1 - \alpha) \pi^{SL}(a); \tag{7}$$

$$\pi^B(a) = \alpha \pi^{BH}(a) + (1 - \alpha) \pi^{BL}(a). \tag{8}$$

We can also determine the profile of trade. In each period $t \in \mathbb{N}$, in state v : (i) a measure $\sigma^v(t + 1, \infty) \beta^v(t)$ of the commodity is traded at p_H ; (ii) a measure $\sigma^v(t) \beta^v(t)$ is traded at p_M ; (iii) a measure $\sigma^v(t) \beta^v(t + 1, \infty)$ is traded at p_L . The rest is carried into period $t + 1$. A measure $\sigma^v(\infty) \beta^v(\infty)$ will never be traded.

3.2 Existence

THEOREM 1. *For any game G , there exists a Nash equilibrium $\phi \in \Phi^6$.*

PROOF. The proof is a variation of Mas-Colell (1984), itself a reformulation of Schmeidler (1973).⁸ The proof is non-trivial, however, in that it resolves some of the difficulties attendant to including ∞ as a possible strategy. Also, we proceed in a slightly different manner, since here expected payoffs depend explicitly on the distribution of types (embodied in x_S and x_B), whereas in Mas-Colell they do not.

Recall that $A \equiv \mathbb{N} \cup \{\infty\}$, where $\mathbb{N} = \{0, 1, 2, \dots\}$. We define the following metric d on A .

$$\begin{aligned}
d(x, y) &= \frac{|x - y|}{(1 + x)(1 + y)} \quad \text{for } x, y \in \mathbb{N}; \\
d(x, \infty) &= d(\infty, x) = \frac{1}{1 + x} \quad \text{for } x \in \mathbb{N}; \\
d(\infty, \infty) &= 0.
\end{aligned} \tag{9}$$

⁸See Khan (1989) for an extension of Mas-Colell's result to nonmetrizable action spaces and upper-semicontinuous (rather than continuous) payoffs. Khan provides detailed proofs for both his and Mas-Colell's results.

This metric generates the following topology \mathcal{T} for A .

$$\mathcal{T} \equiv \{X \subset A \mid \text{if } \infty \in X \text{ then } \{t, t+1, \dots\} \subset X \text{ for some } t \in \mathbb{N}\}.$$

Note that all subsets of \mathbb{N} are elements of \mathcal{T} . As its definition states, \mathcal{T} also contains all sets which contain both ∞ and the numbers $t, t+1, \dots$, for any $t \in \mathbb{N}$. We assume throughout that A is endowed with this topology. Since every subset of A is either open or closed, the Borel σ -algebra of A simply consists of all subsets of A . Since A is a metric space, it is normal.

CLAIM 1. *A is compact.*

PROOF. Consider any open cover of A . One of the sets in the cover must include ∞ and therefore must also include the points $t, t+1, \dots$ for some $t \in \mathbb{N}$. So the other sets in the cover need only cover a finite number of points in A . Hence the open cover has a finite subcover, and A is compact. \square

The set Φ is the set of (Borel) probability measures on A . Since A is metric and compact, Φ is metrizable and weakly compact [see Hildenbrand (1974, p.49)]. Note that Φ is a subset of ℓ^1 , a Hausdorff topological linear space. It is straightforward to show that Φ is convex. Of course Φ^6 shares all these properties.

CLAIM 2. *For any $k \in K$, the mapping $\pi^k : A \times \Phi^6 \rightarrow \mathbb{R}$ is continuous over A .*

PROOF. Consider type-SH agents. Fix any $\phi \in \Phi^6$. Let C be any open subset of \mathbb{R} . If $\pi^{SH}(\infty; \phi) \notin C$ then the inverse image of C (projected onto A) is a subset of \mathbb{N} , hence an element of \mathcal{T} , hence open. Now suppose $\pi^{SH}(\infty; \phi) \in C$. A glance at (3) shows that $\lim_{a \rightarrow \infty} \pi^{SH}(a) = \pi^{SH}(\infty)$. It follows that for t large enough, $\pi^{SH}(a; \phi) \in C$ for all $a \geq t$. Consequently the inverse image of C (projected onto A) contains the points $t, t+1, \dots$ as well as ∞ , and therefore is an element of \mathcal{T} and an open set. The same holds for the other types. This establishes continuity over A . \square

For all $k \in K$, define the best-response correspondence $\psi^k : \Phi^6 \rightrightarrows A$ as $\psi^k(\phi) = \arg \max_{b \in A} \pi^k(b; \phi)$. Since A is compact and π^k is continuous over A , we have by Weierstrass' theorem that for any $\phi \in \Phi^6$, $\psi^k(\phi)$ is nonempty. For all $k \in K$, define $\theta^k : \Phi^6 \rightrightarrows \Phi$ as $\theta^k(\phi) = \{\mu \in \Phi \mid \mu(\psi^k(\phi)) = 1\}$. To be an element of $\theta^k(\phi)$, a probability measure μ must put positive mass only on best-response strategies. $\theta^k(\phi)$ is in a sense a set of optimal mixed strategies (which also includes pure strategies). We must keep in mind, however, that these “probabilities” are really fractions of the initial type- k population. Since $\psi^k(\phi)$ is nonempty for any $\phi \in \Phi^6$, so is $\theta^k(\phi)$. It is straightforward to show that $\theta^k(\phi)$ is also convex-valued.

CLAIM 3. *For all $k \in K$, the correspondence $\theta^k(\phi)$ is upper-hemicontinuous.*

PROOF. Consider a sequence $\{\phi_n\}_{n=1}^\infty$ such that $\phi_n \in \Phi^6$ for all n and $\phi_n \xrightarrow{w} \bar{\phi}$ (where \xrightarrow{w} signifies weak convergence). Consider another sequence $\{\mu_n\}_{n=1}^\infty$ such that $\mu_n \in \theta^k(\phi_n)$ for all n and $\mu_n \xrightarrow{w} \bar{\mu}$. Since Φ is weakly compact, we know $\bar{\mu} \in \Phi$. We need to show $\bar{\mu} \in \theta^k(\bar{\phi})$, i.e. $\bar{\mu}(\psi^k(\bar{\phi})) = 1$.

Since A is normal and $\mu_n(\psi^k(\phi_n)) = 1$ for all n , we have $\bar{\mu}(\limsup_n \psi^k(\phi_n)) = 1$ [see Khan (1989, Lemma 2)].

If we fix a and let ϕ vary, $\pi^k(a; \phi)$ becomes a bounded linear functional on Φ^6 , hence an element of the dual space $(\Phi^6)'$. Since $\phi_n \xrightarrow{w} \bar{\phi}$, we have $\pi^k(a; \phi_n) \rightarrow \pi^k(a; \bar{\phi})$ for all $a \in A$.

Suppose $a \in A \setminus \psi^k(\bar{\phi})$. Then $\pi^k(a; \bar{\phi}) < \pi^k(b; \bar{\phi})$ for some $b \in A$. There must be m such that $\pi^k(a; \phi_n) < \pi^k(b; \phi_n)$ for all $n > m$. This means that $a \in \liminf_n (A \setminus \psi^k(\phi_n))$, and thus $(A \setminus \psi^k(\bar{\phi})) \subset \liminf_n (A \setminus \psi^k(\phi_n))$. And since $\liminf_n (A \setminus \psi^k(\phi_n)) \cap \limsup_n \psi^k(\phi_n) = \emptyset$, it follows that $\limsup_n \psi^k(\phi_n) \subset \psi^k(\bar{\phi})$.

By monotonicity of probability measures, $\bar{\mu}(\limsup_n \psi^k(\phi_n)) \leq \bar{\mu}(\psi^k(\bar{\phi}))$. Consequently $\bar{\mu}(\psi^k(\bar{\phi})) = 1$. \square

Finally, define $\theta : \Phi^6 \Rightarrow \Phi^6$ as $\theta(\phi) = \times_{k \in K} \theta^k(\phi)$. We know that θ^k is upper-hemicontinuous, convex-valued and nonempty-valued, for all $k \in K$. These properties extend to θ . Also, Φ^6 is a nonempty, weakly compact subset of a Hausdorff topological linear space. Hence θ meets the conditions of the Fan-Glicksberg fixed point theorem [Fan (1952), Glicksberg (1952)], and there exists $\phi \in \Phi^6$ such that $\phi \in \theta(\phi)$. Such a ϕ satisfies (2) and is therefore a Nash equilibrium of the game. \square

3.3 Characterization

In the rest of this section, we show that in equilibrium, under certain conditions, the game ends in finite time, with all agents trading.

CLAIM 4. *Suppose $\phi \in \Phi^6$ is a Nash equilibrium of a game G . Then for all $a \in \mathbb{N}$,*

$$(i) \quad \phi^{SH}(a) > 0 \quad \text{implies} \quad \phi^{BH}(0, a-1) = \phi^B(0, a-1) = 1;$$

$$(ii) \quad \phi^{BL}(a) > 0 \quad \text{implies} \quad \phi^{SL}(0, a-1) = \phi^S(0, a-1) = 1.$$

PROOF. We prove the claim for type-SH agents. The proof for type-BL agents is similar.

From (2), $\phi^{SH}(a) > 0$ implies $\pi^{SH}(a) \geq \pi^{SH}(b)$ for all $b \in A$. From (3), we see that $\pi^{SH}(\infty) \geq \pi^{SH}(b)$ for all $b \in A$, and that for any $a \in \mathbb{N}$, $\pi^{SH}(a) = \pi^{SH}(\infty)$ if and only if $\phi^{BH}(0, a-1) = \phi^B(0, a-1) = 1$. \square

The claim states that no one ever meets a type-SH or type-BL agent playing soft. A type-SH or type-BL agent knows he will get a negative payoff if ever he plays soft in a meeting, so he plays tough for as long as his potential opponents are in the market. If his strategy is $a \in \mathbb{N}$, then it must yield him the same payoff as playing ∞ would, and moreover it must be a number large enough that it does not provide any of his potential opponents with an incentive to deviate from their own strategies. In other words, it must be outcome-equivalent to playing ∞ . There is no loss of generality, therefore, in assuming that all type-SH and type-BL agents play ∞ in equilibrium, and we shall do so henceforth.

CLAIM 5. *Suppose $\phi \in \Phi^6$ is a Nash equilibrium of a game G . Then ϕ^{SL} and ϕ^{BH} have finite support and $\phi^{SL}(\infty) = \phi^{BH}(\infty) = 0$.*

PROOF. We prove the claim for type-SL agents. The proof for type-BH agents is similar.

Suppose ϕ^{SL} has infinite support and/or $\phi^{SL}(\infty) > 0$. This means that π^{SL} is maximized at ∞ (or at infinitely many points, which amounts to the same thing, since the payoff function is continuous at ∞). So $\pi^{SL}(\infty) \geq \pi^{SL}(a)$ for all $a \in \mathbb{N}$. By Claim 4 we set $\phi^{BL}(\infty) = 1$. From (4) we calculate

$$\begin{aligned} [\pi^{SL}(\infty) - \pi^{SL}(a)]/\delta^a &= (p_H - p_M)(1 - x_B)\phi^B(a) \\ &\quad + (p_H - c_L)(1 - x_B) \sum_{t=a+1}^{\infty} \delta^{t-a}\phi^B(t) \\ &\quad - (p_L - c_L)(1 - x_B)\phi^B(a+1, \infty) \\ &\quad - (p_L - c_L)x_B. \end{aligned} \tag{10}$$

The first two terms on the right-hand side of (10) necessarily approach zero as a becomes large, since $\lim_t \phi^B(t) = 0$ (even if $\phi^B(\infty) > 0$). The rest of the expression is negative, and the fourth term does not approach zero as a increases. Thus the entire expression becomes negative for $a \in \mathbb{N}$ large enough, which contradicts the supposition. \square

Claim 5 states that there is a finite period in which all remaining informed buyers leave the market if the state is H , and all remaining informed sellers leave the market if the state is L . Roughly, the explanation is as follows. A type-SL agent knows that his opposition consists of: (i) some buyers who play tough all the time (all type-BL agents and possibly some type-B agents as well); and (ii) [as long as $\phi^B(\infty) < 1$] some who eventually play soft. By a large-numbers reasoning, the members of the

latter group drop out of the market at a faster rate than do those of the former. The probability of meeting a buyer playing soft eventually becomes small enough that it is no longer worthwhile for the type-SL agent to stay in the market. The reasoning for type-BH agents is parallel.

CLAIM 6. Suppose $\phi \in \Phi^6$ is a Nash equilibrium of a game G . Define $c_{HL} \equiv \alpha c_H + (1 - \alpha)c_L$ and $u_{HL} \equiv \alpha u_H + (1 - \alpha)u_L$, and consider the following two inequalities:

$$(p_L - c_{HL})(1 - x_B) + (1 - \alpha)(p_L - c_L)x_B > 0; \quad (11)$$

$$(u_{HL} - p_H)(1 - x_S) + \alpha(u_H - p_H)x_S > 0. \quad (12)$$

If (11) and/or (12) is satisfied, then ϕ^S and ϕ^B have finite support and $\phi^S(\infty) = \phi^B(\infty) = 0$.

PROOF. The proof consists of the following two symmetric arguments: (i) if ϕ^S has infinite support and/or $\phi^S(\infty) > 0$, then (11) cannot hold and $\phi^B(\infty) > 0$; (ii) if ϕ^B has infinite support and/or $\phi^B(\infty) > 0$, then (12) cannot hold and $\phi^S(\infty) > 0$.

Suppose ϕ^S has infinite support and/or $\phi^S(\infty) > 0$. Then $\pi^S(\infty) \geq \pi^S(a)$ for all $a \in \mathbb{N}$. From Claim 4 we set $\phi^{BL}(\infty) = 1$. Let a be large enough that all type-BH agents have left the market (see Claim 5). From (3), (4) and (7) we then calculate

$$\begin{aligned} [\pi^S(\infty) - \pi^S(a)]/\delta^a &= (p_H - p_M)(1 - x_B)\phi^B(a) \\ &\quad + (p_H - c_{HL})(1 - x_B) \sum_{t=a+1}^{\infty} \delta^{t-a}\phi^B(t) \\ &\quad - (p_L - c_{HL})(1 - x_B)\phi^B(a + 1, \infty) \\ &\quad - (1 - \alpha)(p_L - c_L)x_B. \end{aligned} \quad (13)$$

As a becomes large, the first two terms on the right-hand side of (13) necessarily approach zero. The rest of the expression approaches $-(p_L - c_{HL})(1 - x_B)\phi^B(\infty) + (1 - \alpha)(p_L - c_L)x_B$. So for $\pi^S(\infty) - \pi^S(a) \geq 0$ to hold for all $a \in \mathbb{N}$ we must have

$$(p_L - c_{HL})(1 - x_B)\phi^B(\infty) + (1 - \alpha)(p_L - c_L)x_B \leq 0. \quad (14)$$

This means $\phi^B(\infty)$ must be positive. It also means (11) cannot hold. This establishes part (i). Now suppose ϕ^B has infinite support and/or $\phi^B(\infty) > 0$. By a similar argument, this requires

$$(u_{HL} - p_H)(1 - x_S)\phi^S(\infty) + \alpha(u_H - p_H)x_S \leq 0. \quad (15)$$

This means that $\phi^S(\infty)$ must be positive, and also that (12) cannot hold, and part (ii) is established. The two arguments together prove the claim. \square

Claim 6 states that if either (11) or (12) is met (or both), then there is a finite period in which all remaining uninformed agents leave the market. What might drive a type-S agent to play tough indefinitely is not the hope of meeting a buyer playing soft: this may be a motivation at first, but eventually it disappears (the argument is similar to that for type-SL agents above). Rather, it is the fear of state H , where playing soft means getting a negative payoff. Consider (14). For the inequality to hold, $(p_L - c_{HL})$ must be negative, $\phi^B(\infty)$ must be positive, and their product must be of sufficient magnitude. If $p_L - c_{HL} < 0$, the parameters are unfavorable: low price, high costs, and a high prior probability that the state is H . If $\phi^B(\infty)$ is small or zero, then eventually the type-B population (in either state) becomes small relative to the type-BL population (in state L): an uninformed seller reasons that if he is still meeting tough buyers at such a time, then the state is probably L .⁹ If $\phi^B(\infty)$ is large enough, on the other hand, the fear that the state might be H will persist. The reasoning for type-B agents is parallel.

It now becomes clear that if at least one of (11) or (12) holds, then the game must end in finite time, with all agents having traded. If the state is H , there is a period in which all remaining buyers play soft; if the state is L , there is a period in which all remaining sellers play soft. We state this formally in Theorem 2.

THEOREM 2. *Suppose (11) and/or (12) holds. Then the game G must end in finite time.*

PROOF. Define $T^k \equiv \min\{a \in \mathbb{N} | \phi^k(0, a) = 1\}$. By Claims 5 and 6, the periods T^{SL} , T^{BH} , T^S and T^B are well-defined. In state H the game ends in period $T_H \equiv \max\{T^{BH}, T^B\}$, since all remaining buyers play soft in that period; in state L it ends in period $T_L \equiv \max\{T^{SL}, T^S\}$, since all remaining sellers play soft then. \square

If neither (11) nor (12) is met, there are equilibria in which some agents do not trade. We show here an example.

EXAMPLE 1. *Suppose neither (11) nor (12) holds. Then the following is an equilibrium: agents of types SH, S, BL, and B all play ∞ ; agents of types SL and BH all play 0.*

⁹From a type-S agent's perspective, the probability of the state being H conditional on period t being reached is $\alpha\beta^H(t, \infty) / [\alpha\beta^H(t, \infty) + (1 - \alpha)\beta^L(t, \infty)]$. This is all foreseeable from the start and is already incorporated in the derivation of π^S .

In this example, all trade takes place in period 0. Afterwards, all remaining agents play tough forever. The fraction of agents not trading is $1 - x_B$ in state H and $1 - x_S$ in state L . These are upper bounds: in any equilibrium, the amount untraded never exceeds $1 - x_B$ in state H and $1 - x_S$ in state L .

4 Information Revelation

In this section, we analyze the degree of information revelation in Nash equilibria of the game as the economy becomes nearly frictionless. As noted in the introduction, the only friction in this model is discounting. So we examine this property of equilibria as discounting is gradually removed.

Formally, we have in mind the following. Fix all parameters except δ , and consider any sequence $\{\delta_n\}_{n=1}^\infty$ satisfying $\delta_n \in (0, 1)$ for all n and $\lim_n \delta_n = 1$. Let G_n be the game consisting of δ_n and the fixed parameters. Finally, consider any sequence $\{\phi_n\}_{n=1}^\infty$ such that $\phi_n \in \Phi^6$ is a Nash equilibrium of G_n for all n . We will show that any such sequence fails to exhibit full revelation of information as $n \rightarrow \infty$. For ease of exposition, we will sometimes speak of limits as $\delta \rightarrow 1$ rather than as $n \rightarrow \infty$.

Consider an uninformed agent who is still in the market in period $t > 0$. This agent has learned no “hard facts” (e.g. opponent’s type, outcomes of other meetings), and he never will. Nonetheless, the fact that he is still active in period t , combined with his knowledge of other agents’ equilibrium strategies, may convey some information to him: given ϕ , meeting t successive tough opponents may be more consistent with the state being H than with the state being L , or vice versa. This information erodes the disadvantage he has vis-à-vis informed agents. The longer he remains in the market, the more accurate the information; and the closer δ is to 1, the longer he can afford to play tough and remain in the market. So it is natural to ask how far this information revelation can go. In the limit, as δ goes to 1, does enough information get revealed to uninformed agents that they end up trading at EPIR prices, i.e. at prices they would have agreed to had they known the state of the world with certainty?

Alternatively, one can have as a benchmark for the analysis the *fully revealing rational expectations equilibrium* (FRREE), a concept associated with centralized trade where all information is revealed thanks to the auctioneer’s choice of the equilibrium price function. We can then compare a sequence of equilibria of the present model of decentralized trade with the FRREE of a model of centralized trade under the same conditions of asymmetric information. It should be clear, given the payoffs described above, that a rational expectations equilibrium of this economy is necessarily a FRREE with a price function satisfying $p(H) \in [c_H, u_H]$,

$p(L) \in [c_L, u_L]$, and where all agents trade.

Let us return now to the decentralized model. In answering the information revelation question, we must verify whether or not the volume of trade taking place at the non-EPIR prices converges to 0 as $\delta \rightarrow 1$.

In the rest of the paper we require several limits to exist. Since the various items whose limits are taken are all defined on compact sets, every sequence $\{\phi_n\}$ has a subsequence for which these limits exist. It should be understood that we restrict our attention to such subsequences.¹⁰ We will index various quantities by n , since they will generally take on different values with each equilibrium in the sequence.

Denote by $VW^v(t)$ the volume of trade taking place at the wrong (non-EPIR) prices in period t in state $v = H, L$. Let $VW^v \equiv \sum_{t \in \mathbb{N}} VW^v(t)$ be the total volume of trade at those prices.¹¹ In state H , trade at a wrong price (p_M or p_L) occurs if and only if a seller plays soft. Therefore

$$VW^H(t) = \sigma^H(t)\beta^H(t, \infty)$$

Summing over all $t \in \mathbb{N}$, and repeating for state L , we have

$$VW^H = \sum_{t \in \mathbb{N}} \sigma^H(t)\beta^H(t, \infty); \quad (16)$$

$$VW^L = \sum_{t \in \mathbb{N}} \sigma^L(t, \infty)\beta^L(t). \quad (17)$$

Let VU^v measure the untraded amount in the economy in state v . This is

$$VU^H = \sigma^H(\infty)\beta^H(\infty); \quad (18)$$

$$VU^L = \sigma^L(\infty)\beta^L(\infty). \quad (19)$$

CLAIM 7. *Consider a sequence of games $\{G_n\}_{n=1}^\infty$ for which all parameters except δ are fixed, and where $\delta_n \rightarrow 1$, $\delta_n \in (0, 1)$ for all n ; and consider a corresponding sequence of Nash equilibria $\{\phi_n\}_{n=1}^\infty$ for which the relevant limits exist. Then we cannot have $\lim_n VU_n^H = \lim_n VU_n^L = \lim_n VW_n^H = \lim_n VW_n^L = 0$.*

¹⁰Note that a sequence $\{\phi_n\}$ does not necessarily converge pointwise to an element of Φ^6 . In that respect, all we can guarantee is that it must have a subsequence which converges *weakly* to a point $\bar{\phi} \in \Phi^6$, since Φ^6 is weakly compact (see proof of Theorem 1). This is irrelevant, however: we are not interested in the properties of any limiting point $\bar{\phi}$, but in those of ϕ_n for n finite and arbitrarily large.

¹¹Here it is important to distinguish between $\sum_{t \in \mathbb{N}}$ and $\sum_{t=0}^\infty$. To illustrate: $\sum_{t \in \mathbb{N}} \phi^k(t) + \phi^k(\infty) = \sum_{t=0}^\infty \phi^k(t)$. For discounted sums there is no difference.

PROOF. Suppose $\lim_n VU_n^H = 0$. From (16), this means, among other things, that

$$\lim_n \sum_{t \in \mathbb{N}} \phi_n^S(t) \phi_n^B(t, \infty) = 0. \quad (20)$$

Similarly, $\lim_n VW_n^L = 0$ implies, via (17), that

$$\lim_n \sum_{t \in \mathbb{N}} \phi_n^S(t, \infty) \phi_n^B(t) = 0. \quad (21)$$

And $\lim_n VU_n^H = 0$ or $\lim_n VW_n^L = 0$ implies, via (18) or (19), that

$$\lim_n \phi_n^S(\infty) \phi_n^B(\infty) = 0. \quad (22)$$

But (20), (21) and (22) cannot all hold, since for any n

$$\begin{aligned} & \sum_{t \in \mathbb{N}} \phi_n^S(t) \phi_n^B(t, \infty) + \sum_{t \in \mathbb{N}} \phi_n^S(t, \infty) \phi_n^B(t) + \phi_n^S(\infty) \phi_n^B(\infty) \\ &= 1 + \sum_{t \in \mathbb{N}} \phi_n^S(t) \phi_n^B(t) \geq 1. \end{aligned} \quad (23)$$

This proves the claim. \square

The proof is illustrated by means of Figure 2, where we have omitted the subscript n . The matrix carries all the products $\phi^S(a) \phi^B(b)$, with $a, b \in A$. These give an account of trade between uninformed agents: e.g. the volume of trade between type-S agents who play strategy 3 and type-B agents who play strategy 5 is $(1 - x_S)(1 - x_B) \phi^S(3) \phi^B(5)$; these trades take place in period 3 at price p_L , since $3 < 5$.¹² All entries are of course non-negative. Region 1 represents trade at p_H , region 2 represents trade at p_M , region 3 represents trade at p_L , and region 4 represents no trade. Essentially, $\lim_n VW_n^H = 0$ requires the sum of all entries in regions 2 and 3 to go to 0 as $n \rightarrow \infty$; likewise $\lim_n VW_n^L = 0$ requires the sum of all entries in regions 1 and 2 to approach 0; and $\lim_n VU_n^H = 0$ (or $\lim_n VU_n^L = 0$) requires the entry in region 4 to approach 0. But the sum of all entries is exactly 1, for any n . The fact that region 2 is counted twice explains the second line of (23).

¹²Some of these agents meet in periods 0, 1 and 2, but those meetings do not result in trade.




Figure 2. Trade between uninformed agents.

REMARK. Claim 7 follows from the structure of the model, not from equilibrium considerations. In other words, it is impossible to construct a strategy profile, even a non-equilibrium one, in which all agents trade at the right prices in both states of the world. This results directly from the presence of uninformed agents on both sides of the market, so the remark applies to the two-sided case only. Equilibrium considerations are invoked in Theorem 3 below. \square

Our focus here is ex-post individual rationality. Under quite general conditions, it is not attained in the limit as $\delta \rightarrow 1$.

THEOREM 3. *Consider a sequence of games $\{G_n\}_{n=1}^\infty$ for which all parameters except δ are fixed, and where $\delta_n \rightarrow 1$, $\delta_n \in (0, 1)$ for all n ; and consider a corresponding sequence of Nash equilibria $\{\phi_n\}_{n=1}^\infty$ for which the relevant limits exist. Suppose (11) and/or (12) holds. Then we cannot have $\lim_n VW_n^H = \lim_n VW_n^L = 0$.*

PROOF. This follows directly from Theorem 2 and Claim 7. \square

The equilibrium results from the previous section (summarized in Theorem 2) tell us when the entry in region 4 (Figure 2) is 0 for all n . That allows us to focus the logic of Claim 7 on the other three regions, i.e. those where trade does occur. So under either of these sufficient conditions, we can say that information is imperfectly revealed in equilibrium.

It is worth emphasizing that (11) and (12) are not overly restrictive. As mentioned in an earlier footnote, the condition “(11) and/or (12)” is weaker than what Wolinsky (1990) requires for existence of a steady-state equilibrium.

5 Efficiency

We turn now to a full-fledged welfare analysis of our model. This is an important question, quite independent of the information revelation issue. Even if the information does not get revealed from the informed to the uninformed agents (as we have shown), the outcome may result in a pure transfer of surplus from those who lack information to those who have it. In such a scenario, one could not conclude that society as a whole is worse off. We shall show that this is not the case. That is, the lack of information revelation of our equilibrium sequences is accompanied by their asymptotic (as $\delta \rightarrow 1$) interim incentive inefficiency.

Again we consider any sequence of discount factors $\{\delta_n\}_{n=1}^\infty$ converging to 1, and any corresponding sequence of Nash equilibria $\{\phi_n\}_{n=1}^\infty$ for which all relevant limits exist.

First, it is obvious that those equilibrium sequences where $\lim_n VU_n^H > 0$ or $\lim_n VU_n^L > 0$ exhibit an interim inefficiency, as a non-negligible fraction of the available gains from trade is not realized. We concentrate, therefore, on those equilibrium sequences for which $\lim_n VU_n^H = \lim_n VU_n^L = 0$.

Recall that $T_n^k \equiv \min\{a \in \mathbb{N} | \phi_n^k(0, a) = 1\}$, where everything is now indexed by n . Also let $T_n \equiv \max\{T_n^{SL}, T_n^S, T_n^{BH}, T_n^B\}$.

CLAIM 8. *Consider a sequence of games $\{G_n\}_{n=1}^\infty$ for which all parameters except δ are fixed, and where $\delta_n \rightarrow 1$, $\delta_n \in (0, 1)$ for all n ; and consider a corresponding sequence of Nash equilibria $\{\phi_n\}_{n=1}^\infty$ for which the relevant limits exist. Then $\lim_n \delta_n^{T_n} < 1$.*

PROOF. Suppose $\lim_n \delta_n^{T_n} = 1$.

STEP 1. We first show that $\lim_n \phi_n^j(T_n^k) = \lim_n \phi_n^j(T_n^k + 1) = 0$ for $(j, k) \in \{(SL, B), (S, BH), (S, B), (BH, S), (B, SL), (B, S)\}$. We calculate

$$\begin{aligned} \lim_n [\pi_n^{SL}(T_n^{SL} + 1) - \pi_n^{SL}(T_n^{SL})] = \\ (1 - x_B) \left[(p_H - p_M) \lim_n \phi_n^B(T_n^{SL}) + (p_M - p_L) \lim_n \phi_n^B(T_n^{SL} + 1) \right]. \end{aligned}$$

This must be non-positive for optimality of T_n^{SL} . It follows that $\lim_n \phi_n^B(T_n^{SL}) = \lim_n \phi_n^B(T_n^{SL} + 1) = 0$ in equilibrium. The other cases are proved similarly.

STEP 2. Let $\hat{\pi}_n^k(v)$ denote an uninformed agent's equilibrium expected payoff for state $v = H, L$; let $\hat{\pi}_n^k$ denote an informed agent's equilibrium expected payoff. For feasibility we must have, for all n ,

$$x_S \hat{\pi}_n^{SH} + (1 - x_S) \hat{\pi}_n^S(H) + x_B \hat{\pi}_n^{BH} + (1 - x_B) \hat{\pi}_n^B(H) \leq u_H - c_H; \quad (24)$$

$$x_S \hat{\pi}_n^{SL} + (1 - x_S) \hat{\pi}_n^S(L) + x_B \hat{\pi}_n^{BL} + (1 - x_B) \hat{\pi}_n^B(L) \leq u_L - c_L; \quad (25)$$

since $u_v - c_v$ is the available surplus in state $v = H, L$.

Using our previous results (including Step 1), and continuing to suppose that $\lim_n \delta_n^{T_n} = 1$, we calculate

$$\lim_n \hat{\pi}_n^{SH} = \lim_n \pi_n^{SH}(\infty) = (p_H - c_H);$$

$$\lim_n \hat{\pi}_n^{BL} = \lim_n \pi_n^{BL}(\infty) = (u_L - p_L);$$

$$\lim_n \hat{\pi}_n^{SL} = \lim_n \pi_n^{SL}(T_n^{SL}) = (p_L - c_L) + (p_H - p_L)(1 - x_B) \lim_n \phi_n^B(0, T_n^{SL} - 1);$$

$$\lim_n \hat{\pi}_n^{BH} = \lim_n \pi_n^{BH}(T_n^{BH}) = (u_H - p_H) + (p_H - p_L)(1 - x_S) \lim_n \phi_n^S(0, T_n^{BH} - 1);$$

$$\lim_n \hat{\pi}_n^S(L) = \lim_n \pi_n^{SL}(T_n^S) = (p_L - c_L) + (p_H - p_L)(1 - x_B) \lim_n \phi_n^B(0, T_n^S - 1);$$

$$\lim_n \hat{\pi}_n^B(H) = \lim_n \pi_n^{BH}(T_n^B) = (u_H - p_H) + (p_H - p_L)(1 - x_S) \lim_n \phi_n^S(0, T_n^B - 1);$$

$$\begin{aligned} \lim_n \hat{\pi}_n^S(H) &= \lim_n \pi_n^{SH}(T_n^S) = (p_H - c_H) \\ &\quad + (p_H - p_L) \left[x_B \lim_n \phi_n^{BH}(0, T_n^S - 1) + (1 - x_B) \lim_n \phi_n^B(0, T_n^S - 1) - 1 \right]; \end{aligned}$$

$$\begin{aligned} \lim_n \hat{\pi}_n^B(L) &= \lim_n \pi_n^{BL}(T_n^B) = (u_L - p_L) \\ &\quad + (p_H - p_L) \left[x_S \lim_n \phi_n^{SL}(0, T_n^B - 1) + (1 - x_S) \lim_n \phi_n^S(0, T_n^B - 1) - 1 \right]. \end{aligned}$$

Substitution into (24) and (25) yields

$$\begin{aligned}
x_B & \left[\lim_n \phi_n^{BH}(0, T_n^S - 1) + \lim_n \phi_n^S(0, T_n^{BH} - 1) \right] \\
& + (1 - x_B) \left[\lim_n \phi_n^B(0, T_n^S - 1) + \lim_n \phi_n^S(0, T_n^B - 1) \right] \leq 1;
\end{aligned} \tag{26}$$

$$\begin{aligned}
x_S & \left[\lim_n \phi_n^{SL}(0, T_n^B - 1) + \lim_n \phi_n^B(0, T_n^{SL} - 1) \right] \\
& + (1 - x_S) \left[\lim_n \phi_n^S(0, T_n^B - 1) + \lim_n \phi_n^B(0, T_n^S - 1) \right] \leq 1.
\end{aligned} \tag{27}$$

The values of the expressions in (26) and (27) depend in part on the ordering of the four T_n^k 's. Clearly each expression in brackets is at least 1. Observe also that, no matter what the discount factor is, no agent will play tough when he knows that his remaining opponents will all play tough forever. This allows us to say, for all n , that $T_n^{SL} \leq T_n^B + 1$, $T_n^{BH} \leq T_n^S + 1$, and $|\max\{T_n^{SL}, T_n^S\} - \max\{T_n^{BH}, T_n^B\}| \leq 1$. The reader can then verify that at least one of the expressions in brackets is equal to 2, hence (26) and (27) cannot both hold. \square

The significance of this is that the uninformed agents' learning strategy ends up having a positive cost. The rate at which they increase their sampling of opponents as $\delta \rightarrow 1$ overcomes the rate at which learning becomes cheaper. This is what ultimately causes the inefficiency of the market, as we show presently.

Consider state L . As noted, the total surplus available in the economy is $u_L - c_L$: each transaction is a split of this amount, and there is a measure 1 of possible transactions. Let $V^L(t)$ denote the volume of trade in period t . In each period, then, the departing agents take with them $(u_L - c_L)V^L(t)$ collectively. The discounted sum of these payoffs is a measure of welfare for this state:

$$W_n^L \equiv (u_L - c_L) \sum_{t=0}^{\infty} \delta_n^t V_n^L(t). \tag{28}$$

If this is less than $u_L - c_L$, then the outcome is Pareto dominated by a split in period 0, and we call the outcome inefficient. W_n^H is defined analogously.

THEOREM 4. *Consider a sequence of games $\{G_n\}_{n=1}^{\infty}$ for which all parameters except δ are fixed, and where $\delta_n \rightarrow 1$, $\delta_n \in (0, 1)$ for all n ; and consider a corresponding sequence of Nash equilibria $\{\phi_n\}_{n=1}^{\infty}$ for which the relevant limits exist. Then we cannot have both $\lim_n W_n^H = u_H - c_H$ and $\lim_n W_n^L = u_L - c_L$.*

PROOF. Suppose the statement is untrue. Then we must have $\lim_n \sum_{t=0}^{\infty} \delta_n^t V_n^H(t) = \lim_n \sum_{t=0}^{\infty} \delta_n^t V_n^L(t) = 1$. In turn this requires the existence of a sequence $\{\lambda_n\}_{n=1}^{\infty}$ such that $\lim_n \delta_n^{\lambda_n} = 1$ and $\lim_n \sum_{t=0}^{\lambda_n} V_n^H(t) = \lim_n \sum_{t=0}^{\lambda_n} V_n^L(t) = 1$. In state L we have

$$\sum_{t=0}^{\lambda_n} V_n^L(t) = 1 - \sigma_n^L(\lambda_n + 1, \infty) \beta_n^L(\lambda_n + 1, \infty).$$

Hence we need

$$\begin{aligned} & \lim_n [x_S \phi_n^{SL}(\lambda_n + 1, \infty) + (1 - x_S) \phi_n^S(\lambda_n + 1, \infty)] \\ & \quad \times [x_B \phi_n^{BL}(\lambda_n + 1, \infty) + (1 - x_B) \phi_n^B(\lambda_n + 1, \infty)] = 0. \end{aligned}$$

Using Claim 4 to simplify, and repeating the analysis for state H , we obtain

$$\lim_n \phi_n^{SL}(\lambda_n + 1, \infty) = \lim_n \phi_n^S(\lambda_n + 1, \infty) = 0; \quad (29)$$

$$\lim_n \phi_n^{BH}(\lambda_n + 1, \infty) = \lim_n \phi_n^B(\lambda_n + 1, \infty) = 0. \quad (30)$$

For one of these four types, $T_n^k = T_n$. Let us say SL is this type. Using our results so far (and continuing to suppose $\lim_n \delta_n^{\lambda_n} = 1$), we have

$$\lim_n [\pi_n^{SL}(T_n^{SL}) - \pi_n^{SL}(\lambda_n + 1)] = (p_L - c_L) \left(\lim_n \delta_n^{T_n} - 1 \right), \quad (31)$$

which must be non-negative by optimality of T_n^{SL} at each n . This requires $\lim_n \delta_n^{T_n} = 1$, in contradiction to the result of Claim 8. The same thing happens if we try type S, BH or B. \square

Basically, the market is asymptotically efficient only if (in the limit) all trade takes place before discounting starts to “bite”. But someone is always willing to play tough until T_n , at which point discounting does bite. This is not optimal, since the agent is incurring a substantial delay cost while not significantly bettering his chances of meeting a soft opponent.

To summarize so far, this is a game of transferable utility, so the Pareto frontier is completely characterized by payoffs that sum to $u_v - c_v$, for state $v = H, L$. Our equilibria do not approach this frontier in the limit.

A Pareto improving allocation can always be implemented by a planner, even if the planner cannot distinguish between informed and uninformed agents.¹³ Each agent is asked to signal his type, and either receives the payoff that the allocation prescribes for the type he signals, or is punished, depending on whether his signal makes sense or not. For example, suppose a type-S agent claims he is type-SH. If the state turns out to be L , then he is caught lying, since no type-SH agent is present in that state, and he is punished. If we make the punishments harsh enough, then the mechanism will induce truth-telling. Our equilibria, therefore, are asymptotically interim *incentive* inefficient.

6 The One-Sided Information Case

So far in the analysis, there have been both uninformed sellers and uninformed buyers. This may be appropriate for some financial and real estate markets where the value of all units traded may depend on some event outside the marketplace, the outcome of which is known to some agents and unknown to others (or forecast more accurately by some agents than by others).

In many economic situations, all agents on one side of the market are informed. There can be uncertainty about the quality of the good itself. We usually think of sellers of commodities as knowing the quality of their product, and of buyers as perhaps lacking information in this regard. In this section, we reprise the analysis of sections 3, 4 and 5, this time with $x_S = 1$ and $x_B \in [0, 1)$. We call this *one-sided information*.

6.1 Equilibrium

Most of the results of Section 3 are easily adaptable to this case. In the expected payoff functions (3)-(8), we eliminate the $\phi^S(\cdot)$ terms and set $x_S = 1$. The existence result (Theorem 1) still holds, since the expected payoff functions π^k are still continuous over A in the new configuration. Claim 4 goes through as before, and again we can say without loss of generality that $\phi^{SH}(\infty) = \phi^{BL}(\infty) = 1$ in equilibrium.

Type-BH agents know their opponents are all informed sellers who play tough all the time. The payoff is $\pi^{BH}(a) = (p_H - c_H)\delta^a$, and the optimal strategy is 0: the inevitable should not be delayed. So $\phi^{BH}(0) = 1$ in equilibrium.

For type-SL agents, we appeal to Claim 5. It tells us that there is a finite period T^{SL} in which all remaining type-SL agents play soft.

After the last type-SL agents have left the market, any remaining type-B agents know that the state is H and that there is no longer any chance of meeting a soft

¹³We assume the planner can distinguish between sellers and buyers.

opponent. Hence $\phi^B(0, T^{SL} + 1) = 1$; that is, $T^B \leq T^{SL} + 1$. [Likewise, after the last type-B agents have left the market, any remaining type-SL agents know that the opposition now consists exclusively of informed (tough) buyers; hence $T^{SL} \leq T^B + 1$.]

The game ends in finite time, with all agents trading. In state H it ends in period T^B ; in state L it ends in period T^{SL} . There are no sufficient conditions analogous to (11) and (12). The statement is true for all parameter values consistent with the definition of G (but with $x_S = 1$).

6.2 Information Revelation

Just as in the two-sided case, information about the state of the world cannot be fully revealed in both states. It *will* be fully revealed in state H , since all trade in that state takes place at p_H , the EPIR price. In state L , however, there will always be a sizeable volume of trade occurring at one or the other of the wrong prices (p_H or p_M). Below we provide a variation of Theorem 3.

Again we consider any sequence of discount factors $\{\delta_n\}_{n=1}^\infty$ converging to 1, and a corresponding sequence of Nash equilibria $\{\phi_n\}_{n=1}^\infty$ for which all relevant limits exist. Let $T_n \equiv \max\{T_n^{SL}, T_n^B\}$. We can show, as we did for the two-sided case, that $\lim_n \delta_n^{T_n} < 1$. The proof is similar, and much easier to carry out.

We can now ascertain whether or not information is fully revealed in state L . The volume of wrong-price trade in that state is

$$VW_n^L = (1 - x_B) \sum_{t \in \mathbb{N}} \phi_n^B(t) \phi_n^{SL}(t, \infty). \quad (32)$$

THEOREM 5. *Consider a sequence of games $\{G_n\}_{n=1}^\infty$ for which all parameters except δ are fixed (with $x_S = 1$), and where $\delta_n \rightarrow 1$, $\delta_n \in (0, 1)$ for all n ; and consider a corresponding sequence of Nash equilibria $\{\phi_n\}_{n=1}^\infty$ for which the relevant limits exist. Then $\lim_n VW_n^L > 0$.*

PROOF. Suppose $\lim_n VW_n^L = 0$. From (32), this requires, for any sequence of integers $\{\tau_n\}_{n=1}^\infty$,

$$\lim_n \phi_n^B(0, \tau_n) \phi_n^{SL}(\tau_n + 1, \infty) = 0. \quad (33)$$

This means there must be a sequence of integers $\{\lambda_n\}_{n=1}^\infty$ such that $\lim_n \phi_n^{SL}(0, \lambda_n) = \lim_n \phi_n^B(\lambda_n + 1, \infty) = 1$. Otherwise there would be, for infinitely many n , an integer τ_n such that $\phi_n^B(0, \tau_n) > \epsilon$ and $\phi_n^{SL}(\tau_n + 1, \infty) > \epsilon$ for some $\epsilon > 0$, in contradiction to (33). See Figure 3.




Figure 3. Densities of ϕ_n^B and ϕ_n^{SL} . The dark lines indicate the ranges where the mass is concentrated. Full information revelation ($\lim_n VW_n^L = 0$) requires a situation like that in panel (a) for all n greater than some m . If the situation is like that in panel (b) for infinitely many n , then $\lim_n VW_n^L > 0$.

Without loss of generality, $\phi_n^{SL}(\lambda_n) > 0$ for all n . [If an integer in the sequence has zero mass, it can be replaced with the next lowest integer that has positive mass without affecting the properties of the sequence.]¹⁴

Strategy λ_n must therefore be optimal for type-SL agents, for all n . In particular we must have $\pi_n^{SL}(\lambda_n) \geq \pi_n^{SL}(0)$ for all n . Strategy T_n^B must by definition be optimal for type-B agents, and in particular we must have $\pi_n^B(T_n^B) \geq \pi_n^B(\lambda_n + 1)$ for all n . We calculate

$$\lim_n [\pi_n^{SL}(\lambda_n) - \pi_n^{SL}(0)] = (p_L - c_L) \left(\lim_n \delta_n^{\lambda_n} - 1 \right); \quad (34)$$

$$\lim_n [\pi_n^B(T_n^B) - \pi_n^B(\lambda_n + 1)] = \alpha(u_H - p_H) \left(\lim_n \delta_n^{T_n} - \lim_n \delta_n^{\lambda_n + 1} \right). \quad (35)$$

From (34), optimality of λ_n for type-SL agents for all n requires $\lim_n \delta_n^{\lambda_n} = 1$. From (35), optimality of T_n^B for type-B agents for all n requires $\lim_n \delta_n^{\lambda_n} = \lim_n \delta_n^{T_n}$. But we know that $\lim_n \delta_n^{T_n} < 1$. This is a contradiction. \square

In words, full information requires the masses of ϕ_n^{SL} and ϕ_n^B to separate as $n \rightarrow \infty$, with the ϕ_n^B mass accumulating to the right of the ϕ_n^{SL} mass. The optimality

¹⁴It can be shown that $\lambda_n \equiv \min\{\arg \max_\tau \phi_n^{SL}(0, \tau) \phi_n^B(\tau + 1, \infty)\}$ forms an appropriate sequence.

which is supposed to be reflected by these distributions comes into question. The situation is optimal for type-SL agents only if waiting from 0 to λ_n becomes costless as $n \rightarrow \infty$. It is optimal for type-B agents only if waiting from $\lambda_n + 1$ to T_n^B becomes costless. Hence waiting from 0 to T_n^B must become costless. But as we saw in Claim 8, this makes for an infeasible outcome.

6.3 Efficiency

The efficiency result is the same as in the two-sided case: the economy is asymptotically interim incentive inefficient. The proof is very similar.

7 Concluding Remarks

Negative results emerge in both the two-sided and one-sided information problems: failure of learning and social inefficiency. Different forces are at work in the two models, however, which is why Theorems 3 and 5 were proved differently. In the two-sided model (assuming all agents trade), force CL competes with forces MI and N, and cannot overcome them. The proof of Theorem 3 does not involve type-SL or type-BH agents (those with an incentive to lie about the state of the world): force N by itself guarantees the result. In contrast, force N is absent from the one-sided model, since there are no meetings between uninformed agents: there force CL loses to force MI. In the two-sided version, force F may prevent some agents from ever trading. In the one-sided version, force F does not appear, and all agents trade. All equilibrium outcomes are asymptotically interim incentive inefficient, due to excessive delay.

Ultimately we would like to isolate the properties which a model of decentralized trade should have in order to yield a positive result. To this end, we discuss briefly the role of three assumptions made in our model.

Anonymity. We know that a key assumption for our results is the strict anonymity of the procedure: an agent not only cannot tell whether his opponent is informed or not, he also cannot recognize an opponent he may have met before. The following example shows why this is important.

EXAMPLE 2. Consider the one-sided information model of section 6. Continue to assume that agents are unable to distinguish between informed and uninformed opponents. They can, however, recognize individuals they have met before. The following is now an equilibrium:

- (i) *each type-SH and type-BL agent plays ∞ ;*
- (ii) *each type-SL and type-BH agent plays 0;*

(iii) *each type-B agent plays tough in period 0 and whenever he meets an opponent he has met before; otherwise, he plays soft.*

If all agents follow their equilibrium strategies in this example, the game ends in period 0 (state L) or 1 (state H). All trade takes place at EPIR prices, and no inefficiency persists in the limit. What makes the equilibrium work is that if a type-SL agent deviates and plays tough in period 0 (and no one else deviates), then he and his opponent will be the only ones left in the game, and they will meet again in period 1 (and 2, and 3, ...).

Equal population size. We have assumed that the measure of buyers and the measure of sellers are equal. Suppose that this is not the case, but that (as before) each agent can only meet one opponent at a time: then in each period there are agents who are unmatched and hence cannot bargain. The two-sided model should not be affected qualitatively by this, since force N should continue to command the outcome. The one-sided model *is* affected, at least when there are more sellers than buyers. The following is an equilibrium, provided that all buyers are matched in period 0: type-SH and type-BL agents play ∞ ; type-SL and type-BH agents play 0; type-B agents play 1. Sellers in state L are forced to play soft: by playing tough, a seller would not trade in period 0, and his probability of being matched to the only remaining buyer in the next period is 0. This is a significant difference with the model of this paper, and it persists even when the measure of sellers is arbitrarily close to that of buyers.

We like to think of our model as a limit of others where the measures of buyers and sellers are different. Consider a situation where the measure of buyers is μ_B , while that of sellers is μ_S , and suppose $\mu_B < \mu_S$. Suppose that there is a friction in the matching process so that the probability for a buyer to meet a seller is $f(\mu_B/\mu_S) < 1$, and so the probability for a seller to meet a buyer is $g(\mu_B/\mu_S) = (\mu_B/\mu_S)f(\mu_B/\mu_S)$. The functions f and g are continuous and such that $f(1) = g(1) = 1$. The sequence of these models as μ_S tends to μ_B would approximate the one in this paper. To the extent that equilibrium payoff correspondences are upper-hemicontinuous, our results suggest that these models would also yield negative results. Note in particular that the equilibrium identified in the previous paragraph ceases to be an equilibrium now. It is not clear, however, what these models would yield, due to the presence of a double limit: $\lim_{\delta \rightarrow 1}$ and $\lim_{\mu_S \rightarrow \mu_B}$. It is possible that the relative speed at which these two frictions are removed may have an impact in the results [see Muthoo (1993)]. Suppose, for example, that all buyers are informed and that $\mu_B < \mu_S$. On the one hand, there is excess supply for the good, so sellers should be playing softer than in our model (this would work against learning). On the other hand, they could learn now from the different frequency of their meetings in the two states (which would work in favor of learning). Clearly, further research

is needed to sort out these complicated effects.

The trading procedure. The two-action simultaneous procedure is a strong assumption. A sequential procedure might be more conducive to information revelation, as information may be revealed in the proposal made by the first mover. However, the analysis of this model appears quite challenging: it would combine elements of signalling and screening in each meeting.

References

- Bergemann, D. and Välimäki, J. (1996). "Learning and strategic pricing." *Econometrica* **64**, 1125-1149.
- Fan, K. (1952). "Fixed-point and minimax theorems in locally convex topological linear spaces." *Proceedings of the National Academy of Sciences, U.S.A.* **38**, 121-126.
- Gale, D. (1986a). "Bargaining and competition part I: characterization." *Econometrica* **54**, 785-806.
- Gale, D. (1986b). "Bargaining and competition part II: existence." *Econometrica* **54**, 807-818.
- Gale, D. (1987). "Limit theorems for markets with sequential bargaining." *Journal of Economic Theory* **43**, 20-54.
- Glicksberg, I.L. (1952). "A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points." *Proceedings of the American Mathematical Society* **3**, 170-174.
- Hildenbrand, W. (1974). *Core and Equilibria of a Large Economy*. Princeton: Princeton University Press.
- Holmström, B. and Myerson, R.B. (1983). "Efficient and durable decision rules with incomplete information." *Econometrica* **51**, 1799-1819.
- Khan, M.A. (1989). "On Cournot-Nash equilibrium distributions for games with a nonmetrizable action space and upper semicontinuous payoffs." *Transactions of the American Mathematical Society* **315**, 127-146.
- Krishna, V. and Morgan, J. (1997). "An analysis of the war of attrition and the all-pay auction." *Journal of Economic Theory* **72**, 343-362.
- Laffont, J.J. (1985). "On the Welfare Analysis of Rational Expectations Equilibria with Asymmetric Information." *Econometrica* **53**, 1-29.
- Mas-Colell, A. (1984). "On a theorem of Schmeidler." *Journal of Mathematical Economics* **13**, 201-206.

- Milgrom, P.R. (1979). "A convergence theorem for competitive bidding with differential information." *Econometrica* **47**, 679-688.
- Milgrom, P.R. (1981). "Rational expectations, information acquisition, and competitive bidding." *Econometrica* **49**, 921-943.
- Muthoo, A. (1993). "Sequential bargaining and competition." *Economic Theory* **3**, 353-363.
- Osborne, M. J. and Rubinstein, A. (1990). *Bargaining and Markets*. San Diego: Academic Press.
- Pesendorfer, W. and Swinkels, J.M. (1997). "The loser's curse and information aggregation in common values auctions." *Econometrica* **65**, 1247-1281.
- Rothschild, M. (1974). "A two-armed bandit theory of market pricing." *Journal of Economic Theory* **9**, 185-202.
- Rubinstein, A. and Wolinsky, A. (1985). "Equilibrium in a market with sequential bargaining." *Econometrica* **53**, 1133-1150.
- Schmeidler, D. (1973). "Equilibrium points of nonatomic games." *Journal of Statistical Physics* **4**, 295-300.
- Serrano, R. and Yosha, O. (1993). "Information revelation in a market with pairwise meetings: the one sided information case." *Economic Theory* **3**, 481-499.
- Serrano, R. and Yosha, O. (1996). "Welfare analysis of a market with pairwise meetings and asymmetric information." *Economic Theory* **8**, 167-175.
- Wilson, R. (1977). "A bidding model of perfect competition." *Review of Economic Studies* **44**, 511-518.
- Wolinsky, A. (1990). "Information revelation in a market with pairwise meetings." *Econometrica* **58**, 1-23.