Research Paper No. 2053

Information Aggregation in Dynamic Markets with Strategic Traders

Michael Ostrovsky

March 2009

RESEARCH PAPER SERIES

STANFORD
GRADUATE SCHOOL OF BUSINESS

Information Aggregation in Dynamic Markets with Strategic Traders

Michael Ostrovsky*
Stanford University

March 11, 2009

Abstract

This paper studies information aggregation in dynamic markets with a finite number of partially informed strategic traders. It shows that for a broad class of securities, information in such markets always gets aggregated. Trading takes place in a bounded time interval, and in every equilibrium, as time approaches the end of the interval, the market price of a "separable" security converges in probability to its expected value conditional on the traders' pooled information. If the security is "non-separable," then there exists a common prior over the states of the world and an equilibrium such that information does not get aggregated. The class of separable securities includes, among others, Arrow-Debreu securities, whose value is one in one state of the world and zero in all others, and "additive" securities, whose value can be interpreted as the sum of traders' signals.

*ostrovsky@gsb.stanford.edu. I am grateful to Jeremy Bulow, Yossi Feinberg, Emir Kamenica, Konstantin Makarychev, Yury Makarychev, Sergey Norin, Rahul Sami, Michael Schwarz, Robert Wilson, Bumin Yenmez, and the participants of Stanford GSB Economics and Finance seminars for helpful comments, suggestions, and discussions.

1 Introduction

The idea that financial markets have the ability to aggregate and reveal dispersed information is an important part of economic thinking. The intuition behind this idea is arbitrage: if the price of a security is wrong, an informed trader will have an incentive to buy or sell this security, thus bringing the price closer to the correct value. This intuition is very compelling when one or more traders are fully informed and know the value of the security. It is also compelling in many cases where each trader is small relative to the market and behaves, in essence, non-strategically, ignoring the effect his trading has on prices and thus revealing all his information. But what happens when there is a small, finite number of large, strategic players, and none of them is fully informed about the value of the security? What if one trader has perfect information about one part of a company and another trader has perfect information about the rest of the company? Will the stock price reflect the true value of the company that the traders could estimate by pooling their information? Or is there a chance that the price will be off? What happens when information is dispersed among many agents in the economy and their knowledge structure is more complex?

This paper shows that for a broad class of securities, information in dynamic markets with partially informed strategic traders always gets aggregated. Trading takes place in a bounded time interval, and in every equilibrium, as time approaches the end of the interval, the market price of a "separable" security converges in probability to its expected value conditional on the traders' pooled information. A security is "separable" if, roughly, for every non-degenerate prior belief about the states of the world, there exists a trader who with positive probability receives an informative signal. If the security is "non-separable," then there exists a prior and an equilibrium such that information does not get aggregated.

The question of information revelation and aggregation in markets has attracted the attention of many economists, beginning with Hayek (1945). Grossman (1976) formally shows that in a market equilibrium, the resulting price aggregates information dispersed among n-types of informed traders, each of whom gets a "piece of information." In his model, individual traders are small relative to the market, strategic foundations for players' behavior are lacking, and the results rely on particular functional forms (e.g., i.i.d. normal errors in signals received by the players; normal prior; etc.). Radner (1979) introduces the concept of Rational Expectations Equilibrium (REE) and shows that generically, a fully revealing REE exists, with prices aggregating all information dispersed among traders. Radner's paper, however, also lacks strategic foundations. A series of papers explore the question of convergence to REE in various dynamic processes (see, e.g., Hellwig, 1982, and Dubey, Geanakoplos, and Shubik, 1987, for the models of centralized trading and Wolinsky, 1990, and Golosov, Lorenzoni, and Tsyvinski, 2008, for the models of decentralized trading). In all of these papers, however, it is assumed that each trader is small relative to the market and therefore ignores the effect of his behavior on the evolution of the trading process, as a result behaving non-strategically along at least one dimension. Proper strategic foundations for the concept of perfect competition with differentially informed agents are offered by the stream of literature studying bidding behavior in single and double auctions (Wilson, 1977; Milgrom, 1981; Pesendorfer and

Swinkels, 1997; Kremer, 2002; and Reny and Perry, 2006). Information aggregation results in these papers, however, rely on the assumption that the market is large, i.e., the number of bidders goes to infinity and individual traders become small relative to the market. They also rely on various symmetry and independence assumptions. No such assumptions are made in the current paper, and the number of traders is finite and fixed.

Kyle (1985) offers a model of dynamic, strategic insider trading, in which the single informed trader takes into account the non-negligible impact of his actions on market prices. In the continuous version of the model, as time approaches the end of the trading interval, the price of the traded security converges to its true value known by the insider. Foster and Viswanathan (1996) and Back, Cao, and Willard (2000) extend the model to the case of multiple, differentially informed strategic traders. In the continuous case, the price of the traded security converges to its expected value conditional on the traders' pooled information. In the discrete case with a finite number of trading periods, convergence is approximate. These models rely on very special functional form assumptions (symmetry, normality, etc.), which allow the authors to construct explicit formulas for particular ("linear") perfect Bayesian equilibria. Laffont and Maskin (1990) criticize this reliance of the results of Kyle (1985) on linear trading strategies; argue that such models inherently have multiple equilibria; present a model of a trading game with a single informed trader and multiple equilibria, in some of which the informed trader's information is not revealed; and conclude that "in a model in which private information is possessed by a trader who is big enough to affect prices, the information efficiency of prices breaks down" and "the efficient market hypothesis may well fail if there is imperfect competition." The results of the current paper show that the conclusions of Kyle (1985), Foster and Viswanathan (1996), and Back, Cao, and Willard (2000) regarding the convergence of the price of a security to its expected value conditional on the traders' pooled information do not in fact depend on the specific functional form assumptions or on the choice of equilibrium: if the traded security is separable, its price converges to its expected value conditional on the pooled information in every perfect Bayesian equilibrium. In the case of a single informed trader, as in Lafont and Maskin (1990), every security is separable, and so information always gets aggregated. The conclusions of Laffont and Maskin are driven by their assumption that trading takes place only once, not by the greater generality of the model they consider. In the case of multiple partially informed traders, the securities considered in Foster and Viswanathan (1996) and Back, Cao, and Willard (2000) have payoffs that are linear in traders' signals, and so as the results of this paper show, information about such securities always gets aggregated as well.

An alternative way to model (and, in fact, to organize in practice) the dynamic trading process is offered by Hanson (2003): the market scoring rule (MSR). In MSR games, there are no noise or liquidity traders and no strategic market makers; the only players are the strategic partially informed traders. There is also an automated market maker. This market maker, in expectation, loses money (though at most a finite, ex ante known amount), facilitating trade and price discovery. (Without a "source" of profits, there would be no trading; see Milgrom and Stokey, 1982, and Sebenius and Geanakoplos, 1983.) Trading proceeds as follows. The uninformed market maker makes an initial, publicly observed prediction about the value of a security. The first informed strategic player can revise that number and make his own prediction, which is also observed by everyone. Then the second player can further modify the prediction, and so on until the last player, after which the first player can again modify the prediction, and the cycle repeats an infinite number of times. The fact that there is an infinite number of trading periods does not mean that the game never ends. Rather, it is a convenient discrete analogue of continuous trading, with trades taking place at times $t_0 < t_1 < \dots$ in a bounded time interval. Sometime after trading is over, the true value of the security is revealed, and each prediction is evaluated according to a strictly proper scoring rule s (e.g., under the quadratic scoring rule, each prediction is penalized by the square of its error; see Section 2.2 for further details). The payoff of a player from each revision is the difference between the score of his prediction and the score of the previous prediction—in essence, the player "buys out" the previous prediction and replaces it with his own. The total payoff of a player in the game is the sum of payoffs from all his revisions. Players are risk-neutral. The discounted MSR (Dimitrov and Sami, 2008) is very similar, except that the total payoff of a player is equal to the discounted sum of payoffs from all his revisions, where the payoff from a revision made in period t_k is multiplied by β^k for some $\beta < 1$.

MSR games look very different from the typical models of trading studied in Economics and Finance, so a natural question is why one would want to study information aggregation in such games, especially since, as I show below, the results are the same as in the more conventional model of trading based on Kyle (1985). There are several reasons. First, MSR can be reinterpreted as a trading process in which agents buy and sell shares (see Hanson, 2003; Pennock, 2006; and Chen and Pennock, 2007), with an automated market maker who continuously adjusts the prices of securities as traders buy and sell them; the "sponsor" of the market can also be reinterpreted as a liquidity (or noise, or risk averse) trader who is willing to trade once even though he knows that in expectation, he will bear some losses. Second, market scoring rules are now widely used in practice to organize prediction markets (firms like Inkling Markets, Consensus Point, and Crowdcast, among others, run MSR-based internal prediction markets for a number of large companies), so it is important to understand whether such mechanisms do in fact aggregate information. Finally, the advantage of using an MSR game as a modeling device is its transparency. MSR-based models bring the question of information aggregation to the forefront, eliminating the need to consider noise traders, strategic market makers, and other aspects of more typical trading games that obscure the intuition behind the main result. Of course, it is important to verify that the results do not in fact depend on the simplifications offered by the MSR, and so this paper contains information aggregation results for both MSR games and trading processes based on Kyle (1985). While the key ideas behind these results are similar, the proof for the latter case is more complicated and less transparent than that for the former.

Two recent papers have studied the equilibrium behavior of traders in MSR games. Chen et al. (2007) consider undiscounted games based on a particular scoring rule—logarithmic (see Section 2.2). In their model, the security can take one of two different values, and the number of revisions is finite. They find that if traders' signals are independent conditional on the value of the security, then it is an equilibrium for each trader in each period to behave myopically, i.e., to make the prediction equal to his posterior belief. They also provide an example of a market in which signals are not conditionally independent and one of the traders has an incentive to behave non-myopically; however, they do not study information aggregation and convergence properties in the cases where myopic behavior is not optimal. Dimitrov and Sami (2008) also consider games based on the logarithmic scoring rule. In their models, in contrast to Chen et al., traders observe independent signals. Each realization of the vector of signals corresponds to a particular value of the security. The number of trading periods is infinite. Dimitrov and Sami find that in that case, in the MSR game with no discounting, myopic behavior is generically not an equilibrium and, moreover, there is no equilibrium in which all uncertainty is guaranteed to get resolved after a finite number of periods. They then introduce a two-player, two-signal MSR game with discounting, and prove that in that game, information gets aggregated in the limit, under the additional assumption that the "complementarity bound" of the security is positive. They report that based on their sample configurations, the bound is not always zero, but do not provide any sufficient conditions for it to be positive. In contrast to Chen et al. (2007) and Dimitrov and Sami (2008), the current paper's information aggregation results (1) do not rely on the independence or conditional independence of signals, allowing instead for general information structures with any number of players; (2) do not depend on discounting; and (3) provide a sharp characterization of securities for which information always gets aggregated and those for which under some priors, price may not converge to the expected value conditional on the traders' pooled information.

The remainder of this paper is organized as follows. Section 2 describes the model of information in the market, two models of trading, and the definitions of information aggregation and separability. Section 3 presents the main result. Section 4 discusses the separability assumption. Section 5 concludes.

2 Setup

There are n players, $i = 1, \ldots, n$. There is a finite set of states of the world, Ω, and a random variable ("security") $X : \Omega \to \mathbb{R}$. As in Aumann (1976), each player i receives information about the true state of the world, $\omega \in \Omega$, according to partition Π_i of Ω (i.e., if the true state is ω, player i observes $\Pi_i(\omega)$). For notational convenience, without loss of generality, assume that the join (the coarsest common refinement) of partitions $\Pi_1, \ldots, \Pi_n$ consists of singleton sets; i.e., for any two states $\omega_1 \neq \omega_2$ there exists player i such that $\Pi_i(\omega_1) \neq \Pi_i(\omega_2)$. $\Pi = (\Pi_1, \ldots, \Pi_n)$ is the *partition structure*. Players have a common prior distribution P over states in Ω.

2.1 Trading: A Model Based on Kyle (1985)

In this model, trading is organized as follows. At time $t_0 = 0$, nature draws a state, ω^*, according to P, and all strategic players i observe their information $\Pi_i(\omega^*)$. At time $t_1 = \frac{1}{2}$, each strategic player i chooses his demand d_1^i. At the same time, there is demand u_1 from noise traders, drawn randomly from the normal distribution with mean zero and variance $t_1 - t_0 = \frac{1}{2}$. Competitive market makers observe the aggregate demand $\sum_i d_1^i + u_1$, form their posterior beliefs about the true state of the world, and set market price y_1 equal to the expected value of the security conditional on that posterior belief. The market clears, and all traders observe price y_1 and aggregate volume $\sum_i d_1^i + u_1$. At time $t_2 = \frac{3}{4}$, the next auction takes place, with each strategic player i choosing demand d_2^i and demand from noise traders u_2 drawn randomly from $N(0, t_2 - t_1 = \frac{1}{4})$. Subsequently, auctions are held at times $t_k = 1 - \frac{1}{2^k}$ with demand from noise traders drawn from $N(0, \frac{1}{2^k})$. The value of the security, $x^* = X(\omega^*)$, is revealed at some time $t^* > 1$. Trader i's payoff is equal to $\sum_{k=1}^{\infty} d_k^i (x^* - y_k)$. The resulting game is denoted $\Gamma^K(\Omega, \Pi, X, P)$.

2.2 Trading: Market Scoring Rules

Under MSR, trading is organized as follows. At time $t_0 = 0$, nature takes a random draw and selects the state, ω^*, according to P. The uninformed market maker makes the initial prediction $y_0 \in \mathbb{R}$ about the value of X (a natural initial value for y_0 is the unconditional expected value of X under P, but it could also be equal to any other real number). At time $t_1 > t_0$, player 1 makes a "revised prediction," y_1. At time $t_2 > t_1$, player 2 makes his prediction, y_2, and so on. At time t_{n+1}, player 1 moves again and makes his new forecast, y_{n+1}, and the whole process repeats until time $t_\infty \equiv \lim_{k\to\infty} t_k = 1$, with players taking turns revising predictions. All predictions y_k are observed by all players. The action space is bounded, but the bounds are wide enough to allow for any prediction consistent with random variable X, i.e., each y_k is a number in an interval $[\underline{y}, \overline{y}]$, where $\underline{y} \leq \min_{\omega\in\Omega} X(\omega) \leq \max_{\omega\in\Omega} X(\omega) \leq \overline{y}$.

At time $t^* > 1$, the true value $x^* = X(\omega^*)$ of the security is revealed. The players' payoffs are computed according to a market scoring rule that is based on a strictly proper single-period scoring rule s. More formally, a single-period scoring rule is a function $s(y, x^*)$, where x^* is a realization of a random variable and y is a prediction. The scoring rule is proper if for any random variable X, the expectation of s is maximized at $y = E[X]$. It is strictly proper if $y = E[X]$ is the unique prediction maximizing the expected value of s. Examples of strictly proper scoring rules include the quadratic scoring rule ($s(y, x^*) = -(x^* - y)^2$), due to Brier (1950), and, when random variable X is bounded (which, of course, is the case here), the logarithmic scoring rule ($s(y, x^*) = (x^* - a)\ln(y - a) + (b - x^*)\ln(b - y)$, for some $a < \underline{y}$ and $b > \overline{y}$), due to Good (1952). Assume that $s(y, x^*)$ is continuous and bounded on $[\underline{y}, \overline{y}] \times [\min_{\omega\in\Omega} X(\omega), \max_{\omega\in\Omega} X(\omega)]$.

Under the basic MSR (introduced by Hanson, 2003, though the idea of repeatedly using a proper scoring rule to help forecasters aggregate information goes back to McKelvey and Page, 1990), players get multiple chances to make predictions, and are paid for each revision. Specifically, for each revision of the prediction from y_{k-1} to y_k, player i is paid $s(y_k, x^*) - s(y_{k-1}, x^*)$. Of course, this number can turn out to be negative, but each player can guarantee himself a zero payment for a revision by simply setting $y_k = y_{k-1}$, i.e., by not revising the forecast. Note also that if each player behaves myopically in each period, the prediction that he will make is his posterior belief about the expected value of the security, given his initial information and the history of revisions up to that point, and thus the "game" turns into the communication process of Geanakoplos and Polemarchakis (1982).

A slight modification of the game above, introduced by Dimitrov and Sami (2008), is a discounted MSR: it is the same as the basic MSR, except that the payment for the revision from y_{k-1} to y_k is equal to $\beta^k(s(y_k, x^*) - s(y_{k-1}, x^*))$, $0 < \beta \leq 1$. When $\beta = 1$, this rule coincides with the basic MSR. The total payoff of each player is the sum of all payments for revisions. The players are risk-neutral. The resulting game is denoted $\Gamma^{MSR}(\Omega, \Pi, X, P, y_0, \underline{y}, \overline{y}, s, \beta)$.

2.3 Information Aggregation

Definition 1 *In a perfect Bayesian equilibrium of game Γ^K or Γ^{MSR},* information gets aggregated *if sequence y_k converges in probability to random variable $X(\omega^*)$.*

Since the number of possible states of the world is finite, this definition is equivalent to saying that for any $\epsilon > 0$ and $\delta > 0$, there exists K such that for any $k > K$, for any realization of the nature's draw $\omega^* \in \Omega$, the probability that $|y_k - X(\omega^*)| > \epsilon$ is less than δ.

2.4 Separability

Consider the following example from Geanakoplos and Polemarchakis (1982).

Example 1 *There are two agents, 1 and 2, and four states of the world, $\Omega = \{A, B, C, D\}$. The prior is $P(\omega) = \frac{1}{4}$ for every $\omega \in \Omega$. The security is $X(A) = X(D) = 1$ and $X(B) = X(C) = -1$. Partitions are $\Pi_1 = \{\{A, B\}, \{C, D\}\}$ and $\Pi_2 = \{\{A, C\}, \{B, D\}\}$.*

In the example, by construction, it is common knowledge that each player's expectation of the value of the security is zero, even though it is also common knowledge that the actual value of the security is not zero, and that the traders' pooled information would be sufficient to determine the security's value. Thus, even if the traders repeatedly and truthfully announce their posteriors, as in Geanakoplos and Polemarchakis (1982), they will never learn the true value of the security. Dutta and Morris (1997) and DeMarzo and Skiadas (1998, 1999) study competitive equilibria with information structures similar to that of Example 1 and show that they give rise to the generic existence of "Common Beliefs Equilibria"/"partially informative REE" in which, in contrast to the fully revealing REE of Radner (1979), equilibrium prices do not fully aggregate traders' information.

DeMarzo and Skiadas (1998, 1999) show that competitive equilibrium prices are guaranteed to fully aggregate information if and only if securities and information structures like that of Example 1 are ruled out, i.e., in their terminology, the function mapping traders' signals into fully informative equilibrium prices is "separably oriented." Adapted to the current paper's setup, this condition translates into the following definition of separability, which, as Theorem 1 below shows, plays a key role in this paper's results.

Definition 2 *Security X is non-separable under partition structure Π if there exist distribution P on the underlying state space Ω and value $v \in \mathbb{R}$ such that:*

1. *$P(\omega)$ is positive on at least one state ω in which $X(\omega) \neq v$;*
2. *For every player i and every state ω with $P(\omega) > 0$,*

$$E[X|\Pi_i(\omega)] = \frac{\sum_{\omega' \in \Pi_i(\omega)} P(\omega')X(\omega')}{\sum_{\omega' \in \Pi_i(\omega)} P(\omega')} = v.$$

Otherwise, security X is separable.

Note that non-separable securities are not degenerate (e.g., for any security with payoffs close to the ones in Example 1, there is a distribution P that would satisfy the requirements in Definition 2, and thus all such securities are non-separable). Note also that if there is only one informed trader in the market, then every security is separable. Section 4 discusses the separability condition in more detail.

3 Main Result

The main result of this paper is that information about separable securities always gets aggregated, while for non-separable securities that is not the case.

Theorem 1 *Consider state space Ω, security X, and partition structure Π.*

1. *If security X is separable under Π, then for any prior distribution P:*
   - *in any PBE of the corresponding game Γ^K information gets aggregated;*
   - *for any strictly proper scoring rule s, initial value y_0, bounds $\underline{y}$ and $\overline{y}$, and discount factor $\beta \in (0, 1]$, in any PBE of the corresponding game Γ^{MSR} information also gets aggregated.*
2. *If security X is non-separable under Π, then there exists prior P such that:*
   - *there exists a PBE of the corresponding game Γ^K in which information does not get aggregated;*
   - *for any s, y_0, $\underline{y}$, $\overline{y}$, and β, there exists a PBE of the corresponding game Γ^{MSR} in which information does not get aggregated*

Proof. The proof of the second statement, that for non-separable securities information does not always get aggregated, is straightforward. Consider prior P and value v that satisfy requirements 1 and 2 of Definition 2. Then in game Γ^{MSR}, it is an equilibrium for all traders to make the same prediction $y_k = v$ in every period t_k after any history, and in game Γ^K, it is an equilibrium for the traders to always submit zero demand and for the competitive market makers to set price $y_k = v$ in every period t_k after any history (beliefs in the equilibria of both games are never updated from the priors). In these equilibria, information does not get aggregated. The proof of the first statement is in the Appendix. ■

The intuition behind the proof of the first statement of Theorem 1 is as follows. In any equilibrium of game Γ^{MSR}, consider an uninformed outside observer who has the same prior as the informed traders, receives no direct information about the state of the world, and observes all predictions made by the traders. Consider the stochastic process that corresponds to the observer's vector of posterior beliefs about the likelihoods of the states of the world after each revision. By construction, this process is a bounded martingale, and therefore, by the martingale convergence theorem, converges to some vector-valued random variable Q_∞. If Q_∞ puts positive weights on two states of the world in which the value of the security is different, then separability implies that there is a player who can, in expectation, make a non-vanishing positive profit by revising the prediction in any sufficiently late period. This, in turn, can be shown to imply that the player is not maximizing his payoff (because he never actually makes that deviation), which is impossible in equilibrium. Thus, with probability 1, Q_∞ has to put all weight on states in which the value of the security is the same. Since the beliefs have to be on average correct, this is only possible if this value is the correct one with probability 1. Now, if Q_∞ does put all weight on the states with the correct value of the security, but the prediction does not converge to the same value, then even the uninformed observer could make a profitable revision in infinitely many periods, and thus any informed player could make such revisions as well, again contradicting the assumption of profit-maximizing behavior. Therefore, the outside observer's posterior beliefs, in the limit, have to put all weight on the states with the correct value of the security, and the prediction has to converge to the same value. In game Γ^K, the intuition is similar, although the statement that the lack of information aggregation implies the existence of non-vanishing profitable arbitrage for at least one trader becomes more delicate and requires a more elaborate proof.

4 Separable Securities

In light of Theorem 1, it is important to understand the restrictions the separability condition places on securities. This section describes several important classes of separable securities and gives an alternative "dual" characterization of the condition. Some of these results are equivalent to or are corollaries of the more general results about the "separably oriented" condition in DeMarzo and Skiadas (1998, 1999), but in the current paper's setting they also have short, self-contained proofs, presented below for completeness.

4.1 Dual Characterization of Separability

While Definition 2 is very intuitive, there is an alternative, "dual" way to characterize separable securities. This alternative characterization is convenient in applications, as Corollaries 1 and 2 below illustrate.

Theorem 2 *Security X is separable under partition structure Π if and only if for every $v \in \mathbb{R}$, there exist numbers λ_π corresponding to the elements π of partitions Π_i of all players such that for every state ω with $X(\omega) \neq v$,*

$$(X(\omega) - v)\left(\sum_{i=1,\ldots,n} \lambda_{\Pi_i(\omega)}\right) > 0.$$

Proof. The "if" direction of the theorem is proved in Proposition 7 of DeMarzo and Skiadas (1998) using the following "adding-up" argument, which makes the condition more intuitive. Suppose X is non-separable and take P and v satisfying the requirements of non-separability. Take numbers λ_π as in the statement of the theorem. Consider the unconditional expectation $E[(X(\omega) - v)\sum_i \lambda_{\Pi_i(\omega)}]$ under P. On the one hand, by the choice of parameters λ_π, the expectation is strictly positive. On the other hand,

$$\begin{aligned} E[(X(\omega) - v)\sum_i \lambda_{\Pi_i(\omega)}] &= \sum_i E[(X(\omega) - v)\lambda_{\Pi_i(\omega)}] \\ &= \sum_i \sum_{\pi \in \Pi_i} P(\pi) E[(X(\omega) - v)\lambda_\pi | \pi] = 0, \end{aligned}$$

where the last equality follows from requirement 2 of Definition 2.

The "only if" direction is proved in the Appendix. ■

4.2 Order Statistics

The first corollary of Theorem 2 shows that securities that can be represented as order statistics of traders' signals (minimum, maximum, median, etc.) are separable.

Corollary 1 *Suppose security X can be expressed as an order statistic of players' signals: $X(\omega) = x_{(j)}(\omega)$ (i.e., the j^{th} lowest signal x), where $x_i(\omega) \equiv x_i(\Pi_i(\omega))$ is the "signal" observed by player i in state ω. Then X is separable.*

Proof. Take any $v \in \mathbb{R}$. For every i and ω, set $\lambda_{\Pi_i(\omega)}$ equal to 1 if $x_i(\omega) \geq v$ and to $-\frac{n-j+1}{j-0.5}$ if $x_i(\omega) < v$. Then $X(\omega) > v \Rightarrow x_{(j)} > v \Rightarrow \sum_i \lambda_{\Pi_i(\omega)} \geq (n - j + 1) - \frac{n-j+1}{j-0.5}(j - 1) > 0$ and $X(\omega) < v \Rightarrow x_{(j)} < v \Rightarrow \sum_i \lambda_{\Pi_i(\omega)} \leq (n - j) - \frac{n-j+1}{j-0.5} j < 0$. ■

Corollary 1 implies that any Arrow-Debreu security, i.e., random variable X that is equal to 1 in one state of the world and to 0 in all other states, is separable. Thus, by Theorem 1, information about Arrow-Debreu securities always gets aggregated. While the analysis of markets with multiple securities is beyond the scope of this paper, this result suggests that in complete markets, information also always gets aggregated.

4.3 Monotone Transformations of Additive Payoffs

The second corollary of Theorem 2 shows that monotone transformations of additive securities (e.g., additive securities, positive multiplicative securities, call or put options on additive or positive multiplicative securities, and so on) are separable, where a security is "additive" if it can be expressed as the sum of traders' signals. This assumption, of course, also includes a seemingly more general case of $X(\omega)$ being a linear function of the signals, such as the average, because signals can be rescaled, or a stochastically monotone function (Nielsen et al., 1990), for the same reason.

Corollary 2 *Suppose security X can be expressed as $X(\omega) = f\left(\sum_i x_i(\Pi_i(\omega))\right)$, where $x_i(\Pi_i(\omega))$ is the "signal" observed by player i in state ω and f is a monotone function. Then X is separable.*

Proof. Assume f is increasing (the other case is completely analogous) and continuous (since X takes only a finite number of values, this is w.l.o.g.). Take any $v \in [\min_{\omega\in\Omega} X(\omega), \max_{\omega\in\Omega} X(\omega)]$ (the case where v is outside this interval is trivial). Take any z such that $f(z) = v$. Setting $\lambda_{\Pi_i(\omega)} = x_i(\Pi_i(\omega)) - \frac{z}{n}$ for every player i and state ω and applying Theorem 2 proves the result. ■

Thus, information about securities with additive payoffs and their monotone transformations always gets aggregated, for every distribution of priors, correlation structure of signals, and so on.

4.4 Increasing Payoffs

A generalization of securities with additive payoffs are securities with "increasing" payoffs, which place only one restriction on the information structure: each player's signals can be ordered from better to worse. It turns out that for such securities, separability generally depends on the number of traders.

Claim 1 *Suppose each player's information $\Pi_i(\omega)$ can be interpreted as a signal $x_i(\omega) \equiv x_i(\Pi_i(\omega)) \in \mathbb{R}$ in such a way that $X(\omega_1) \geq X(\omega_2)$ whenever $x_i(\omega_1) \geq x_i(\omega_2)$ for all i. If there are two traders in the market, then any such security is separable.*

Proof. Suppose the security is non-separable and take P and v that verify its non-separability. Take the lowest signal x_1 of player 1 and the corresponding element π_1 of partition Π_1 such that there exist states $\omega \in \pi_1$ with $P(\omega) > 0$ and $X(\omega) \neq v$. Among ω in π_1 for which $P(\omega) > 0$, let ω' be the one with the largest corresponding signal of player 2. Then $X(\omega') > v$ and, by construction and monotonicity, all other states in $\Pi_2(\omega')$ that occur with positive probability under P have associated values of X greater than or equal to v, contradicting requirement 2 of Definition 2. ■

With three or more players, it is no longer true that any security with increasing payoffs is separable. Indeed, consider the following example.

Example 2 *There are three players, 1, 2, and 3. Each player's signal is equal to -1, 0, or 1. If the sum of signals is less than 0, then the value of security X is -1. If the sum is greater than 0, the value is 1. If the sum of signals is equal to 0, then the value of security X is equal to:*

- -1 *if the vector of signals is* $(-1,0,1)$*,* $(0,1,-1)$*, or* $(1,-1,0)$*;*

- 1 *if the vector of signals is* $(1,0,-1)$*,* $(0,-1,1)$*, or* $(-1,1,0)$*; and*

- 0 *if the vector of signals is* $(0,0,0)$*.*

It is straightforward to check that value $v = 0$ and prior probability P that places probability $\frac{1}{6}$ on every permutation of $(-1,0,1)$ and 0 on all other states satisfy the requirements for the non-separability of security X.

5 Concluding Remarks

This paper leaves several important questions for future research. The first one is under what conditions, perfect Bayesian equilibria exist in the games studied in the paper. A discrete version of the discounted MSR game, in which players are only allowed to pick predictions from a finite set of values, is continuous at infinity, and therefore has a perfect Bayesian equilibrium (Fudenberg and Levine, 1983). In the current paper's model, however, action spaces are continuous, and so equilibrium existence is an open question, for both discounted and undiscounted MSR games. For the case of game Γ^K in which the value of the security is drawn from a normal distribution and the signals are distributed symmetrically and normally, the techniques of Foster and Viswanathan (1996) can be used to prove the existence of a linear equilibrium, but for the finite-state case considered in the current paper, equilibrium existence is also an open question.

Another important question is what happens when the traded security is non-separable and the traders' common prior is generic. For instance, suppose the security and the partition structure are as in Example 1, but the prior is a small generic perturbation of the one in the example. Then if the players simply announced their posterior beliefs truthfully, as in Geanakoplos and Polemarchakis (1982), information would get aggregated. What happens in the strategic trading game? Does there exist an equilibrium in which information gets aggregated with probability 1? Is there an equilibrium in which with positive probability information does not get aggregated, and instead in the limit, players get "stuck" at (or converge to) a prediction and a profile of beliefs under which none of them can make a profitable revision? Are the answers the same for all non-separable securities, scoring rules, and other parameters of the game? Using the techniques of the current paper, one can show that in general, prices have to converge to a random variable that is a "common knowledge/common belief" equilibrium of the corresponding economy (Dutta and Morris, 1997; DeMarzo and Skiadas, 1998, 1999), but it is unclear to which of the multiple equilibria they will converge.

There are also several interesting questions that go beyond the current paper's model. First, the intuition behind the main result of the paper suggests that it is very general and should continue to hold in many other market microstructure models, with liquidity-driven or noise traders, strategic market makers, and so on. Nevertheless, the details of the trading process may turn out to matter for the results, and so it is important to consider formally other dynamic microstructure models and to check in which of them similar conclusions hold.

Second, the model in the current paper assumes that traders are risk neutral. What happens if they are not? What happens if their utility functions are different?

Finally, this paper assumes that traders already possess the information, and the only concern is whether this information will get aggregated in the market. But markets are also often viewed as an incentive mechanism for traders to gather costly information, not just as an aggregation mechanism. What happens when traders can acquire information at a cost? How well is information extracted and aggregated in that case? Are some mechanisms better than others?

Appendix: Proofs

Proof of the First Statement of Theorem 1 for Game Γ^{MSR}

Let $w = 1, \ldots, |\Omega|$ index the states in Ω. Let $r = (r^1, r^2, \ldots, r^{|\Omega|})$ be a probability distribution over the states (with $r^w \equiv r(\omega_w)$ denoting the probability of state ω_w) and let z be any real number. Define *instant opportunity* of player i as the highest expected payoff the player can receive from making only one change to the forecast if the state is drawn according to distribution r and the current prediction is z. Formally, the instant opportunity of player i given r and z is equal to

$$\sum_{\omega \in \Omega} r(\omega) \left(s(E_r[X|\Pi_i(\omega)], X(\omega)) - s(z, X(\omega)) \right).$$

We first make an auxiliary observation. Let Δ be the set of probability distributions r such that there are states ω_a and ω_b with $r(\omega_a) > 0$, $r(\omega_b) > 0$, and $X(\omega_a) \neq X(\omega_b)$.

Lemma 1 *If security X is separable, then for all $r \in \Delta$ there exist $\phi > 0$ and $i \in \{1, 2, \ldots, n\}$ such that for any $z \in [\underline{y}, \overline{y}]$, the instant opportunity of player i given r and z is greater than ϕ.*

Proof. Consider separable security X and $r \in \Delta$ such that for any trader i and any $\phi > 0$, there exists $z \in [\underline{y}, \overline{y}]$ such that the instant opportunity of i given r and z is less than ϕ. By continuity of score function s, for every trader i there exists $z_i \in [\underline{y}, \overline{y}]$ such that the instant opportunity of i given r and z_i is equal to zero. Since s is strictly proper, this implies that for any ω with $r(\omega) > 0$, $E_r[X|\Pi_i(\omega)] = z_i$, which in turn implies that every z_i is equal to $E_r[X]$. This contradicts the assumption that security X is separable. ■

Now, let $q_0^w = P(\omega_w)$, i.e., the prior probability of state ω_w. Take a perfect Bayesian equilibrium of game Γ and consider the following stochastic process Y in $\mathbb{R}^{|\Omega|+1}$. Y_0 is deterministic and is equal to $(y_0, q_0^1, q_0^2, \ldots, q_0^{|\Omega|})$. Then nature draws state ω at random, according to distribution P, and each player i observes $\Pi_i(\omega)$. After that, player 1 plays according to his equilibrium strategy and makes forecast y_1. Based on this forecast y_1, the equilibrium strategy of player 1, and the prior P, a Bayesian outside observer, who shares prior P with the traders and observes all forecasts y_k but does not directly observe any information about the realized state ω, can form posterior beliefs about the probability of each state ω_w. Denote this probability by q_1^w. Y_1 is equal to $(y_1, q_1^1, q_1^2, \ldots, q_1^{|\Omega|})$. The rest of the process is constructed analogously: $Y_k = (y_k, q_k^1, q_k^2, \ldots, q_k^{|\Omega|})$, where y_k is the forecast made at time t_k and q_k^w is the posterior belief of the Bayesian outside observer about the probability of state ω_w, given his prior P, equilibrium strategies of players, and their history of forecasts up to and including time t_k.

Of course, $Q = \{(q_k^1, \ldots, q_k^{|\Omega|})\}_{k=0,1,\ldots}$ is also a stochastic process in $\mathbb{R}^{|\Omega|}$. The key idea of the proof is that this process is a martingale, by the law of iterated expectations. And by the martingale convergence theorem, it has to converge to a random variable, $Q_\infty = (q_\infty^1, \ldots, q_\infty^{|\Omega|})$.[1]

Suppose the statement of Theorem 1 does not hold for this equilibrium. Consider the limit random variable Q_∞ and two possible cases.

Case 1

Suppose there is a positive probability that Q_∞ assigns positive likelihoods to two states ω_a and ω_b with $X(\omega_a) \neq X(\omega_b)$. This implies that there is a vector of posterior probabilities $r = (r^1, \ldots, r^{|\Omega|})$ such that $r^a > 0$, $r^b > 0$, and for any $\epsilon > 0$, the probability that Q_∞ is in the ϵ-neighborhood of r is positive. Since Q_k converges to Q_∞, for any $\epsilon > 0$, there exists K and $\zeta > 0$ such that for any $k > K$, the probability that Q_k is in the ϵ-neighborhood of r is greater than ζ.

Now, by Lemma 1, for some player i and $\phi > 0$, the instant opportunity of player i is greater than ϕ given r and any $z \in [\underline{y}, \overline{y}]$. By continuity, this implies that for some $\epsilon > 0$, the instant opportunity of player i is greater than ϕ for any $z \in [\underline{y}, \overline{y}]$ and any vector of probabilities r' in the ϵ-neighborhood of r.

Therefore, for some i, t_K, and $\eta > 0$, the expected (over all realizations of stochastic process Q) instant opportunity of player i at any time $t_{n\kappa+i} > t_K$ is greater than η.

Case 2

Now suppose there is zero probability that Q_∞ assigns positive likelihoods to two states ω_a and ω_b with $X(\omega_a) \neq X(\omega_b)$. In other words, in the limit, the outside observer always believes that with probability 1, the value of the security is equal to some x^*, and places zero likelihood on all other possible values. Suppose the true state is ω, which has a positive prior probability of occurring. Then by Bayes' rule, it can only happen with zero probability that the outside observer's posterior beliefs Q_∞ place zero likelihood on the value of the security being equal to $X(\omega)$. (To see this, let H be the set of histories $(y_1, y_2, \ldots)$ after which the outside observer places zero likelihood on the true state being ω. Then $Prob(\omega|H) = 0$. But then $Prob(H|\omega) = Prob(\omega|H)Prob(H)/Prob(\omega) = 0$.) Hence, for every realization ω of nature's draw, with probability 1, Q_∞ will place likelihood 1 on the value of the security being equal to $X(\omega)$, i.e., in the limit, the outside observer's belief about the value of the security converges to its true value (even though his belief about the state of the world itself does not have to converge to the truth, if there are multiple states in which the security has the same value).

[1] Since the process is bounded, and thus uniformly integrable, convergence is both with probability 1 and in L^1. See, e.g., Øksendal (1995, Appendix C) for details.

Suppose now that process y_k does not converge in probability to the true value of the security. That is, there exist state ω and numbers $\epsilon > 0$ and $\delta > 0$ such that after state ω is drawn by nature, for any K, there exists $k > K$ such that $Prob(|y_k - X(\omega)| > \epsilon) > \delta$. This, together with the fact that even for the uninformed outsider the belief about the value of the security converges to the correct one with probability 1, implies that for some player i and $\eta > 0$, for any K, there exists time $t_{n\kappa+i} > t_K$ at which the expected instant opportunity of player i is greater than η.

Crucially, in both Case 1 and Case 2, there exist player i^* and value $\eta^* > 0$ such that there is an infinite number of times $t_{n\kappa+i^*}$ in which the expected instant opportunity of player i^* is greater than η^*. Fix i^* and η^*.

Now, let S_k be the expected score of prediction y_k (where the expectation is over all draws of nature and moves by players). The expected payoff to the player who moves in period t_k (it is always the same player) from the forecast revision made in that period is $\beta^k(S_k - S_{k-1})$.

The rest of the proof is split into two parts: $\beta < 1$ and $\beta = 1$.

Part "$\beta < 1$"

Take any period t_k. Let Ψ_k be the sum of all players' expected payoffs from the revisions made in periods t_k and later, divided by β^k: $\Psi_k = (S_k - S_{k-1}) + \beta(S_{k+1} - S_k) + \beta^2(S_{k+2} - S_{k+1}) + \dots$. We can make two observations about Ψ_k. First, it is non-negative, because each player can guarantee himself a payoff of zero. Second, for a similar reason, it is greater than or equal to the expected instant opportunity of the player who makes the forecast at time t_k.

Consider now $\lim_{K\to\infty} \sum_{k=1}^{K} \Psi_k$. On the one hand, under both Case 1 and Case 2, this limit has to be infinite, because each term Ψ_k is non-negative, and an infinite number of them are greater than η^*. On the other hand, for any K, $\sum_{k=1}^{K} \Psi_k =$

$$
\begin{array}{lccccccc}
 & (S_1 - S_0) & + & \beta(S_2 - S_1) & + & \beta^2(S_3 - S_2) & + & \dots \\
+ & & & & & & & \\
 & (S_2 - S_1) & + & \beta(S_3 - S_2) & + & \beta^2(S_4 - S_3) & + & \dots \\
+ & & & & & & & \\
 & \vdots & & & & & & \\
+ & & & & & & & \\
 & (S_K - S_{K-1}) & + & \beta(S_{K+1} - S_K) & + & \beta^2(S_{K+2} - S_{K+1}) & + & \dots
\end{array}
$$

$= \sum_{k=0}^{\infty} \beta^k (S_{k+K} - S_k) \le \frac{2M}{1-\beta}$, where $M = \max_{\{y\in[\underline{y},\overline{y}],\, \omega\in\Omega\}} |s(y, X(\omega))|$. Hence, both Cases 1 and 2 are impossible, and so y_k converges to the true value of security X.

Part "$\beta = 1$"

Take any player i. His expected payoff is equal to $\sum_{j=1}^{\infty}(S_{i+nj} - S_{i+nj-1})$. In equilibrium, the players' expected payoffs exist and are finite, so the infinite sum has to converge. Therefore, for any $\epsilon > 0$, there exists J such that for any $j > J$, $|\sum_{j'=j}^{\infty}(S_{i+nj'} - S_{i+nj'-1})| < \epsilon$. But in both Case 1 and Case 2, that contradicts the assumption that players are profit-maximizing after any history. To see that, it is enough to consider player i^* and some period t_{nj+i^*} such that the expected instant opportunity of i^* is greater than η^* and $|\sum_{j'=j}^{\infty}(S_{i^*+nj'} - S_{i^*+nj'-1})|$ is less than η^*.

Proof of the First Statement of Theorem 1 for Game Γ^K

Step 0. Consider a perfect Bayesian equilibrium of Γ^K.

Step 1. Take any $k \geq 0$ and let Ψ_k denote the unconditional expected total payoff of noise traders arriving after period t_k, multiplied by $(-\sqrt{2}^k)$, i.e.,

$$\Psi_k = (-\sqrt{2}^k) E\left[\sum_{k'=k+1}^{\infty} (x^* - y_{k'}) u_{k'}\right].$$

Since each strategic trader's expected continuation payoff after any period t_k is non-negative (because a strategic trader can always guarantee himself a payoff of zero by simply not trading), and the expected continuation payoff of market makers is zero (by construction), the expected payoff of noise traders arriving after period t_k cannot be positive, and so $\Psi_k \geq 0$. Moreover, for any $k' > k$, since y_k is independent of $u_{k'}$ and $E[y_k] = E[x^*]$, we have $E[(x^* - y_k)u_{k'}] = 0$ and so

$$\begin{aligned}
\Psi_k &= (-\sqrt{2}^k) E\left[\sum_{k'=k+1}^{\infty} (x^* - y_{k'}) u_{k'}\right] \\
&= (-\sqrt{2}^k) E\left[\sum_{k'=k+1}^{\infty} (x^* - y_k) u_{k'}\right] + (-\sqrt{2}^k) E\left[\sum_{k'=k+1}^{\infty} (y_k - y_{k'}) u_{k'}\right] \\
&= (-\sqrt{2}^k) E\left[\sum_{k'=k+1}^{\infty} (y_k - y_{k'}) u_{k'}\right] \\
&= (-\sqrt{2}^k) \sum_{k'=k+1}^{\infty} E\left[(y_k - y_{k'}) u_{k'}\right] \\
&\leq \sqrt{2}^k \sum_{k'=k+1}^{\infty} \sqrt{E\left[(y_k - y_{k'})^2\right] E\left[(u_{k'})^2\right]},
\end{aligned}$$

by the Cauchy-Schwarz inequality. Since process y_k is a uniformly bounded martingale, by the martingale convergence theorem, for any $\epsilon > 0$ there exists K such that for any $k > K$ and $k' > k$, $E[(y_k - y_{k'})^2] < \epsilon^2$, and so

$$\begin{aligned}
\Psi_k &\leq \sqrt{2}^k \sum_{k'=k+1}^{\infty} \sqrt{E\left[(y_k - y_{k'})^2\right] E\left[(u_{k'})^2\right]} \\
&\leq \sqrt{2}^k \sum_{k'=k+1}^{\infty} \epsilon \sqrt{Var(u_{k'})} \\
&= \sqrt{2}^k \epsilon \sum_{k'=k+1}^{\infty} \frac{1}{\sqrt{2}^{k'}} \\
&= \epsilon(1 + \sqrt{2}).
\end{aligned}$$

Therefore, Ψ_k converges to zero as t_k goes to 1. Note that for any strategic trader, the unconditional expected continuation payoff after period t_k is at most $\Psi_k/\sqrt{2}^k$, because in expectation, in the continuation game, noise traders lose $\Psi_k/\sqrt{2}^k$, market makers break even, and other strategic players do not lose money.

Step 2. Let Q_k be the stochastic process in $\mathbb{R}^{|\Omega|}$ denoting the posterior belief of an uninformed outside observer (or, in this case, a competitive market maker) about the true state of the world. Note that (i) Q is a uniformly bounded martingale and (ii) by construction, for each $k \geq 1$, y_k is equal to the expected value of X under Q_k. Take the limit random variable Q_∞. If with probability 1, it places all weight on states in which the value of the security is the same, then by the same argument as in the proof of the theorem for game Γ^{MSR}, this value has to be the true one with probability 1, and we are done.

Step 3. Suppose instead that there is a positive probability that Q_∞ places positive weights on two states in which the value of security X is different. Then there exist states ω and $\tilde{\omega}$ and distribution r over states in Ω such that $X(\omega) \neq X(\tilde{\omega})$, $r(\omega) > 0$, $r(\tilde{\omega}) > 0$, and for any $\epsilon > 0$ there exist $\delta > 0$ and K such that for any $k > K$, the probability that Q_k is in the ϵ-neighborhood of r is greater than δ. (An ϵ-neighborhood of r is the set of probability distributions r' on Ω such that for every state $\omega \in \Omega$, $|r(\omega) - r'(\omega)| \leq \epsilon$.) Fix distribution r for the rest of the proof and let $x_r = E_r[X]$.

Step 4. Since security X is separable, there exist trader i and elements π_a and π_b of his partition such that $r(\pi_a) > 0$, $r(\pi_b) > 0$, and $x_a \equiv E_r[X|\pi_a] < x_r < x_b \equiv E_r[X|\pi_b]$. Let $\tau = \min(x_b - x_r, x_r - x_a)/3$, $\rho = \min(r(\pi_a), r(\pi_b))/2$, and let $\phi > 0$ be such that for any r' in the ϕ-neighborhood of r, differences $|x_r - E_{r'}[X]|$, $|x_a - E_{r'}[X|\pi_a]|$, and $|x_b - E_{r'}[X|\pi_b]|$ are all less than τ and probabilities $r'(\pi_a)$ and $r'(\pi_b)$ are both greater than ρ (such ϕ exists by continuity). By the choice of r in Step 3, there exist $\delta_1 > 0$ and K_1 such that for any $k > K_1$, the probability that Q_k is in the $\frac{\phi}{2}$-neighborhood of r is greater than δ_1. Fix ϕ and δ_1 for the rest of the proof.

Step 5. Since Q_k converges to Q_∞, for any $\epsilon > 0$ there exists $K_2(\epsilon) \geq K_1$ such that for any $k > K_2$, the probability that $\max_\omega(|Q_k(\omega) - Q_{k+1}(\omega)|) > \frac{\phi}{2}$ is less than ϵ. Let $K_3(\epsilon) = K_2(\epsilon\delta_1\rho)$ (which is greater than or equal to $K_2(\epsilon)$). By construction, for any $k > K_3(\epsilon)$:

1. the probability that Q_k is in the $\frac{\phi}{2}$-neighborhood of r is at least δ_1;
2. conditional on point 1, the probability that player i's information is π_a is at least ρ;
3. conditional on points 1 and 2, the probability that Q_{k+1} is in the ϕ-neighborhood of r is at least $(1 - \epsilon)$, because the unconditional probability of Q_{k+1} being outside that neighborhood is at most $\epsilon\delta_1\rho$ and the probability of points 1 and 2 is at least $\delta_1\rho$.

Step 6. Note that when Q_{k+1} is in the ϕ-neighborhood of r, player i makes, in expectation, at least $(x_r - \tau) - (x_a + \tau) \geq \tau$ per unit of security purchased at time t_{k+1} (or loses at least that much per sold unit, if $d^i_{k+1} < 0$). By the arguments above, the total expected continuation payoff of player i after period t_k is bounded by $\Psi_k/\sqrt{2}^k$, and Ψ_k converges to 0. These observations imply that for any $\lambda > 0$, conditional on player i's information being π_a and Q_k being in the $\frac{\phi}{2}$-neighborhood of r, the probability that player i buys more than $\lambda/\sqrt{2}^k$ units of the security at time t_{k+1} must converge to 0 as t_k goes to 1. (Otherwise, consider the following "alternative" continuation strategy for player i: in periods t_{k+2} and later; or in period t_{k+1} when his information is not π_a or Q_k is not in the $\frac{\phi}{2}$-neighborhood of r; or in period t_{k+1} when the original strategy called for selling units or for buying less than $\lambda/\sqrt{2}^k$—in all of these circumstances, do not trade; otherwise, follow the original strategy. For a sufficiently small ϵ, this "alternative" strategy gives player i an expected continuation payoff close to or greater than $\tau\lambda\delta_1\rho/\sqrt{2}^k$, which is greater than $\Psi_k/\sqrt{2}^k$ for any sufficiently large k.) Similarly, the probability that player i sells more than $\lambda/\sqrt{2}^k$ units of the security at time t_{k+1}, conditional on his information being π_a and Q_k being in the $\frac{\phi}{2}$-neighborhood of r, must also converge to 0 as t_k goes to 1. (Otherwise, for a sufficiently large k, the expected losses incurred by player i in period t_{k+1} following that set of histories cannot be offset by the possible gains in periods t_{k+2} and later, and so he would be better off not trading at all, in any period $t_{k'} > t_k$, following that set of histories.)

Step 7.

Lemma 2 *For any $\mu > 0$ there exists $\nu > 0$ such that for any $z \in [-1, 1]$ and any set $S \subset \mathbb{R}$ whose probability under $N(0, \sigma^2)$ is less than or equal to ν, the probability of set $(S - \frac{z}{\sigma})$ under $N(0, \sigma^2)$ is less than or equal to μ.*

Proof. By rescaling, it is enough to prove the lemma for $\sigma = 1$. Let $\Phi(\cdot)$ denote the cdf of $N(0, 1)$ and, without loss of generality, assume that $\mu \leq 1$. Let u be the solution of equation $\Phi(u-1) = 1 - \frac{\mu}{4}$. Let $\nu_0 = \frac{\mu}{2}e^{-\frac{u^2}{2}}$. For the rest of the proof, "the probability of a set" refers to its probability under $N(0, 1)$ and, slightly abusing notation, is denoted "Φ."

Take any set S whose probability is less than or equal to ν_0. Take any $z \in [-1, 1]$. Then $\Phi(S - z) = \Phi((S - z) \cap [-u + 1, u - 1]) + \Phi((S - z) \cap ((-\infty, -u + 1) \cup (u - 1, \infty)))$.

By the choice of u, $\Phi((S-z) \cap ((-\infty, -u+1) \cup (u-1, \infty))) \leq \Phi((-\infty, -u+1) \cup (u-1, \infty)) = \frac{\mu}{2}$.

Also, $\Phi(S \cap [-u, u]) \leq \Phi(S) \leq \nu_0$. The density of the normal distribution is *at least* $\frac{1}{\sqrt{2\pi}}e^{-\frac{u^2}{2}}$ everywhere on $[-u, u]$, and so the Lebesgue measure of $S \cap [-u, u]$ is *at most* $\frac{\nu_0}{\frac{1}{\sqrt{2\pi}}e^{-\frac{u^2}{2}}}$. Hence, the Lebesgue measure of $((S - z) \cap [-u + 1, u - 1]) \subset ((S - z) \cap [-u - z, u - z]) = ((S \cap [-u, u]) - z)$ is also at most $\frac{\nu_0}{\frac{1}{\sqrt{2\pi}}e^{-\frac{u^2}{2}}}$. But then, since the density of the normal distribution is at most $\frac{1}{\sqrt{2\pi}}$ everywhere, the probability of $(S - z) \cap [-u + 1, u - 1]$ under the normal distribution is at most

$$\frac{1}{\sqrt{2\pi}} \frac{\nu_0}{\frac{1}{\sqrt{2\pi}}e^{-\frac{u^2}{2}}} = \frac{\nu_0}{e^{-\frac{u^2}{2}}} = \frac{\mu}{2}.$$

Therefore, $\Phi(S-z) = \Phi((S-z)\cap[-u+1,u-1]) + \Phi((S-z)\cap((-\infty,-u+1)\cup(u-1,\infty))) \leq \frac{\mu}{2}+\frac{\mu}{2}=\mu$. ■

Step 8. Let $\lambda = \frac{1}{4}$. Take a small $\epsilon > 0$ and a large k such that (i) the probability that player i buys or sells more than $\lambda/\sqrt{2}^{k+1}$ units of the security at time t_{k+1}, conditional on his information being π_a and Q_k being in the $\frac{\phi}{2}$-neighborhood of r, is less than ϵ and (ii) in the equilibrium under consideration, the expected continuation payoff of player i after period t_k is less than $\epsilon/\sqrt{2}^k$. Consider the following "alternative" continuation strategy for player i: in periods t_{k+2} and later, or in period t_{k+1} when his information is not π_a or Q_k is not in the $\frac{\phi}{2}$-neighborhood of r, do not trade; otherwise, buy $\frac{1}{2\sqrt{2}^{k+1}}$ units of the security. Note that when player i buys $\frac{1}{2\sqrt{2}^{k+1}}$ units of the security under this alternative strategy, with probability close to one, he was buying or selling at most $\frac{1}{4\sqrt{2}^{k+1}}$ units of the security under the original strategy, and so in that case, the total amount demanded by the strategic traders in period t_{k+1} changes by less than $\frac{1}{\sqrt{2}^{k+1}}$. Since the demand from noise traders in period t_{k+1} is distributed as $N(0, \frac{1}{2^{k+1}})$ and since under the original strategy of player i, market makers' belief Q_{k+1} was in the ϕ-neighborhood of r with probability at least $1-\epsilon$ (conditional on player i's information being π_a and Q_k being in the $\frac{\phi}{2}$-neighborhood of r), by Lemma 2, under the alternative strategy of player i, the probability that Q_{k+1} is in the ϕ-neighborhood of r is also close to 1, for a sufficiently small ϵ. But then the expected continuation payoff of player i from this alternative strategy is close to or greater than $\tau\lambda\delta_1\rho/\sqrt{2}^k$, which is greater than $\epsilon/\sqrt{2}^k$, contradicting the assumption that the original strategy was optimal for player i in the continuation game following every history.

Proof of Theorem 2: "Only If" Direction[2]

Suppose security X is separable. Take any $v \in \mathbb{R}$. Ignore all states ω with $X(\omega) = v$ and let $w = 1,\dots,W$ index the remaining states. Let $m = 1,\dots,M$ index all elements π of all players' partitions. Construct an $M \times W$ matrix A as follows. If state ω_w is in subset π_m, then the element in row m and column w of the matrix is equal to $X(\omega_w) - v$. Otherwise, it is equal to zero.

By Gordan's Transposition Theorem,[3] exactly one of the following two systems of equations and inequalities has a solution:

1. $Ax = 0$, $x \geq 0$, $x \neq 0$ (where $x \in \mathbb{R}^W$);

2. $A^T\lambda > 0$ (where $\lambda \in \mathbb{R}^M$).

Note that if system (1) has a solution, then the security is non-separable (Take solution x of (1); rescale it so that its elements sum to 1; and use rescaled probabilities as the common prior P.) Hence, if security X is separable, system (1) does not have a solution, which in turn implies that system (2) does, which is exactly the statement of the "only if" direction of Theorem 2.

[2]I am grateful to Yury Makarychev for this proof.

[3]http://eom.springer.de/m/m130240.htm

References

[1] Aumann, Robert J. (1976). "Agreeing to Disagree," *Annals of Statistics*, 4(6), pp. 1236–1239.

[2] Back, Kerry, C. Henry Cao, and Gregory A. Willard (2000). "Imperfect Competition among Informed Traders," *Journal of Finance*, 55(5), pp. 2117–2155.

[3] Brier, Glenn W. (1950). "Verification of Forecasts Expressed in Terms of Probability," *Monthly Weather Review*, 78(1), pp. 1–3.

[4] Chen, Yiling, and David M. Pennock (2007). "A Utility Framework for Bounded-Loss Market Makers," *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pp. 49–56.

[5] Chen, Yiling, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen (2007). "Bluffing and Strategic Reticence in Prediction Markets," *Proceedings of the 3rd International Workshop on Internet and Network Economics*, pp. 70–81.

[6] DeMarzo, Peter, and Costis Skiadas (1998). "Aggregation, Determinacy, and Informational Efficiency for a Class of Economies with Asymmetric Information," *Journal of Economic Theory*, 80(1), pp. 123–152.

[7] DeMarzo, Peter, and Costis Skiadas (1999). "On the Uniqueness of Fully Informative Rational Expectations Equilibria," *Economic Theory*, 13(1), pp. 1–24.

[8] Dimitrov, Stanko, and Rahul Sami (2008). "Non-myopic Strategies in Prediction Markets," *Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 200–209.

[9] Dubey, Pradeep, John Geanakoplos, and Martin Shubik (1987). "The Revelation of Information in Strategic Market Games," *Journal of Mathematical Economics*, 16(2), pp. 105–137.

[10] Dutta, Jayasri, and Stephen Morris (1997). "The Revelation of Information and Self-Fulfilling Beliefs," *Journal of Economic Theory*, 73(1), pp. 231–244.

[11] Foster, F. Douglas, and S. Viswanathan (1996). "Strategic Trading When Agents Forecast the Forecasts of Others," *Journal of Finance*, 51(4), pp. 1437–1478.

[12] Fudenberg, Drew, and David Levine (1983). "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory*, 31(2), pp. 251–268.

[13] Geanakoplos, John D., and Heraklis M. Polemarchakis (1982). "We Can't Disagree Forever," *Journal of Economic Theory*, 28(1), pp. 192–200.

[14] Golosov, Michael, Guido Lorenzoni, and Aleh Tsyvinski (2008). "Decentralized Trading with Private Information," working paper.

[15] Good, I. J. (1952). "Rational Decisions," *Journal of the Royal Statistical Society, Series B*, 14(1), pp. 107–114.

[16] Grossman, Sanford (1976). "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information," *Journal of Finance*, 31(2), pp. 573–585.

[17] Hanson, Robin (2003). "Combinatorial Information Market Design," *Information Systems Frontiers*, 5(1), pp. 107–119.

[18] Hayek, Friedrich A. (1945). "The Use of Knowledge in Society," *American Economic Review*, 35(4), pp. 519–530.

[19] Hellwig, Martin F. (1982). "Rational Expectations Equilibrium with Conditioning on Past Prices: A Mean-Variance Example," *Journal of Economic Theory*, 26(2), pp. 279–312.

[20] Kremer, Ilan (2002). "Information Aggregation in Common Value Auctions," *Econometrica*, 70(4), pp. 1675–1682.

[21] Kyle, Albert S. (1985). "Continuous Auctions and Insider Trading," *Econometrica*, 53(6), pp. 1315–1336.

[22] Laffont, Jean-Jacques, and Eric S. Maskin (1990). "The Efficient Market Hypothesis and Insider Trading on the Stock Market," *Journal of Political Economy*, 98(1), pp. 70–93.

[23] McKelvey, Richard D., and Talbot Page (1990). "Public and Private Information: An Experimental Study of Information Pooling," *Econometrica*, 58(6), pp. 1321–1339.

[24] Milgrom, Paul R. (1981). "Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica*, 49(4), pp. 921–943.

[25] Milgrom, Paul, and Nancy Stokey (1982). "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26(1), pp. 17–27.

[26] Nielsen, Lars T., Adam Brandenburger, John Geanakoplos, Richard McKelvey, and Talbot Page (1990). "Common Knowledge of an Aggregate of Expectations," *Econometrica*, 58(5), pp. 1235–1239.

[27] Øksendal, Bernt (1995). *Stochastic Differential Equations: An Introduction with Applications,* 4th ed., Springer.

[28] Pennock, David M. (2006). "Implementing Hanson's Market Maker," http://blog.oddhead.com/2006/10/30/, accessed December 25, 2008.

[29] Pesendorfer, Wolfgang, and Jeroen M. Swinkels (1997). "The Loser's Curse and Information Aggregation in Common Value Auctions," *Econometrica*, 65(6), pp. 1247–1282.

[30] Radner, Roy (1979). "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica*, 47(3), pp. 655–678.

[31] Reny, Philip J., and Motty Perry (2006). "Toward a Strategic Foundation for Rational Expectations Equilibrium," *Econometrica*, 74(5), pp. 1231–1269.

[32] Sebenius, James K., and John Geanakoplos (1983). "Don't Bet on It: Contingent Agreements with Asymmetric Information," *Journal of the American Statistical Association*, 78(382), pp. 424–426.

[33] Wilson, Robert (1977). "A Bidding Model of Perfect Competition," *Review of Economic Studies*, 44(3), pp. 511–518.

[34] Wolinsky, Asher (1990). "Information Revelation in a Market with Pairwise Meetings," *Econometrica*, 58(1), pp. 1–23.