

Federal Reserve Bank of Minneapolis  
Research Department Staff Report 144

August 1991

## ELICITING TRADERS' KNOWLEDGE IN "FRICTIONLESS" ASSET MARKET

Edward J. Green\*

University of Minnesota and  
Federal Reserve Bank of Minneapolis

### ABSTRACT

---

Intuitively, a patient trader should be able to make his trading partners compete to reveal whatever information is relevant to their transactions with him. This possibility is examined in the context of a model resembling that of Gale (1986). The main result is that, under assumptions having to do with asset structure and spanning, incentive-compatible elicitation of trading partners' knowledge is feasible.

---

\*This paper will appear in the proceedings of the International Conference on Game Theory that was held at the Indian Statistical Institute, Delhi Centre, in December 1990. My interest in this subject was originally stimulated by discussions with Charles Plott and Shyam Sunder regarding subjects' behavior in experimental asset markets. More recently Eric Maskin, Steven Matthews, Nancy Stokey, and Asher Wolinsky have all offered suggestions that are reflected in the present formulation and interpretation of these results.

The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

## 1. Introduction

Consider an uninformed trader who is very patient, and who has opportunities to bargain with many informed trading partners. Intuitively this trader should be able to make these partners compete with one another in their dealings with him, and one aspect of such competition would be that they should reveal whatever information is relevant to their transactions with him. Asher Wolinsky (1990) has shown that this intuition is not necessarily correct, though. Wolinsky considers noncooperative equilibrium in a partial-equilibrium model of trade in an indivisible good of random quality that is unobserved by some of the traders. He shows that a fraction of these traders make uninformed transactions, and that this fraction has a positive lower bound even as the parameters of the market approach a “frictionless-bargaining” limit.

Here I will consider this issue further. Specifically, this paper concerns how a patient trader might elicit the knowledge of trading partners in an incentive-compatible way. This possibility will be examined in the context of a model resembling that of Douglas Gale (1986). This model is chosen for two reasons. First, in the complete-information case, Gale has related the noncooperative equilibrium of this model very transparently to Walras equilibrium. Since an important aspect of Wolinsky’s work is its bearing on an intuition about perfect competition, this transparency is a virtue. Second, additional structure suggested by the theory of financial markets can be imposed on the model. This additional structure has immediate relevance to the question at hand.

The main result is that, under these assumptions having to do with asset structure and spanning, incentive-compatible elicitation of trading partners’ knowledge is feasible. Moreover there is a sense in which a trader need not be very sophisticated in order to interpret trading partners’ responses to his elicitation efforts. Specifically, although learning from experience in the market generally requires the application of Bayes’ Theorem, this form of elicitation works in a more statistically straightforward way.

It seems plausible that strategies incorporating such elicitation could implement full-information Walras equilibrium if small deviations from optimal behavior are tolerated.

## 2. The economy

Let there be  $\ell$  divisible goods and a nonatomic measure space  $(I, \mathcal{A}, \nu)$  of traders, and let each trader have a closed convex cone  $\mathbf{X} \in \mathbb{R}^\ell$  having nonempty interior as his consumption set. Assume that  $\nu$  is normalized to be a probability measure.

There is a measurable space  $(\Omega, \mathcal{B})$  of states of the world.

There are  $T$  types of trader, with a positive measure of individual traders belonging to each type. For each type  $t$ , let  $A^t \in \mathcal{A}$  be the set of traders of type  $t$ . Assume that  $\nu(A^t)$  is a rational number for each  $t$ . A trader's type determines his endowment, utility function, prior beliefs about the state of nature, and information partition on  $\Omega$ . All traders of type  $t$  receive endowment  $e^t \in \text{int}(\mathbf{X})$ ; they have the strictly concave, continuously differentiable utility function  $U^t : \mathbf{X} \times \Omega \rightarrow [0, 1]$ ; they share prior beliefs about the state of nature represented by a probability measure  $\pi_t : \mathcal{B} \rightarrow [0, 1]$ ; and the information that they are able to observe directly about the state of nature is determined by a finite algebra  $\mathcal{F}_t \subseteq \mathcal{B}$ . Each trader of type  $t$  initially maximizes the expectation of  $U^t$  with respect to  $\pi_t$  conditioned on  $\mathcal{F}_t$ . Let  $\mathcal{F}$  denote the smallest algebra that contains  $\mathcal{F}_t$  for every  $t$ . That is,  $\mathcal{F}$  is the algebra of events on which traders will condition their expectations if they fully share their information with each other.  $\mathcal{F}(\omega)$  will denote the atom of  $\mathcal{F}$  that contains state of nature  $\omega$ .

There is a finite algebra  $\mathcal{P} \subseteq \mathcal{B}$  of payoff-relevant events. That is, each utility function  $U^t$  is measurable with respect to the  $\sigma$ -algebra  $\mathcal{E} \times \mathcal{P}$ , where  $\mathcal{E}$  is the Borel  $\sigma$ -algebra on  $\mathbf{X}$ . For convenience, I will assume that this statement remains true for an extension of  $U^t$  to the Cartesian product with  $\Omega$  of an open set  $Y$  containing  $\mathbf{X}$ , that this extension of  $U^t$  is continuously differentiable in consumption for each  $\omega$ , and that  $U^t(x, \omega) = 0$  if and only if  $x$  is a boundary point of  $\mathbf{X}$ . The gradient of  $U^t(x, \omega)$  as a function of  $x$  (with  $\omega$  treated as a parameter) will be denoted by  $\nabla U^t(x, \omega)$ .

Transactions take place in random pairwise meetings that occur in discrete time, beginning at date 1. This random-meetings technology is specified by two functions:  $M : I \times \mathbb{N}_+ \times \Omega \rightarrow I$  and  $P : I \times \mathbb{N}_+ \times \Omega \rightarrow \{0, 1\}$ . Each of these functions represents an *i.i.d.*

family of random variables indexed by  $I \times \mathbb{N}_+$ . Each function is measurable separately in  $I$  and in  $\Omega$  for every combination of values of the other arguments. (Green (1988) shows that the functions cannot be jointly measurable in  $I \times \Omega$ , given the *i.i.d.* assumption.)

These functions have the following intuitive interpretation. At each date  $n$ , trader  $i$  is randomly matched with another trader  $j = M(i, n, \omega)$ . It is assumed that  $i$  and  $j$  have never met before, and in fact that they are not even indirectly acquainted with one another. (*e.g.* there is no trader  $k$  who has previously traded first with  $i$  and subsequently with  $j$ , and whose behavior towards  $j$  might therefore affect how  $j$  will currently be disposed to deal with  $i$ .) The matched traders do not observe anything about one another's past history, either regarding type or beliefs. One of these two traders is required to propose a net trade (which may be the zero net trade) to the other, and this trade must be feasible for the proposer to execute. For every trader  $k$ ,  $P(k, n, \omega) = 1$  if and only if  $k$  is the proposer at date  $n$  in state  $\omega$ . A transaction takes place if the trading partner accepts this proposal. In any event, the two traders are separated after the partner replies to the proposal.

In state  $\omega$ , each trader  $i$  successively holds an infinite sequence of commodity bundles  $x_0, x_1, x_2, \dots$ . The bundle  $x_0$  is  $e^t$  if  $i$  is of type  $t$ , and the successive bundles are determined by the successive net trades that  $i$  transacts. I will assume that the utility enjoyed by  $i$  in state  $\omega$  is  $\liminf_{n \rightarrow \infty} U^t(x_n)$ . That is, commodity bundles may be interpreted as yielding a flow of services, and traders may be viewed as maximizing the asymptotic value of this flow. This specification differs from the one used in most of the literature on "frictionless markets" where it is required that each trader must eventually consume the commodity bundle that he holds, and must leave the market at that point. The present specification is adopted because it is technically simpler than the usual one—for example, it is not needed to assume that there is a constant flow of new traders into the market. No essential feature of the results of this paper depends on this flow-of-services feature of the model. However, the no-discounting specification is important. I will discuss its importance after I have presented the analysis of the model.

The foregoing discussion will probably be sufficiently clear for most readers. For the

record, though, here is a more formal description of the model. Relative to each  $\pi_t$ , the following assumptions are satisfied by the random variables determined by  $M$  and  $P$  for all values of  $i$  and  $n$ . These random variables are independent of one another and of all of the algebras  $\mathcal{F}_t$  and  $\mathcal{P}$ . For all  $i$  and  $n$ ,  $\pi_t(\{\omega | P(i, n, \omega) = 1\}) = 1/2$ . For every  $i, n, t, u$ , and  $\omega$  (almost surely with respect to every  $\pi_t$ ),  $i \neq M(i, n, \omega)$  and  $i = M(M(i, n, \omega), n, \omega)$  (*i.e.* the matching relation is irreflexive and symmetric);  $P(i, n, \omega) = 1 - P(M(i, n, \omega), n, \omega)$  (*i.e.* exactly one member of each matched pair is the proposer);  $\pi_t(\{\omega | M(i, n, \omega) \in A_u\}) = \nu(A_u)$  (*i.e.* various types of trader are matched *ex ante*);  $\nu(\{i | i \in A_u \text{ and } M(i, n, \omega) \in A_w\}) = \nu(A_u)\nu(A_w)$  (*i.e.* various types of trader are randomly matched *ex post*); and, almost surely, traders who have previously met or have been related indirectly in the sense explained above are not matched with one another again.

### 3. Strategies, histories and the evolution of traders' beliefs

A strategy is a rule that prescribes, at each date and in each state of the world, how a trader should behave as a function of what he has observed prior to choosing his current action. If the trader must make a proposal, then what he has observed is simply the history of his own prior trading. If the trader is responding to a proposal, then the current proposal is also part of his observational basis for choice. Thus at each date, two actions—the proposal and the response—occur in sequence. The definition of noncooperative equilibrium will reflect this sequential structure.

At date  $n$ , a trader's history of observations at prior dates can be described as a sequence  $h^n = (h_0, \dots, h_{n-1})$ . If trader  $i$  is of type  $t$  and if  $h$  is  $i$ 's history prior to  $n$  in state  $\omega$ , then  $h_0$  will be the ordered pair  $(t, F)$ , where  $F$  is the atom of  $\mathcal{F}_t$  that contains  $\omega$ . For  $m \in \{1, \dots, n-1\}$ ,  $h_m$  will be an ordered pair specifying the actions of  $i$  and his trading partner at  $m$ . (That is, the first element of the ordered pair will be  $i$ 's action, and the second will be his partner's action. One of these will always be a proposed trade, and the other an acceptance or rejection.)

For each  $n \in \mathbb{N}_+$ , let  $H_n^t$  be the set of such histories prior to date  $n$  for a trader of type  $t$ .

At date  $n$ , each trader observes either that he is a proposer or else that a net trade  $z \in \mathbb{R}^{\ell}$  has been proposed for his response. If the trader is of type  $t$ , then the information on which his decision at date  $n$  must be based is an element of the set  $\mathcal{H}_n^t = H_n^t \times [\{\text{propose}\} \cup \mathbb{R}^{\ell}]$ . (Elements of  $\mathcal{H}_n^t$  will be denoted by  $h = (h_0, \dots, h_{n-1}, z)$ .) Let  $\mathcal{H}^t = \bigcup_{n \in \mathbb{N}_+} \mathcal{H}_n^t$ . Then a strategy for a trader of type  $t$  is a function  $\sigma^t : \mathcal{H}^t \rightarrow \mathbb{R}^{\ell} \cup 2^{\mathbb{R}^{\ell}}$  that satisfies  $\sigma^t(h, z) \in \mathbb{R}^{\ell} \iff z = \text{propose}$  for all  $(h, z) \in \mathcal{H}^t$ . That is,  $\sigma^t(h)$  specifies the offer that will be made if the trader is required to propose, and it specifies the set of offers that the trader would accept if his partner is the proposer.

I will be concerned only with equilibrium strategy profiles that are symmetric in each type, so that a strategy profile can be represented as a vector  $\sigma = (\sigma^1, \dots, \sigma^T)$  of strategies for all types of trader. Equivalently, if  $H^* = \bigcup_{t \leq T} \bigcup_{n \in \mathbb{N}} H_n^t$ , then a symmetric strategy profile can be represented by a function  $\sigma : H^* \rightarrow \mathbb{R}^{\ell} \times 2^{\mathbb{R}^{\ell}}$ . Given such a symmetric strategy profile, it is easy to see that only finitely many histories in each  $H_n^t$  can be generated by play according to the profile. Using this fact, it can routinely be shown that no measurability difficulties arise in defining the stochastic process of transactions or in making expected-utility calculations based on this stochastic process.

**Lemma 1.** *Let  $H_n(\sigma)$  and  $\mathcal{H}_n(\sigma)$  be the set of the histories in  $\bigcup_{t \leq T} H_n^t$  and  $\bigcup_{t \leq T} \mathcal{H}_n^t$  respectively that can occur with positive probability when  $\sigma$  is played. (Note that  $\mathcal{H}_n(\sigma)$  is not defined for  $n = 0$ .) Then  $H_n(\sigma)$  and  $\mathcal{H}_n(\sigma)$  are finite for every strategy profile  $\sigma$  and for every date  $n$ .*

**Proof:** To begin,  $H_0(\sigma) = \{(t, F) | t \leq T \text{ and } F \in \mathcal{F}_t\}$ . Recursively,  $H_{n+1}(\sigma) = \{(h^n, (z, \text{accept})) | z = \sigma_1(h^n) \text{ and } \exists h' \in H_n(\sigma) [z \in \sigma_2(h')]\} \cup \{(h^n, (z, \text{reject})) | z = \sigma_1(h^n) \text{ and } \exists h' \in H_n(\sigma) [z \notin \sigma_2(h')]\} \cup \{(h^n, (\text{accept}, z)) | z \in \sigma_2(h^n) \text{ and } \exists h' \in H_n(\sigma) [z = \sigma_1(h')]\} \cup \{(h^n, (\text{reject}, z)) | z \notin \sigma_2(h^n) \text{ and } \exists h' \in H_n(\sigma) [z = \sigma_1(h')]\}$ . By induction, each of these sets  $H_n(\sigma)$  is finite. The finiteness of  $\mathcal{H}_n(\sigma)$  is proved by observing that  $\mathcal{H}_{n+1}(\sigma) = H_n(\sigma) \times [\{\text{propose}\} \cup \{\sigma_1(h) | h \in H_n(\sigma)\}]$ . *Q.E.D.*

How a trader should update his beliefs about the state of nature, based on his history, is now considered. The ultimate objective of the trader is to infer from the history to a

posterior probability measure over  $\Omega$ . To understand how this inference should correctly be made, consider first the case of a trader  $i$  who has never been the proposer prior to date  $n$ . Let us make a “full-revelation” assumption that, for each  $m \in \mathbb{N}$ , the history-contingent trade proposals recommended by the strategies of  $\sigma$  map distinct elements of  $H_m(\sigma)$  to distinct elements of  $\mathbb{R}^\ell$ . Theorem 1 will establish that in this case (with analogous assumptions that  $i$ ’s trading partners, their trading partners, and so forth have never been proposers before the matchings that connect them with  $i$ ), if trader  $i$ ’s history is consistent with  $\sigma$ , then  $i$  can determine the history in  $H^*$  of his partner at each prior date, of each of his past partners’ partners up to the date when they were matched, and so forth. Denote this set of histories by  $\theta(h^n, \sigma)$ , if  $h^n$  is  $i$ ’s history in  $H_n(\sigma)$ .

Formally, define  $\Theta(h^n, \sigma)$  to be the smallest set among whose elements are all subsets  $\theta \subseteq H^*$  that satisfy two conditions. First,  $h^n \in \theta$ . Second, if  $h^{m+1} = (h_0, \dots, h_m) \in \theta$  and  $h^m = (h_0, \dots, h_{m-1})$ , then there is a history  $h' = (h'_0, \dots, h'_{m-1}) \in \theta$  such that  $h_m$  is the outcome of the proposal and response  $\sigma(h^m)$  and  $\sigma(h')$ . [Note that the ordering of proposal and response in  $h_m$  determines which of the previous histories belongs to the proposer at date  $m$ , and which to the responder.] Every element  $\theta \in \Theta(h^n, \sigma)$  can be regarded as a theoretical explanation of how  $h^n$  would be generated in the play of  $\sigma$ . That is,  $\theta(h^n, \sigma)$  is a set of histories that provides a complete explanation of how  $h^n$  came to be observed in the context of the strategy profile  $\sigma$ , and no proper subset of  $\theta(h^n, \sigma)$  would provide a complete explanation.

**Theorem 1.** *If history  $h^n \in H_n(\sigma)$ , if for each  $m < n$  the restriction of  $\sigma_1$  to  $H_m(\sigma)$  is 1-1 into  $\mathbb{R}^\ell$ , if  $h^n$  specifies that  $i$  has been the responder at every date prior to  $n$ , and if analogously  $i$ ’s trading partners, their trading partners, and so forth have never been proposers before the matchings that connect them with  $i$ , then  $\Theta(h^n, \sigma)$  consists of a single element (to be denoted by  $\theta(h^n, \sigma)$ ).*

**Proof:** The proof is by induction. At date 1,  $i$  knows his own his type and the element of his observation algebra that he has observed (*i.e.* his history prior to date 1), and he has had no previous trading partner. This knowledge is all that is asserted for  $n = 1$ . Now

consider the induction hypothesis that each trader can infer from his history  $h^n$  everything about his trading partners' histories, their partners' histories prior to the dates when  $i$  traded with them, *et cetera*—*i.e.* that  $\theta(h^n, \sigma)$  is uniquely defined. Then  $h^{n+1}$  also reveals this information regarding partners before date  $n + 1$ , since  $h^n$  is the initial segment of  $h^{n+1}$ . By the full-revelation assumption,  $h^{n+1}$  also determines uniquely the history  $h'$  of  $i$ 's trading partner at date  $n$ . Applying the induction hypothesis to this trading partner, then, we see that  $i$  can infer the histories of his date- $n$  trading partner's previous trading partners, their histories prior to the dates when  $i$ 's date- $n$  trading partner traded with them, *et cetera*—*i.e.* that  $\theta(h', \sigma)$  is uniquely defined. Then  $\theta(h^{n+1}, \sigma) = \theta(h^n, \sigma) \cup \theta(h', \sigma) \cup \{h_{n+1}\}$ . This establishes the induction step, so the assertion holds for all dates. *Q.E.D.*

Each history in  $\theta(h^n, \sigma)$  belongs to  $i$  or to some other trader with whom  $i$  has been directly or indirectly involved in the history of transactions. Now it is straightforward to describe what  $i$  should do. He knows that each history in  $\theta(h^n, \sigma)$  begins with a direct observation of an atom of one of the information algebras  $\mathcal{F}_t$ , and that these observations constitute the whole of the direct evidence about the state of nature possessed by anyone with whom he has ever been directly or indirectly associated. Therefore  $i$  (as a good Bayesian) should condition his prior probability distribution on the intersection of these atoms. Formally, let  $\Delta : H^* \rightarrow 2^\Omega$  specify the direct observation that begins each history. (That is if  $h_0 = (t, F)$ , then  $\Delta(h) = F$ .) If trader  $i$  is of type  $t$ , then his posterior beliefs at the beginning of date  $n$  should be represented by  $\pi_t$  conditioned on the event  $\bigcap \{\Delta(h) | h \in \theta(h^n, \sigma)\}$ . This beginning-of-date posterior probability measure represents the beliefs of a trader at the time when he must make a proposal. However, if the trader is a responder, then he conditions his beliefs both on his own history  $h^n \in H_n(\sigma)$  and on the observation of the net trade  $z$  that is offered to him at date  $n$ . That is, in order to cover the case of both a proposer and a responder, the trader's posterior probability must be defined in terms of histories in  $\mathcal{H}_n(\sigma)$  rather than of histories in  $H_n(\sigma)$ . If  $\sigma_1$  is 1-1 on  $H_n(\sigma)$ , then the responder will infer that the proposer's history is  $\sigma_1^{-1}(z)$ . Therefore the beliefs of a responder of type  $t$  with history  $(h^n, z)$  should be represented by  $\pi_t$  conditioned on the event  $\bigcap \{\Delta(h) | h \in \theta(h^n, \sigma) \cup \theta(\sigma_1^{-1}(z), \sigma)\}$ .



In general, an element of  $\Theta(h^n, \sigma)$  will be called a theory of  $h^n$ . Without the “full-revelation” assumption (*e.g.* if  $i$  had been a proposer at some previous date, and if there are several theories that would account (consistently with the strategy profile  $\sigma$ ) for the partner’s response to the proposal), the probability-updating problem is more complicated than when the assumption holds. In this general case, there may be several alternative theories  $\theta \in \Theta(h^n, \sigma)$ . Each of these theories would be consistent with the history that trader  $i$  has observed. Based on his prior probability beliefs about the random-matchings process,  $i$  can assign relative likelihoods to these alternative theories. Then for each theory he should compute a posterior probability measure by conditionalization as above, and he should weight these conditional probabilities by the likelihoods of the respective theories. (That is, he should apply Bayes’ Theorem.) The formal details of this likelihood-weighted conditionalization will not be needed in this paper, even though a noncooperative Bayesian equilibrium concept will be studied. The avoidance of appeal to Bayes’ Theorem may be interpreted as showing that traders are able to act optimally according to a Bayesian criterion—and to recognize that they are acting optimally—without having to solve the fairly complex computational problem that Bayesian inference usually entails. This interpretation is one of the most important aspects of this paper, linking it with the study of “bounded rationality” in economic equilibrium.

The other complication that might arise would be that no theory with positive prior probability would be consistent with equilibrium behavior that would generate a trader’s history. In that case Bayesian revision of the prior probability measure is undefined, and posterior probability must be defined in some alternative way. In the theory of finite-player, extensive-form games, specifying a reasonable assignment of posterior probabilities is an important, difficult, and somewhat controversial problem. What makes the problem so deep is the potential for one player strategically to manipulate another player’s beliefs, in order to profit from that manipulation later in the game. Since the transactions technology of the present model provides no opportunity for such strategic manipulation, this issue of conditioning on observations of aberrant behavior will not arise here. To complete the specification of probability updating, assume that a trader who observes aberrant behavior on the part of a trading partner will simply ignore it and retain his prior probability

assessment. However, a trader who deviates from his own strategy in the profile  $\sigma$  will revise his probability assessment on the basis of his trading partner's action just as he would ordinarily do.

In this section, the evolution of traders' beliefs about the state of nature has been specified. Formally, let  $\mathcal{M}$  denote the set of probability measures defined on the measurable space  $(\Omega, \mathcal{B})$  and also define  $\mathcal{H}^* = \bigcup_{t \leq T} \mathcal{H}^t$ . It has been shown that there is a function  $\Pi : \mathcal{H}^* \times \sigma \rightarrow \mathcal{M}$  such that  $\Pi(h, \sigma)$  is the posterior probability measure that represents the beliefs of a trader whose history is  $h$ , if the trader is certain that all trading partners are behaving as specified by  $\sigma$ .

#### 4. Trading, perfect Bayesian equilibrium, and full-information Walras equilibrium

Recall that, in state  $\omega$ , each trader  $i$  successively holds an infinite sequence of commodity bundles  $x_0, x_1, x_2, \dots$ . The bundle  $x_0$  is  $e^t$  if  $i$  is of type  $t$ , and the successive bundles are determined by the successive net trades that  $i$  transacts. Clearly a trader's successive commodity holdings can be described in terms of his history, that is, by a function  $X^* : H^* \rightarrow \mathbf{X}$ . If  $h_0 = (t, F)$ , then  $X^*(h^1) = e^t$ . Let  $z \in \mathbb{R}^l$ . If  $h_n = (z, \text{accept})$ , then  $X^*(h^{n+1}) = X^*(h^n) - z$ . If  $h_n = (\text{accept}, z)$ , then  $X^*(h^{n+1}) = X^*(h^n) + z$ . If  $h_n = (z, \text{reject})$  or  $h_n = (\text{reject}, z)$ , then  $X^*(h^{n+1}) = X^*(h^n)$ .

Note that a trader's history in any state of nature and at any date is determined by his identity, the date, the state of nature, and the strategy profile that is being played. That is, letting  $\Sigma$  denote the set of strategy profiles, there is a function  $\eta : I \times \mathbb{N}_+ \times \Omega \times \Sigma \rightarrow H^*$  such that, if  $h^n$  is  $i$ 's history at date  $n$  in state of nature  $\omega$  when  $\sigma$  is played, then  $h^n = \eta(i, n, \omega, \sigma)$ . Thus the function  $X : I \times \mathbb{N}_+ \times \Omega \times \Sigma \rightarrow \mathbf{X}$  defined by  $X = X^* \circ \eta$  characterizes the random evolution of all traders' commodity holdings as a function of the strategy profile.

It is clear that each trader in this environment can be represented as solving a stochastic dynamic programming problem, when the strategies of other traders are given. Fol-

lowing Tirole (1988), I will refer to a strategy vector for which each trader's strategy is an optimal solution of that trader's dynamic programming program as a perfect Bayesian equilibrium. To state this equilibrium concept explicitly in the context of the present model, let  $\tau(h)$  denote the unique type (encoded in  $h_0$ ) of trader for whom  $h \in \mathcal{H}^t$ . Define strategy profile  $\sigma'$  to be a deviation after history  $h^n$  from strategy profile  $\sigma$  if (a)  $h^n \in H^{\tau(h)}$ , and (b)  $\sigma'(h^m) = \sigma(h^m)$  for all  $m < n$ , and (c)  $\sigma'(h) = \sigma(h)$  for all  $h \notin H^{\tau(h)}$ . The strategy profile  $\sigma$  is a perfect Bayesian equilibrium if, whenever  $h \in H^t$  and  $i \in A^t$  and  $\sigma'$  is a deviation after  $h$  from  $\sigma$ , the expectation according to  $\Pi(h)$  of  $\liminf_{n \rightarrow \infty} U^t(X(i, n, \omega, \sigma'))$  does not exceed the corresponding expectation of  $\liminf_{n \rightarrow \infty} U^t(X(i, n, \omega, \sigma))$ .

Note that, by what has been assumed about the stochastic process of random meetings, the value of the expectation of  $\liminf_{n \rightarrow \infty} U^t(X(i, n, \omega, \sigma))$  according to  $\Pi(h)$  does not depend on which representative  $i$  of type  $\tau(h)$  is taken (except possibly for a set of  $i$  having  $\nu$ -measure zero). Therefore this conditional expected utility can be denoted by  $\bar{U}(h, \sigma)$ . That is,  $\bar{U}(h, \sigma) = \int_{\Omega} \liminf_{n \rightarrow \infty} U^{\tau(h)}(X(i, n, \omega, \sigma)) d\Pi(h)$ . Now the definition of equilibrium can be restated:  $\sigma$  is a perfect Bayesian equilibrium if, whenever  $h \in H^*$  and  $\sigma'$  is a deviation after  $h$  from  $\sigma$ ,  $\bar{U}(h, \sigma') \leq \bar{U}(h, \sigma)$ .

Define  $w : I \times \Omega \rightarrow \text{int}(\mathbf{X})$  to be an interior full-information Walras allocation if two conditions hold. First, there must be an  $\mathcal{F}$ -measurable price function  $p : \Omega \rightarrow \mathbb{R}^t \setminus \{0\}$  such that, for a set of traders  $i$  having  $\nu$ -measure 1, almost surely (with respect to each of the measures  $\pi_t$ ), the commodity bundle  $w(i, \omega)$  maximizes on the budget set  $\{x | x \in \mathbf{X} \text{ and } p(\omega) \cdot x \leq p(\omega) \cdot e^t\}$  the expectation of  $U^t$  with respect to  $\pi_t$  conditioned on the event  $\mathcal{F}(\omega)$ . (Here it is understood that  $i \in A^t$ .) Second, the materials-balance condition  $\int_I w(i, \omega) d\nu = \sum_{t \leq T} e^t \nu(A_t)$  must hold almost surely.

Denote traders' asymptotic commodity holdings by  $\chi(i, \omega, \sigma) = \lim_{n \rightarrow \infty} X(i, n, \omega, \sigma)$  if this limit exists, and let  $\chi(i, \omega, \sigma) = 0$  otherwise. If  $w$  is an interior full-information Walras allocation and  $\sigma$  is a perfect Bayesian equilibrium, then define  $\sigma$  to implement  $w$  if  $\forall i \forall \omega w(i, \omega) = \chi(i, \omega, \sigma)$ .

In general, optimality of a trader's strategy will require the trader both to trade optimally and to learn optimally about the true state of nature. However, if all traders have the

same information algebra  $\mathcal{F}$ , then there is nothing to be learned beyond what each trader observes directly. In this degenerate case, the environment defined here is essentially identical to that studied by Douglas Gale (1986) and by McLennan and Sonnenschein (1989). That is, in each state of nature  $\omega$ , the preferences of each trader of type  $t$  are specified by the conditional expectation of  $U^t$  with respect to  $\pi_t$  conditioned on  $\mathcal{F}$ . The traders' endowments and these preferences in  $\omega$  determine a nonempty set of Walras allocations in  $\omega$ . Gale (1986) has proved a result that is tantamount to the following theorem.

**Theorem 2.** *If all traders have the same information algebra (i.e.  $\forall t \mathcal{F}_t = \mathcal{F}$ ), then every interior Walras allocation is implemented by a perfect Bayesian equilibrium. There is some  $\epsilon > 0$  such that, with probability 1 on the equilibrium path, no trader's commodity holding  $X(i, n, \omega, \sigma)$  is ever closer than  $\epsilon$  to the boundary of the consumption set.*

**Proof:** Gale's proof concerns a model that is different from the present one in several technical respects. The theorem can be proved for the present model by making use of the trading sequence that Ostroy and Starr (1974) prove to exist. I will briefly outline this proof here. Ostroy and Starr construct a sequence of net trades, all of which have zero value at the Walras equilibrium price. After this trading has been completed, all traders hold their Walras-allocation commodity bundles. Inspection of Ostroy and Starr's proof shows that it is possible to specify that the net trades in the sequence should be distinct from one another. (This specification may require the sequence constructed by Ostroy and Starr to be lengthened.) Furthermore there is some  $\epsilon > 0$  such that no trader's commodity holding during trading sequence is ever closer than  $\epsilon$  from the boundary of the consumption set.

In the trading sequence, then, each trader makes a prescribed sequence of trades. Although Ostroy and Starr consider a deterministic trading sequence in an economy with finitely many traders, their argument can be adapted routinely to the present setting. (The assumption that  $\nu(A^t)$  is a rational number for each  $t$  facilitates this adaptation.) Here are the essential features of a strategy profile  $\sigma$  that implements the interior Walras allocation. At each trading date a trader proposes the first net trade in his prescribed sequence that

he has not yet executed, and he accepts the negative of that net trade if it is offered by his trading partner. (When a net trade proposed in accordance with  $\sigma$  is accepted, the proposer must be of the type specified by Ostroy and Starr since all of the net trades between different pairs of types are distinct from one another.) Eventually every trader is randomly paired with all of the partners who are appropriate to complete his prescribed trading sequence. The market clears, and every trader eventually holds his Walras-allocation commodity bundle. Because all of the net trades offered on the equilibrium path keep a trader within his budget set and no offer that would move the proposer's commodity holding out of his budget set is ever accepted, and because the Walras-equilibrium commodity bundle is each trader's most preferred commodity bundle within his budget set, a trader has no incentive to make out-of-equilibrium proposals or to make out-of-equilibrium responses to proposals that occur with positive probability on the equilibrium path. As far as histories that occur with positive probability on the equilibrium path, then, this strategy profile satisfies the criterion for being a perfect Bayesian equilibrium. The specification of  $\sigma(h)$  for out-of-equilibrium histories can be made in such a way that  $\sigma$  is indeed a perfect Bayesian equilibrium. *Q.E.D.*

## 5. Special classes of perfect Bayesian equilibria

If  $z$  is a net trade and  $h \in H^*$  is a history, then define  $\sigma^{h,z}$  to be the strategy profile that is identical to  $\sigma$  except that  $z \in \sigma_2^{h,z}(h) \iff z \notin \sigma_2(h)$ . If  $z$  is offered to a trader with history  $h$ , then perfect Bayesian equilibrium requires that  $\bar{U}((h, z), \sigma^{h,z}) \leq \bar{U}((h, z), \sigma)$ . At this point, I will define several properties that a strategy profile might possess. I will appeal to these properties later in the paper to establish the strict version of this inequality. It is a plausible conjecture, although not a proven fact, that Theorem 2 could be proved by construction of a strategy profile that possesses all of these properties.

A complication that I will assume not to occur is that  $\sigma^{h,z}$  might specify an infeasible net trade and thus not be an element of  $\Sigma$ . Specifically, it is possible that  $z \notin Z^h$  and that  $X^*(h) + z \notin X$ . Then  $z$  cannot feasibly be accepted, so  $\sigma^{h,z} \notin \Sigma$ . (It is also possible for

changing the response to a proposal to  $z$  to affect the feasibility  $\sigma$  at subsequent histories.) The following property of a strategy profile rules out this possibility for sufficiently small net trades. Define  $\sigma$  to be robust at  $h$  if, for some  $\epsilon > 0$ ,  $\sigma^{h,z} \in \Sigma$  for every  $z$  satisfying  $\|z\| < \epsilon$ . Also, define  $\sigma$  to be uniformly interior if, for some  $\epsilon > 0$  and for every  $h \in H^*$ ,  $X^*(h)$  has distance at least  $\epsilon$  from the exterior of  $\mathbf{X}$ .

Suppose that  $\sigma$  is uniformly interior and implements a full-information Walras allocation. If a trader deviates from  $\sigma$  by accepting an infinitesimal proposal  $z$ , then the microeconomic theory of the consumer suggests that he can subsequently do no better than to pretend that he had not deviated to and “carry along” the discrepancy of  $z$  in his commodity holdings from what they would have been had he not deviated. Formally, let  $\phi(h, z, \sigma) = 1$  if  $z \in \sigma_2(h)$  and  $\phi(h, z, \sigma) = -1$  if  $z \notin \sigma_2(h)$ , and define  $\sigma$  to possess the envelope property at  $h$  if it is robust and uniformly interior and, for almost all traders  $i$  of type  $\tau(h)$ ,

$$\begin{aligned} \lim_{\epsilon \searrow 0} \epsilon^{-1} [\bar{U}^{\tau(h)}(h, \sigma) - \bar{U}^{\tau(h)}(h, \sigma^{h, \epsilon z})] \\ = \phi(h, z, \sigma) \int_{\Omega} \nabla U^{\tau(h)}(\chi(i, \omega, \sigma))(z) d\Pi(h)(\omega). \end{aligned} \tag{1}$$

In the strategy profiles constructed in Gale (1986) and discussed in the proof of Theorem 2, a trader could acquire his Walras-allocation consumption bundle almost surely even if he were to deviate by refusing to participate in trade for an arbitrarily long (but finite) amount of time. In general, define a perfect Bayesian equilibrium  $\sigma$  to display persistent opportunities at  $h \in \mathcal{H}_n^t$  if, for every  $m \in \mathbb{N}$  and for every atom  $F \in \mathcal{F}_t$ , there is a deviation  $\sigma'$  after  $h$  from  $\sigma$  such that  $\bar{U}(h, \sigma') = \bar{U}(h, \sigma)$  and  $\forall m' \leq m \forall h \in H_{n+m'}^t, \sigma(h) = (0, \emptyset)$ .

## 6. Definition of a test portfolio

The remainder of this paper will be concerned with the analysis of responses to deviations from a perfect Bayesian equilibrium  $\sigma$  that implements a full-information Walras allocation. The deviations in question will be proposals of net trades that are designed to elicit the knowledge of trading partners regarding events in  $\mathcal{F}$ .

Suppose that a trader  $i$  is paired with a trading partner  $j$  at some date  $n$ , and that all traders' decisions are in accord with a strategy profile  $\sigma$ . Trader  $i$  does not observe  $j$ 's history  $h = \eta(j, n, \omega, \sigma) \in H_n(\sigma)$  directly. For some event  $B \in \mathcal{F}$ , suppose that  $i$  wants to discover whether or not  $j$  assigns probability 1 to  $B$ . Is there some net trade that  $i$  can offer to  $j$  which  $j$  will accept if  $\Pi(h)(B) = 1$  but reject if  $\Pi(h)(B) < 1$ ? If so,  $i$  could propose it to learn what  $j$  knows.

Such a net trade will be called a test portfolio for  $B$  at  $n$  in  $\sigma$  if such response decisions are based on strict expected-utility comparisons. (If indifference were possible, then it would be possible for a net trade—even the zero net trade—to be offered in order for acceptance or rejection to provide a purely conventional signal of the responder's knowledge.) That is,  $z$  is a test portfolio for  $B$  at  $n$  in  $\sigma$  if

$$\begin{aligned} \forall h \in H^n(\sigma) \ [ \sigma^{h,z} \in \Sigma \text{ and } \bar{U}((h, z), \sigma^{h,z}) < \bar{U}((h, z), \sigma) \text{ and} \\ [z \in \sigma_2(h) \iff \Pi(h, z)(B) = 1] ]. \end{aligned} \quad (2)$$

With this definition, an existence theorem can be stated for test portfolios in an important class of economies that I will now define.

## 7. Financial assets and spanning

At this point, I am going to modify the model introduced in Section 2 to apply specifically to trading in financial assets. I will call this model an asset-trading economy. Let the traders of each type  $t$  have a utility function  $V_t : \mathbb{R}_+ \rightarrow [0, 1]$  for wealth. Assume that  $V_t$  is strictly increasing, strictly concave, and continuously differentiable. Let each good  $h$  be an asset that is a state-contingent claim to wealth. Represent this state-contingent claim by a  $\mathcal{P}$ -measurable random variable  $a^h : \Omega \rightarrow \mathbb{R}$ . On this interpretation, a commodity bundle is a portfolio of assets. If a trader holds portfolio  $x \in \mathbb{R}^\ell$ , then his wealth in state  $\omega$  is  $W(x, \omega) = \sum_{h=1}^\ell x_h a^h(\omega)$ . The consumption set  $\mathbf{X}$  is the cone of portfolios that yield nonnegative wealth in every state of nature. The state-contingent utility functions  $U^t$  of Section 2 are to be interpreted as  $U^t(x, \omega) = V_t(W(x, \omega))$ .

Define the assets  $a^1, \dots, a^\ell$  to span  $\mathcal{P}$  if, when  $W(x, \cdot)$  is regarded as a function from  $\Omega$  to  $\mathbb{R}$  indexed by  $x$ , the set  $\{W(x, \cdot) | x \in \mathbb{R}^\ell\}$  is the set of all  $\mathcal{P}$ -measurable real-valued functions on  $\Omega$ . Since  $\mathcal{P}$  is finite, suppose that it has  $\kappa$  atoms  $\{P_1, \dots, P_\kappa\}$  which partition  $\Omega$ . For each  $k \leq \kappa$ , let  $\omega^k \in P_k$ . Define  $\psi(x) = (W(x, \omega^1), \dots, W(x, \omega^\kappa))$ . Then the spanning assumption is equivalent to the assumption that  $\psi$  is onto  $\mathbb{R}^\kappa$ .

The assumption that the assets span  $\mathcal{P}$  is a strong one. Another strong assumption will also be used here: that  $\mathcal{F} \subseteq \mathcal{P}$ . To see exactly why this assumption is restrictive, imagine that  $\mathcal{F}$  were generated by some random variable that traders were able to observe. It might be supposed that this random variable is informative about a financial asset because it is closely statistically correlated with its returns, even though it might not be a financial variable strictly speaking. For example, weather forecasts are useful for predicting the prices of agricultural futures contracts at maturity. However, if  $\mathcal{F} \subseteq \mathcal{P}$  and the assets span  $\mathcal{P}$ , then the random variable in question must coincide precisely with the state-contingent value of some portfolio. Thus the possibility that  $\mathcal{F}$  reflects information such as weather forecasts is virtually ruled out. That is, the formal assumption that  $\mathcal{F} \subseteq \mathcal{P}$  reflects the idea that the information dispersed among the various types of trader is information regarding asset returns *per se*.

## 8. Existence of a test portfolio

Now I show that, for perfect Bayesian equilibria possessing the envelope property in asset-trading economies with spanning, test portfolios exist for all events in  $\mathcal{F}$ . The proof requires a finite-dimensional, inner-product representation of the expected marginal utility of a net trade. For each  $h \in \mathcal{H}_n(\sigma)$ , let  $i^h$  be a trader of type  $\tau(h)$ . For each  $h \in \mathcal{H}_n(\sigma)$  and  $k \leq \kappa$ , define  $\Gamma_k^h = \int_{P_k} V'_{\tau(h)}(W(\chi(i^h, \omega, \sigma), \omega)) d\Pi(h)(\omega)$ .

**Lemma 2.** *For every history  $h \in H_n(\sigma)$  and every portfolio  $z \in \mathbb{R}^\ell$ ,  $\Gamma^h \cdot \psi(z) = \int_{\Omega} \frac{\partial}{\partial \epsilon} V_i(W(\chi(i^h, \omega, \sigma) + \epsilon z), \omega)|_{\epsilon=0} d\Pi(h)(\omega) = \int_{\Omega} \nabla U^{\tau(h)}(\chi(i^h, \omega, \sigma))(z) d\Pi(h)(\omega)$ , which is the marginal expected value to a trader of type  $\tau(h)$ , conditional on having observed*



history  $h$ , of adding a unit of portfolio  $z$  to the random portfolio that he will asymptotically hold according to  $\sigma$ .

**Proof:** The first equation is derived by applying the chain rule and expressing the result in terms of  $\Gamma^h$ . The second equation follows directly from the interpretation that  $U^t(x, \omega) = V_t(W(x, \omega))$  in an asset-trading economy. The integral is seen to be a marginal expected value by applying Leibniz' Rule for commuting integration with partial differentiation. *Q.E.D.*

Lemma 2 expresses  $\int_{\Omega} \nabla U^{\tau(h)}(\chi(i^h, \omega, \sigma))(z) d\Pi(h)(\omega)$ , which the envelope property associates with the marginal value to a trader with history  $h$  of accepting a "small" proposal  $z$ , as a multiplicatively separable function of  $h$  and  $z$ . This separability condition is a special feature of asset-trading economies which exchange economies in general do not possess. Now I will show, making essential use of the separability condition, that the incentive-compatibility condition (2) in the definition of a test portfolio can be satisfied in an asset-trading economy where the spanning condition holds.

**Theorem 3.** *Consider a perfect Bayesian equilibrium  $\sigma$  in an asset-trading economy where the assets span  $\mathcal{P}$  and where  $\ell \geq 3$ . If  $\sigma$  possesses the envelope property at every  $h \in H_n(\sigma) \times \{\text{propose}\}$  and  $B \in \mathcal{F} \subseteq \mathcal{P}$ , and displays persistent opportunities at every history at date  $n$ , then there is a portfolio  $z$  that satisfies condition (2).*

**Proof:** Define  $Q = \{\Gamma^h | h \in H_n(\sigma) \times \{\text{propose}\} \text{ and } \Pi(h)(B) = 1\}$  and define  $R = \{\Gamma^h | h \in H_n(\sigma) \times \{\text{propose}\} \text{ and } \Pi(h)(B) < 1\}$ . The convex hulls of  $Q$  and  $R$  are disjoint on account of two factors. First,  $\mathcal{F} \subseteq \mathcal{P}$ . Second, every  $\Gamma^h \in Q$  satisfies  $\Gamma_k^h = 0$  for every  $k$  such that  $B \cap P_k = \emptyset$  while no  $\Gamma^h \in R$  satisfies this condition. These two convex hulls are disjoint polyhedral convex sets by Lemma 1, so a separation theorem guarantees that there is a vector  $y \in \mathbb{R}^{\kappa}$  such that

$$\forall \Gamma^h \in Q \quad y \cdot \Gamma^h > 0 \quad \text{and} \quad \forall \Gamma^h \in R \quad y \cdot \Gamma^h < 0. \quad (3)$$

The spanning condition implies that there is a net trade  $z^*$  such that  $\psi(z^*) = y$ . Since the set of  $y$  satisfying (3) is open and the set of net trades in  $\sigma_1(H_n(\sigma))$  is finite (by Lemma 1),  $z^*$  can be chosen not to be a scalar multiple of any of these net trades since  $\ell \geq 3$ . If  $h = (h', z^*)$ , then  $\Pi(h) = \Pi(h', \text{propose})$  by the assumption (at the end of Section 3) regarding how deviations are ignored in probability updating. Making these substitutions in (3) and applying Lemma 2 yields that

$$\int_{\Omega} \nabla U^{\tau(h)}(\chi(i^h, \omega, \sigma))(z^*), \omega) d\Pi(h', \text{propose})(\omega) > 0 \text{ if } \Gamma^h \in Q$$

and

$$\int_{\Omega} \nabla U^{\tau(h)}(\chi(i^h, \omega, \sigma))(z^*), \omega) d\Pi(h', \text{propose})(\omega) < 0 \text{ if } \Gamma^h \in R.$$

These two inequalities together with (1) establish that  $z = \epsilon z^*$  satisfies (2) if  $\epsilon > 0$  is taken to be sufficiently small.

In this proof, computations for history  $(h, \text{propose})$  have been used in place of those for the actual history  $(h, z)$ . The assumption that  $\sigma$  displays persistent opportunities, along with the envelope property, can be used to justify this substitution for a net trade  $z$  that is sufficiently small. *Q.E.D.*

## 9. Conclusion

Theorem 2 shows that, under assumptions having to do with asset structure and spanning, incentive-compatible elicitation of trading partners' knowledge is feasible. Moreover there is a sense in which a trader need not be very sophisticated in order to interpret trading partners' responses to his elicitation efforts. Specifically, although learning from experience in the market generally requires the application of Bayes' Theorem, this form of elicitation works in a more statistically straightforward way.

One might conjecture that there would be a perfect Bayesian equilibrium in which each trader would propose test portfolios until he had become fully informed (*i.e.* until some atom of  $\mathcal{F}$  had posterior probability 1), after which he would trade as in Theorem 1 to acquire his Walras-equilibrium commodity bundle. To the extent that traders' preferences for net trades are close to being identical contingent on their beliefs, though, the following problem exists. There is a sense in which it is disadvantageous for a trader to propose a test portfolio. If the proposal is accepted, it must be because the responder knows that it is desirable. When this acceptance is given, the offerer acquires the knowledge on which it is based. Therefore the offerer will realize that the portfolio is desirable and will regret having given away so much. Since the test portfolio constructed in the proof of Theorem 2 can be halved to produce another test portfolio for the same event, offering a nonzero test portfolio can never be optimal. The zero portfolio is not a test portfolio, though, since the strict-inequality conditions (2) are not satisfied. A perfect Bayesian equilibrium can be constructed in which uninformed traders propose the zero portfolio and in which the acceptance or rejection of that portfolio at each date depends on whether or not the responder's posterior probability of an atom of  $\mathcal{F}$  assigned to that date is 1. The responder, who will never meet the proposer again or even have any further indirect connection with the proposer, has no incentive to "lie" about his knowledge. However, it may seem arbitrary to suggest that traders tell the truth to one another because they are indifferent between truthful and false reporting.

Eric Maskin has pointed out to me that this last statement should not be interpreted as meaning that there necessarily exist "sensible" perfect Bayesian equilibria in which the zero portfolio is offered to obtain information, but in which "untruthful" responses are sometimes received. Rather than offering the zero portfolio and receiving an untrustworthy response, the proposer could offer a test portfolio of negligible size by Theorem 3. The trader making this proposal would have only a trivial loss to regret if it were accepted, and it seems intuitive that he would receive a response that would definitely be "truthful." (If this intuition were correct, then (2) would hold and the portfolio would be a test portfolio.) Thus a trivial expected loss would be suffered for the sake of an appreciable expected gain. Considering this trade-off, to propose the zero portfolio could not be consistent with

equilibrium if a “truthful” response would not be received with certainty.

In combination with the conjecture that all traders must become fully informed in every perfect Bayesian equilibrium, the results of Gale (1986) and especially of McLennan and Sonnenschein (1989) would suggest that full-information Walras allocations are the only implementable ones. To the extent that one considers the form of communication in perfect Bayesian equilibrium that has just been discussed to be counterintuitive as a description of information transmission in an anonymous trading environment (because traders behave “too well” towards strangers from whom they will receive no reward), such a characterization theorem constitutes a puzzle.

One solution of this puzzle would be to relax the concept of perfect Bayesian equilibrium to tolerate small deviations from rationality. (A closely related solution, in which traders use nonsequential search strategies rather than conducting optimal sequential search for information, has been proposed by Gale (1987).) Because test portfolios can be made arbitrarily small, such a relaxation would allow them to be proposed in equilibrium. However, if the test portfolios are small, then “wrong” decisions about whether or not to accept them may also be consistent with equilibrium. Thus the full-information Walras allocations might well be a proper subset of the allocations that could be implemented with respect to such an equilibrium concept.

## References

- Gale, D., (1986), "Bargaining and Competition" (Parts I and II), *Econometrica*, **54**, 787-818.
- Gale, D., (1987), "Information Revelation in a Model of Nonsequential Search," mimeo, Boston University.
- Green, E., (1988), "Individual-Level Randomness in a Nonatomic Population," mimeo, University of Minnesota.
- McLennan, A., and H. Sonnenschein, (1989), "Sequential Bargaining as a Noncooperative Foundation for Walrasian Equilibrium," mimeo, University of Minnesota.
- Ostroy, J., and R. Starr, (1974), "Money and the Decentralization of Exchange," *Econometrica*, **42**, 1093-1113.
- Tirole, J., (1988), *The Theory of Industrial Organization*, Cambridge, USA: MIT Press, 1988.
- Wolinsky, A., (1990), "Information Revelation in a Market with Pairwise Meetings," *Econometrica*, **58**, 1-23.