

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<http://www.econometricsociety.org/>

Econometrica, Vol. 80, No. 6 (November, 2012), 2595–2647

INFORMATION AGGREGATION IN DYNAMIC MARKETS WITH STRATEGIC TRADERS

MICHAEL OSTROVSKY

Graduate School of Business, Stanford University, Stanford, CA 94305, U.S.A.

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

INFORMATION AGGREGATION IN DYNAMIC MARKETS WITH STRATEGIC TRADERS

BY MICHAEL OSTROVSKY¹

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 1 in one state of the world and 0 in all others, and “additive” securities, whose value can be interpreted as the sum of traders’ signals.

KEYWORDS: Information aggregation, efficient market hypothesis, rational expectations equilibrium, arbitrage, prediction markets.

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?

¹I am grateful to Jeremy Bulow, Yeon-Koo Che, Peter DeMarzo, Yossi Feinberg, Emir Kamenica, Yair Livne, George Mailath, Konstantin Makarychev, Yury Makarychev, Stephen Morris, Sergey Norin, Marek Pycia, Philip Reny, Rahul Sami, Michael Schwarz, Costis Skiadas, Robert Wilson, Bumin Yenmez, and the co-editor and anonymous referees for helpful comments, suggestions, and discussions. I am grateful to the Hoover Institution for its hospitality during the year when the first draft of this paper was completed, and to the NSF for financial support.

1.1. *Main Result*

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 nondegenerate 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 showed 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) introduced the concept of Rational Expectations Equilibrium (REE) and showed 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 explored the question of convergence to REE in various dynamic processes (see, e.g., Hellwig (1982), and Dubey, Geanakoplos, and Shubik (1987), for models of centralized trading; Wolinsky (1990), and Golosov, Lorenzoni, and Tsyvinski (2011), for models of decentralized trading; and McKelvey and Page (1986), and Nielsen, Brandenburger, Geanakoplos, McKelvey, and Page (1990), for models that extend the basic communication process of Geanakoplos and Polemarchakis (1982) to more complex settings in which agents’ beliefs are iteratively updated in response to repeated public observations of summary statistics of their actions). In all of these papers, however, it is assumed that each trader 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), Reny and Perry (2006)). Information aggregation results in these papers, however, rely on the assumption that the market is large, that is, the number of bidders goes to infinity and individual traders become small relative to the market. They also rely on various symmetry assumptions. No such assumptions are made in the current paper, and the number of traders is finite and fixed.

Kyle (1985) offered a model of dynamic insider trading, in which the single informed trader takes into account the nonnegligible 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) extended 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") equilibria. Laffont and Maskin (1990) criticized this reliance of the results of Kyle (1985) on linear trading strategies; argued that such models inherently have multiple equilibria; presented 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 concluded 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 equilibrium. In the case of a single informed trader, as in Laffont and Maskin (1990), every security is separable, and so information always gets aggregated. The contrasting conclusions of Laffont and Maskin are thus driven by the specifics of their model, not by the fact that they explicitly consider multiple equilibria. 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.

1.2. Market Scoring Rule

The framework of Kyle (1985) is the basis for the main model of trading in the current paper. However, in that framework, the question of information *aggregation*—that is, of multiple partially informed traders learning from each other and pooling their information over time—is intertwined with the question of information *revelation*—that is, of an informed trader taking advantage of his information and eventually moving the price of the security to its correct value. Even for the case of only one informed trader, when the issue of information aggregation does not arise, it is far from obvious what equilibria look

like and whether the informed trader's information will be revealed by the end of trading (and these questions are the focus of the Kyle (1985) paper and its criticisms in, e.g., Laffont and Maskin (1990)). Thus, the proof of the main result needs to address both of these issues, obscuring the intuition behind the information aggregation part.

To be able to illustrate that intuition more transparently, I also consider an auxiliary model of dynamic trading, in which the issue of information revelation does not arise: in the single-trader case, by construction, revelation is straightforward and immediate. This model is based on the market scoring rule (MSR) of Hanson (2003, 2007). 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 the 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 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 at time t_k is multiplied by β^k for some $\beta < 1$.

While my primary reason for studying this model is to illustrate the intuition behind information aggregation in the main model and thus make that result more transparent, information aggregation in MSR-based models is also of independent interest, for several reasons. First, such models can be viewed as generalizations of the communication processes of Geanakoplos and Polemarchakis (1982) and other papers in this tradition, in which several differentially informed agents sequentially announce their beliefs about the value of a random variable (or the probability of an event), and those beliefs eventually

converge to a common posterior. In those papers, it is assumed that the agents make truthful announcements, and strategic issues are ignored. Discounted MSR includes this truthful process as a special case, $\beta = 0$ (strictly speaking, the case $\beta = 0$ is ruled out in this paper, but it is easy to show that as β becomes small, in any equilibrium, players will behave almost myopically, i.e., will reveal their expectations almost truthfully), and at the same time makes it possible to examine the role of strategic behavior (for $\beta > 0$). The results of this paper show that, for separable securities, information aggregation does not depend on whether agents behave strategically or myopically.

Second, the MSR model includes as a special case a basic model of trading with an automated inventory-based market maker who offers to buy or sell shares in the security at price p that is a function of the (possibly negative) net inventory the market maker holds at that moment. Specifically, suppose the market maker starts with zero net inventory, sets the price for the security as a continuous decreasing function $p(z)$, where z is the total amount of shares he holds in his inventory (i.e., the more he holds, the less he is willing to pay for additional shares), and commits to buying or selling shares according to that price schedule. Thus, if his current inventory is z_0 , and a trader decides to sell $(z_1 - z_0)$ units of the security to the market maker, the market maker will pay that trader $\int_{z_0}^{z_1} p(\tilde{z}) d\tilde{z}$ for the $(z_1 - z_0)$ units. The current price of the security will move from $p(z_0)$ to $p(z_1)$. If the true value of the security then turns out to be equal to x , then the trader's payoff from this transaction will be equal to $\int_{z_0}^{z_1} p(\tilde{z}) d\tilde{z} - x(z_1 - z_0) = \int_{z_0}^{z_1} (p(\tilde{z}) - x) d\tilde{z}$. Thus, it is strictly optimal (myopically) for the trader to pick z_1 in such a way that $p(z_1)$ is equal to his belief about the value of the security (assuming the image of function $p(\cdot)$ includes that value). His payoff from this transaction is equal to his payoff from moving the forecast from y_0 to y_1 in the MSR model with the strictly proper scoring rule $s(y, x) = \int_0^{z_y} (p(\tilde{z}) - x) d\tilde{z}$, where $z_y = p^{-1}(y)$, that is, $p(z_y) = y$.²

Finally, note that while the market maker in this setting expects to lose money, the worst possible loss is bounded and can be controlled by adjusting the parameters of the rule. Another attractive feature of MSR in practice (relative to, say, continuous double auctions) is that a player can instantaneously make his prediction/trade at any time, without having to wait for another player who is willing to take the other side of the trade or to submit a limit order and wait for it to be filled. These features make MSR attractive for use in internal corporate prediction markets, and it is in fact used for that purpose: companies like Consensus Point and InKling Markets operate MSR-based prediction

²The correspondence between trading shares and eliciting beliefs from a single agent by the means of scoring rules was first noted by Savage (1971), who also provided additional technical details. Hanson (2003), Pennock (2006), and Chen and Pennock (2007) discussed this correspondence for the case of MSR. The study of automated market makers goes back to Black (1971a, 1971b), while formal analysis of inventory-based market makers goes back to Amihud and Mendelson (1980).

markets for Ford, Chevron, Best Buy, General Electric, and many other large corporations.³ Thus, the question of whether information in MSR-based prediction markets gets aggregated has direct practical implications.

Two recent papers have studied the equilibrium behavior of traders in MSR games.⁴ Chen, Reeves, Pennock, Hanson, Fortnow, and Gonen (2007) considered 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 found 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, that is, to make the prediction equal to his posterior belief. They also provided an example of a market in which signals are not conditionally independent and one of the traders has an incentive to behave non-myopically. Dimitrov and Sami (2008) also considered 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 found 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 introduced a two-player, two-signal MSR game with discounting, and proved that in that game, information gets aggregated in the limit, under the additional assumption that the "complementarity bound" of the security is positive. They reported that, based on their sample configurations, the bound is not always zero, but did 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.

1.3. *Paper Structure and Additional Results*

The remainder of this paper is organized as follows. Section 2 describes the model of information in the market, two models of trading (the main model based on the framework of Kyle (1985) and the auxiliary model based on the Market Scoring Rule), and the definitions of information aggregation and separability.

³Hanson (2009), Consensus Point (2011a, 2011b), Inking Markets (2011a, 2011b).

⁴These papers have been subsequently combined in Chen, Dimitrov, Sami, Reeves, Pennock, Hanson, Fortnow, and Gonen (2010).

Section 3 presents the main result.

Section 4 explores a discretized version of the main model. In the main model, traders' action spaces are continuous and unbounded, and it is not clear whether an equilibrium always exists in that environment. The results of Section 4.1 show that the main model can be discretized in a natural way, so that an equilibrium is guaranteed to exist, and in any equilibrium, information gets aggregated.

Section 5 explores the robustness of the main result to another variation. In the main model, even though trading happens in a finite interval, there are infinitely many trading periods whose size shrinks to zero over time. In the model of Section 5, there is a finite number of equally long trading periods, and thus instead of asking whether information gets aggregated as time approaches the end of the trading interval, the question is whether, for a sequence of games with an increasing number of trading periods, the price of a separable security in the last period becomes arbitrarily close, in expectation, to the true value of the security. The answer to that question is "yes" if there is one strategic trader; if there are several strategic traders, the answer is "yes" under an additional condition on equilibria.

Section 6 discusses the separability assumption and presents two important classes of separable securities: securities that can be represented as order statistics of traders' signals (including Arrow–Debreu securities) and those that can be represented as monotone transformations of linear functions of traders' signals. Most of the results in this section follow from the results of DeMarzo and Skiadas (1998, 1999), who explored a more general "separably oriented" condition in a richer environment, but I also include short self-contained proofs for the current paper's setting.

Section 7 concludes.

2. SETUP

There are n players, $i = 1, \dots, n$. There is a finite set of states of the world, Ω , and a random variable ("security") $X : \Omega \rightarrow \mathbb{R}$. 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 Π_1, \dots, Π_n consists of singleton sets; that is, for any two states $\omega_1 \neq \omega_2$, there exists player i such that $\Pi_i(\omega_1) \neq \Pi_i(\omega_2)$. $\Pi = (\Pi_1, \dots, \Pi_n)$ is the *partition structure*. Players have a common prior distribution P over states in Ω .

2.1. Trading: Main Model

In the main model, based on Kyle (1985), 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 $t_1 - t_0 = \frac{1}{2}$.⁵ Competitive market makers observe the aggregate demand $v_1 = \sum_i d_1^i + u_1$, 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.⁶ The market clears, and all traders observe price y_1 and aggregate demand v_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 $\Gamma^K(\Omega, \Pi, X, P)$.⁷

An *equilibrium* in game Γ^K is a profile of players' (possibly mixed) strategies S_i and the corresponding market makers' pricing rule Y such that

(i) price y_k set by the market makers in every period t_k is equal to the expected value of security X conditional on the observed aggregate demands in the market up to time t_k and profile S of players' strategies; and

(ii) for every player i , the expected payoff from following strategy S_i is at least as high as that from following any alternative strategy S'_i , given pricing rule Y and the profile of strategies of other players S_{-i} .

2.2. Trading: Auxiliary Model

In the auxiliary model, based on the market scoring rule of Hanson (2003, 2007), 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

⁵As in Kyle (1985), the idea behind this assumption is that during the trading period, $[0, 1]$, demand from noise traders arrives continuously according to a Brownian motion, and thus the demand that accumulates between times t and t' is distributed normally, with mean zero and variance proportional to $t' - t$. Assuming that this variance is in fact equal to $t' - t$ is just a normalization.

⁶More formally, pricing rule Y determines price y_1 as a function of v_1 , and an equilibrium condition requires this price to be equal to the expected value of security X .

⁷This model is mathematically isomorphic to a model of trading with an infinite number of periods, in which the demand from noise traders is distributed *identically* in every period, but also in every period, with probability $\rho = 1/\sqrt{2}$, the true value of the security is revealed and trading ends. To see this, one only needs to rescale the trades of all informed and noise traders by $\sqrt{2}^{k-1}$ in every period k ; then the expected payoff of any trader in the original game with the original strategies is equal to his expected payoff in the "reinterpreted" game with the "rescaled" strategies. This isomorphism also holds for the proofs of information aggregation results in these models; they would remain valid with virtually no changes for any $\rho \in (0, 1)$. Thus, these information aggregation results can also be interpreted as results on information aggregation in markets with a random deadline, which were studied by Back and Baruch (2004) and Caldentey and Stacchetti (2010) in the case of a single informed trader.

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 $t_\infty \equiv \lim_{k \rightarrow \infty} t_k = 1$, 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 , that is, each y_k is a number in an interval $[\underline{y}, \bar{y}]$, where $\underline{y} \leq \min_{\omega \in \Omega} X(\omega) \leq \max_{\omega \in \Omega} X(\omega) \leq \bar{y}$.

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 in the current setting), the logarithmic scoring rule ($s(y, x^*) = (x^* - a) \ln(y - a) + (b - x^*) \ln(b - y)$, for some $a < \underline{y}$ and $b > \bar{y}$), due to [Good \(1952\)](#).

Under the basic MSR (introduced by [Hanson \(2003, 2007\)](#), 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 $y_k = y_{k-1}$, that is, 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 $\Gamma^{\text{MSR}}(\Omega, \Pi, X, P, y_0, \underline{y}, \bar{y}, s, \beta)$.

2.3. Information Aggregation

DEFINITION 1: Under a profile of players' strategies in game Γ^{MSR} , or under a profile of players' strategies and the corresponding market makers' pricing rule in game Γ^K , we say that *information gets aggregated* if sequence y_k converges in probability to random variable $X(\omega^*)$.

Since the set of possible states Ω is finite, this definition is equivalent to saying that, for any $\varepsilon > 0$ 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 $|y_k - X(\omega^*)| > \varepsilon$ is less than δ . (Note that in addition to exogenous events, such as the realizations of demand from noise traders in game Γ^K , this probability may depend on the choices of strategic players, who may use mixed strategies.)

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) studied competitive equilibria with information structures similar to that of Example 1 and showed 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) showed 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; that is, in their terminology, the function mapping traders' signals to fully informative equilibrium prices is "separably oriented." Adapted to the current paper's setup, this condition translates into the following definition of separability, which plays a key role in subsequent 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:

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

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

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 perfectly informed trader in the market, then every security is separable.⁸ Section 6 discusses the separability condition in more detail and describes some natural classes of securities satisfying this condition.

3. MAIN RESULT

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

THEOREM 1: Consider state space Ω , security X , and partition structure Π .

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

⁸By assumption, the join of strategic traders' information partitions consists of singleton sets, and so if there is only one strategic trader in the market, he has to be perfectly informed.

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 (i) and (ii) 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 Appendix A. *Q.E.D.*

The intuition behind the proof of the first statement of Theorem 1 for game Γ^{MSR} is as follows. Fix an equilibrium and 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 (and knows their strategies). 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.^{9,10}

⁹The proof only considers players' behavior on the equilibrium path, which is why the result can be stated for Nash equilibrium rather than the more restrictive solution concepts like perfect Bayesian equilibrium or sequential equilibrium (although, of course, it holds for those more restrictive concepts as well). In fact, the result would also hold for even weaker solution concepts, such as the self-confirming equilibrium, in which players may hold incorrect beliefs about the actions of others at information sets off the equilibrium path.

¹⁰Back and Baruch (2004) also used the martingale convergence theorem to pin down the limit beliefs of an outside observer in a strategic dynamic trading environment, although in their

In game Γ^K , the principle behind the result is the same, but the statement that the lack of information aggregation implies the existence of non-vanishing profitable arbitrage opportunities that some trader can actually take advantage of becomes more delicate and requires a more elaborate proof.

4. EQUILIBRIUM EXISTENCE AND INFORMATION AGGREGATION IN A MODEL WITH DISCRETE ACTION SPACES

While Theorem 1 ensures that information gets aggregated in every equilibrium, it does not guarantee that an equilibrium does in fact exist. Also, the standard equilibrium existence results do not apply to games Γ^K and Γ^{MSR} , for a number of reasons (in particular, action spaces in those games are infinite). This section presents a discretized version of game Γ^K and shows that, in this version, equilibrium is guaranteed to exist and information always gets aggregated.¹¹

4.1. Model

Game Γ is as follows. As before, there are n strategic traders; a finite set of states of the world Ω ; and a random variable $X : \Omega \rightarrow \mathbb{R}$. Each trader i receives information about the true state of the world, $\omega \in \Omega$, according to partition Π_i of Ω . The join of partitions Π_1, \dots, Π_n consists of singleton sets. $\Pi = (\Pi_1, \dots, \Pi_n)$ is the partition structure. Strategic traders have a common prior distribution P over states in Ω .

Trading is organized as follows. At time $t_0 = 0$, nature draws a state, ω^* , according to P , and all strategic traders i observe their information $\Pi_i(\omega^*)$. At time $t_1 > 0$, each strategic trader i chooses his demand d_1^i from $2M_1 + 1$ possible actions $\{-M_1d_1, -(M_1 - 1)d_1, \dots, M_1d_1\}$, where unit of discretization d_1 is a positive real number and bound M_1 is a positive integer. At the same time, there is demand u_1 from noise traders, drawn randomly from a “discretized” normal distribution with mean zero, “variance parameter” $\sigma_1^2 > 0$, and the same unit of discretization d_1 . Formally, random variable \tilde{u}_1 is drawn from $N(0, \sigma_1^2)$, and then u_1 is equal to \tilde{u}_1 rounded to the nearest multiple of d_1 (i.e.,

setting there is only one informed strategic trader and the underlying security has only two possible values. The theorem has also been used for this purpose in the literatures on information herding (Smith and Sørensen (2000)) and long-run reputation persistence (Cripps, Mailath, and Samuelson (2004, 2007)).

¹¹For game Γ^{MSR} , a discrete version of the discounted ($\beta < 1$) game, in which players are only allowed to pick predictions from a finite set of values in every period, is continuous at infinity, and therefore has a perfect Bayesian equilibrium (Fudenberg and Levine (1983)). The proof that information in this game gets approximately aggregated (i.e., that for a sufficiently fine grid, the expected limit of the difference between the price and the true value of security has to be small) is very similar to the proof of Theorem 1 for game Γ^{MSR} , and is therefore omitted. It is an open question whether equilibria are guaranteed to exist for the case $\beta = 1$.

if $\tilde{u}_1 \in [-0.5d_1, 0.5d_1)$, then $u_1 = 0$; if $\tilde{u}_1 \in [0.5d_1, 1.5d_1)$, then $u_1 = d_1$; and so on). Competitive market makers observe aggregate demand $v_1 = \sum_i d_1^i + u_1$, 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 demand v_1 .

At time $t_2 > t_1$, the next auction takes place, with each strategic trader i choosing demand d_2^i from a set of $2M_2 + 1$ possible actions, with the unit of discretization d_2 , and with demand from noise traders, u_2 , drawn randomly from the discretized normal distribution with mean zero, variance parameter σ_2^2 , and unit of discretization d_2 . Subsequently, auctions are held at times t_k , with demand from noise traders drawn from the discretized normal distribution with variance parameter σ_k^2 and unit of discretization d_k , either until the last auction, which takes place at time $t_K \leq 1$ (in the case of finite games), or for an infinite number of periods $t_k \leq 1$ (in the case of infinite games; in that case we say $K = \infty$). The value of the security, $x^* = X(\omega^*)$, is revealed at some time $t^* > 1$. Trader i 's payoff is equal to $\sum_{k=1}^K d_k^i (x^* - y_k)$.

Strategy S_i of trader i is a set of functions $S_{ik}(\pi, d_1^i, \dots, d_{k-1}^i, v_1, \dots, v_{k-1})$ with values in the $2M_k$ -simplex, denoting the probability distribution over the $2M_k + 1$ possible actions of trader i in period k after he observed partition π , submitted demands d_1^i, \dots, d_{k-1}^i in the first $k - 1$ periods, and observed aggregate demands v_1, \dots, v_{k-1} in those periods.¹² Strategy profile S is a set of strategies of individual traders. Pricing rule Y is a set of functions $y_k(v_1, \dots, v_k)$, for all k .

As before, an *equilibrium* in game Γ is a profile of strategies S and the corresponding pricing rule Y such that

(i) for all k , price y_k is equal to the expected value of security X conditional on the observed aggregate demands in the market up to time t_k and profile S of players' strategies; and

(ii) for every player i , the expected payoff from following strategy S_i is at least as high as that from following any alternative strategy S'_i , given pricing rule Y and the strategies of other players S_{-i} .

4.2. Results

The first result shows that for any *finite* ($K < \infty$) game Γ , an equilibrium always exists. Note that this result does not immediately follow from standard

¹²These strategies do not formally depend on prices $y_{k'}$, because those prices are uniquely determined by prior aggregate demands $v_{k'}$, and thus including them as arguments of functions S_{ik} would be redundant. Note also that throughout the paper, it is assumed that aggregate demands $v_{k'}$ are observed by strategic traders. Under an alternative assumption that only prices $y_{k'}$ are observed by strategic traders but aggregate demands $v_{k'}$ are not observed, the results would be the same.

equilibrium existence results, because, first, “game” Γ is not a game in the strict game-theoretic sense (market makers are Bayesian but not strategic), and second, while the set of possible actions of each player is finite, the set of strategies is not: there are infinitely many possible information sets (due to infinitely many possible demands from noise traders). The proofs of all results in this section are in Appendix B.

THEOREM 2: *For any trading game Γ with $K < \infty$, an equilibrium exists.*

We now construct the following game Γ^∞ , with infinitely many trading periods, which is essentially a discretized version of the trading game defined in the main model (Section 2.1). Take any integer $M > 0$ and any positive real numbers d and σ . For every k , set $t_k = 1 - \frac{1}{2^k}$, $M_k = M$, $\sigma_k = \sigma \sqrt{t_k - t_{k-1}} = \frac{\sigma}{\sqrt{2^k}}$, and $d_k = \frac{d}{\sqrt{2^k}}$.

For this game, an equilibrium always exists, and in every equilibrium, if the traded security is separable, information always gets aggregated.

THEOREM 3: *For every M, d , and σ , there exists an equilibrium of game Γ^∞ .*

THEOREM 4: *Take any equilibrium of game Γ^∞ and consider the stochastic process y_k arising from that equilibrium. If security X is separable, then, for any $\varepsilon > 0$ and $\delta > 0$, there exists K' such that, for any $k > K'$ and any realization ω of the nature’s draw, the probability that $|y_k - X(\omega)| > \varepsilon$ is less than δ .*

The proof of Theorem 3 proceeds by constructing a sequence of “approximations” of an equilibrium, by considering a sequence of “truncated” finite games, for which equilibrium existence was established in Theorem 2. The proof of Theorem 4 is essentially the same as the proof of Theorem 1.

Note that while Theorems 3 and 4 would continue to hold for considerably more general models (e.g., it is not essential that the number of possible actions is the same in every period and for every strategic trader, or that time intervals decline as powers of 2), the proofs do rely on the assumption that the sizes of the largest possible trades by strategic traders are comparable to the standard deviations of noise traders’ demand. If the largest possible trades were much smaller than those standard deviations (more precisely, if as k increased, they went to zero faster than the standard deviations), then the information aggregation result might fail simply due to the fact that any moves by strategic players would be vanishingly small relative to those by noise traders, and thus would not noticeably impact prices. On the other hand, the proof of the equilibrium existence result relies on the fact that the size of the largest possible trade is limited, ensuring that game Γ^∞ is continuous at infinity, that is, that potential continuation profits after period t_k go to zero as k goes to infinity.

5. INFORMATION AGGREGATION IN A MODEL
WITH A FINITE NUMBER OF TRADING PERIODS

One feature that is different in the main model of this paper compared to earlier work (such as Kyle (1985) and Foster and Viswanathan (1996)) is the timing of trading. In this paper's main model, there are countably many trading periods, shrinking in length, and information aggregation occurs in a single game, along the sequence of trading periods. In the earlier work, there are finitely many trading periods of equal size, and results like information aggregation obtain only approximately for a given game—but the approximation becomes arbitrarily accurate if one considers a sequence of trading games, in which the number of trading periods increases and each trading period becomes small.¹³ It is thus a natural question to what extent this paper's results depend on having an infinite sequence of shrinking trading periods and whether they would continue to hold in a model similar to that in the earlier work, and I discuss this question below.

Consider the following model. There are n strategic traders and a finite set of states of the world Ω , with security X , partition structure Π , and common prior P , as before. Trading is organized as follows. There are K trading periods. At time $t_0 = 0$, nature draws a state, ω^* , according to P , and the strategic traders observe their partition elements. At time $t_1 = \frac{1}{K}$, each strategic trader i chooses his demand $d_1^i \in \mathbb{R}$. At the same time, there is demand u_1 from noise traders, drawn randomly from the normal distribution with mean zero and variance $\frac{\sigma^2}{K}$. Competitive market makers observe total demand $v_1 = \sum_i d_1^i + u_1$, 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. Trading takes place, and price y_1 and aggregate volume v_1 are publicly observed.

At time $t_2 = \frac{2}{K}$, the next auction takes place, with each strategic trader i choosing demand $d_2^i \in \mathbb{R}$, and with demand from noise traders, u_2 , again drawn randomly from the normal distribution with mean zero and variance $\frac{\sigma^2}{K}$. Subsequently, auctions are held at times $t_k = \frac{k}{K}$, $k \leq K$, with demand from noise traders drawn from the normal distribution with variance $\frac{\sigma^2}{K}$. The value of the security, $x^* = X(\omega^*)$, is revealed at some time $t^* > 1$. Trader i 's payoff is equal to $\sum_{k=1}^K d_k^i (x^* - y_k)$.

Strategy S_i of trader i is a set of functions S_{ik} denoting the probability distribution over the set of his possible actions (\mathbb{R}) in period k as a function of the element of his partition that he observed, his submitted demands

¹³See, for example, Theorem 4 in Kyle (1985) and numerical results in Section VI in Foster and Viswanathan (1996). Back (1992) and Back, Cao, and Willard (2000) provided analogous results in models with continuous trading, where information aggregation obtains exactly. However, some earlier papers (Back and Baruch (2004), Caldentey and Stacchetti (2010)) also considered models with, in essence, shrinking trading periods; see footnote 7 for discussion.

d_1^i, \dots, d_{k-1}^i in the first $k - 1$ periods, and the publicly observed aggregate demands v_1, \dots, v_{k-1} in those periods.

Pricing rule Y corresponding to strategy profile S is a set of functions $y_k(v_1, \dots, v_k)$ such that $y_k(v_1, \dots, v_k)$ is equal to the conditional expectation of the value of security X given that the true state was drawn according to prior probability P , strategic traders played according to strategy profile S , and the observed aggregate demands in the first k periods were v_1, \dots, v_k .

The resulting trading game is denoted Γ_K . As before, strategy profile S and pricing rule Y form an equilibrium of game Γ_K if Y is the pricing rule corresponding to S , and under S , each strategic trader is behaving optimally given pricing rule Y and the strategies of other players.

The results for this model are as follows. If there is only one strategic trader, then as the number of periods increases, the expected difference between the true value of the security, x^* , and the market price in the last period, y_K , always converges to zero.

THEOREM 5: *If $n = 1$, then for any $\varepsilon > 0$, there exists K_ε such that, for any $K > K_\varepsilon$, for any equilibrium (S^*, Y^*) of game Γ_K , $E[|y_K - x^*|] < \varepsilon$.*

See Appendix C for the proof.

For the case $n > 1$, an additional assumption is needed. Take any K , any equilibrium (S^*, Y^*) of game Γ_K , any $k < K$, and any player i . Take any history $h_{i,k}$ that player i could have observed up to period k (this includes his original signal, his actions, and the aggregate volumes of trade in periods before and including k). Define $\bar{x}(h_{i,k})$ as the expected value of security X conditional on history $h_{i,k}$. Define $\bar{y}(h_{i,k})$ as the average expected market price of the security in periods $k + 1$ and later, conditional on $h_{i,k}$ (i.e., $\bar{y}(h_{i,k}) = E[\frac{1}{K-k} \sum_{k'=k+1}^K y_{k'} | h_{i,k}]$). Finally, let $\bar{y}(0, h_{i,k})$ be the following “hypothetical” expectation: it is equal to the average expected market price in periods $k + 1$ and later, conditional on history $h_{i,k}$, if, following this history, player i completely stops trading and withdraws from the market, instead of following his prescribed equilibrium strategy (without other players or market makers knowing about this withdrawal). Take any $D \geq 1$. We say that equilibrium (S^*, Y^*) is D -destructive after history $h_{i,k}$ if $|\bar{x}(h_{i,k}) - \bar{y}(h_{i,k})| > D|\bar{x}(h_{i,k}) - \bar{y}(0, h_{i,k})|$, that is, if the active presence of trader i in the market keeps prices further away (by more than a factor of D) from the expected value of the security, compared to where these prices would have been if the trader had not traded at all.

Now consider a sequence of equilibria (S_m^*, Y_m^*) of games Γ_{K_m} (where $K_1 < K_2 < \dots$). We say that this sequence is *infinitely destructive* if, for some player i and some $\xi > 0$, for any $D \geq 1$, one can find index m and $k < K_m$ such that on the path of play, the measure of histories $h_{i,k}$ after which equilibrium (S_m^*, Y_m^*) is D -destructive is greater than ξ . In other words, no matter how large factor D is, in some game Γ_{K_m} , with a non-vanishing probability (greater than ξ), the actions of player i distort the average prices away from the correct expected value by more than a factor of D .

THEOREM 6: *Consider a sequence of equilibria (S_m^*, Y_m^*) of games Γ_{K_m} . If security X is separable and the sequence is not infinitely destructive, then, for any $\varepsilon > 0$, there exists m_ε such that, for any $m \geq m_\varepsilon$, in equilibrium (S_m^*, Y_m^*) , $E[|y_{K_m} - x^*|] < \varepsilon$.*

See the Supplemental Material (Ostrovsky (2012)) for the proof.

The proofs of Theorems 5 and 6 follow the same outline as the proof of Theorem 1 for the Kyle-based trading game. First, I establish an upper bound on the losses of noise traders: As time approaches the end of the trading interval, in the continuation game following that time these losses become arbitrarily small relative to the standard deviation of noise traders' demand over that time. The proof is complicated by the fact that instead of considering a single game and its equilibrium, I now need to consider a sequence of games and their equilibria, and thus statements like the one in the previous sentence require a more delicate formalization (Steps 0 and 1). Also, since there is no longer an infinite sequence of prices in any of these games, the martingale convergence theorem does not apply, and hence I show various convergence results directly, by decomposing the variance of $X(\omega)$ into a sum of variances of price changes, thus obtaining bounds on the latter. Those bounds are then used to provide a bound on the losses of noise traders (Step 2). While not important for the proof, it is worth noting that these bounds do not rely on the counterfactual assumption of information non-aggregation or, in the case of Theorem 6, on the assumption of the sequence of equilibria not being infinitely destructive; that is, they apply to any sequence of equilibria (S_m^*, Y_m^*) of games Γ_{K_m} .

Second, I establish a lower bound on the continuation profits of some strategic trader for an equilibrium in which information does not get aggregated. If information does not get aggregated and the traded security is separable, then at least one trader, in any arbitrarily late trading period, has a belief about the value of the security that is substantially different from the current (and expected future) prices. This trader, by trading an amount proportional to the standard deviation of noise traders' demand, would not move future prices too much, and thus could make continuation profits proportional to that standard deviation. Those profits provide a lower bound on his expected continuation profits in equilibrium, and for a sufficiently late trading period $t < 1$, exceed the expected losses of noise traders (which, by the first part of the proof, become arbitrarily small relative to the standard deviation of noise traders' demand). But this cannot be the case, because continuation profits of market makers are by construction zero, continuation profits of other strategic traders (in the case of Theorem 6) are nonnegative, and the sum of all continuation profits in the economy is zero.

The reason why the second part of the proof in this case is different from the corresponding part in the proof of Theorem 1 is that now the "counterfactual" trading by the arbitrageur has to take place over multiple periods, whereas in the case of Theorem 1 a one-shot deviation was sufficient (due to the shrinking nature of the sizes of time periods, each trading period was "large" relative

to the subsequent continuation game). The extra assumption of the sequence of equilibria not being infinitely destructive is used in showing that the arbitrage described above does in fact work. If a strategic trader's *withdrawal* from trading were to bring prices very close to his expected value of the security, then trading a fraction of a standard deviation of noise traders' demand would not necessarily generate profits proportional to that standard deviation. Intuitively, infinitely destructive sequences of equilibria appear to be quite unnatural, since prices in the model are set by Bayesian agents whom one would expect to become less rather than more accurate if the actual behavior of the players is different from the one these agents base their beliefs on. However, it is still an open question whether such sequences can in fact exist, and if they do, whether the result of Theorem 6 holds for them as well.¹⁴

In the case of Theorem 5, this additional assumption is not needed, because if prices do not converge to the correct value of the security (but are still unbiased estimates of that value conditional on publicly available information), then one can look at where the prices would be if the strategic trader stopped trading, and the answer would have to be substantially different from the true value of the security, with positive probability (see Step 3 and the first paragraph of Step 4 for details).

6. SEPARABLE SECURITIES

In light of the results on information aggregation, it is important to understand the restrictions the separability condition places on securities. This section describes two natural classes of separable securities and gives an alternative "dual" characterization of the condition. This dual characterization, as well as the result on one of the two classes (monotone transformations of additive payoffs) is due to DeMarzo and Skiadas (1998, 1999), who studied fully and partially informative rational expectation equilibria in rich settings with (potentially) multiple securities and infinite state spaces, and provided a variety of general results on the "separably oriented" condition, which is equivalent to the separability of security X in my setting. However, the results presented below also have short self-contained proofs for the current setting, which are included for completeness. The result on the other class of separable securities (order statistics) is new. All proofs are in Appendix D.

¹⁴One difficulty in showing that infinitely destructive sequences do not exist is the following. Suppose, for example, that there are two strategic traders in the market, and they are both perfectly informed. From their point of view, as the number of trading periods becomes large, they are facing a long repeated game. This repeated game may in principle have a large set of equilibria, including some convoluted ones in which the two traders collude to keep prices away from the correct value and make moderate profits, but if one of them deviates, for example, to non-trading, the other one quickly brings the price to the correct value. It is an open question whether such equilibria do in fact exist.

The first result of this section is a “dual” characterization of separability. While this characterization appears less intuitive than Definition 2, it is very convenient in applications, as Corollaries 1 and 2 below illustrate.

THEOREM 7: *Security X is separable under partition structure Π if and only if, for every $v \in \mathbb{R}$, there exist functions $\lambda_i: \Pi_i \rightarrow \mathbb{R}$ for $i = 1, \dots, n$ such that, for every state ω with $X(\omega) \neq v$,*

$$(X(\omega) - v) \sum_{i=1, \dots, n} \lambda_i(\Pi_i(\omega)) > 0.$$

In other words, the condition says that, for every $v \in \mathbb{R}$, we can pick multipliers $\lambda_i(\pi)$ for all elements π of partitions Π_i of all players i in such a way that, for all ω at which $X(\omega) \neq v$, the sign of $(X(\omega) - v)$ is the same as the sign of $\sum_i \lambda_i(\Pi_i(\omega))$.

The first corollary of Theorem 7 shows that securities that can be represented as order statistics of traders’ signals (minimum, maximum, median, etc.) are separable. Of course, there is no notion of a “signal” in the model of this paper; at least not one for which we can talk about relations like “lower,” “higher,” etc. Thus, the result applies to securities for which each observation by each player (i.e., each element of his information partition) can be meaningfully interpreted as a numerical signal. Note that these signals are allowed to be less informative than the original partitions; that is, two different elements of a partition are allowed to induce the same signal. The formal statement is as follows.

COROLLARY 1: *Consider any security X and suppose there exist functions $x_i: \Pi_i \rightarrow \mathbb{R}$ for $i = 1, \dots, n$ such that, for some $j \leq n$, for all $\omega \in \Omega$, $X(\omega)$ is equal to the j th lowest of the n numbers $x_1(\Pi_1(\omega)), x_2(\Pi_2(\omega)), \dots, x_n(\Pi_n(\omega))$. Then X is separable.*

Corollary 1 implies that any Arrow–Debreu security, that is, random variable X that is equal to 1 in one state of the world and to 0 in all other states, is separable, and hence by Theorem 1, information about Arrow–Debreu securities gets aggregated. To see that, consider security X that is equal to 1 in state ω^* and to 0 in all other states. For all i and $\pi \in \Pi_i$, let $x_i(\pi) = 1$ if $\pi = \Pi_i(\omega^*)$ and $x_i(\pi) = 0$ if $\pi \neq \Pi_i(\omega^*)$. Then $X(\omega) \equiv \min_i \{x_i(\Pi_i(\omega))\}$, and thus by Corollary 1, security X is separable.

The second corollary of Theorem 7 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 (where again a “signal” is a numerical interpretation of a trader’s information). Additive securities, of course, include a seemingly more

general case of securities that are linear (e.g., the average) or stochastically monotone (McKelvey and Page (1986), Bergin and Brandenburger (1990)) functions of signals, because signals can be rescaled.

COROLLARY 2: *Consider any security X and suppose there exist functions $x_i : \Pi_i \rightarrow \mathbb{R}$ for $i = 1, \dots, n$ and a monotone function $f : \mathbb{R} \rightarrow \mathbb{R}$ such that, for all $\omega \in \Omega$, $X(\omega) = f(\sum_i x_i(\Pi_i(\omega)))$. Then X is separable.*

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.

7. CONCLUDING REMARKS

This paper leaves several important questions for future research. One such 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. Note that in game Γ^{MSR} , for a generic prior, information in any *pure-strategy* equilibrium will get aggregated even for non-separable securities.^{15,16} It is not clear, however, whether this result can be extended in any form to game Γ^K or mixed-strategy equilibria of game Γ^{MSR} .

¹⁵I am grateful to an anonymous referee for this observation.

¹⁶To see this, note that, in any pure strategy equilibrium, after any history, the knowledge of every agent (players and outside observers) is simply a set of states ω that are consistent with the observed history; the posterior probabilities of these states are proportional to their probabilities under the original prior. The set of subsets of set Ω is finite, and thus for a generic prior, each subset induces a different expected value of security X . Also, in any pure-strategy equilibrium, all learning is guaranteed to stop before some time t . Consider the beliefs of players and the outside observer at time t . If the outside observer (and thus all strategic players) knows the true state of the world ω , by the same argument as in the proof of Theorem 1, forecast y_t has to be equal to the value of the security. If he does not know the true state of the world, his information is some set

Two other open questions in the context of the current paper's models have been raised earlier in the paper: the existence of an equilibrium for the game Γ^{MSR} without discounting (either in the original model or in a discretized version) and the necessity of the "not infinitely destructive" condition in Theorem 6.

There are also several interesting questions that go beyond the current paper's models. Note that while there are numerous mathematical differences between these models, the underlying principle behind the information aggregation result is the same in all of them: In a dynamic market, as time approaches the end of trading, the beliefs of an outside observer have to converge *somewhere*, and thus as we get to the end of trading, they become stable. If there is a positive probability that these limit beliefs place a positive likelihood on a state with a wrong value of the security, then the security must also be "mispriced" with positive probability in any sufficiently late period, and separability implies that at least one trader is aware of this mispricing. Thus, any trading mechanism that allows such a trader to make more money by taking advantage of this mispricing than by following his supposed equilibrium strategy has to lead to information aggregation in the limit.

This principle should continue to hold in many other market microstructure models: with strategic or automated market makers; with or without "external" noise; with one or many securities¹⁷; 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. Also, this paper focuses on information aggregation properties of dynamic markets and abstracts away from issues of allocative efficiency: in the models, all traders are risk-neutral, and have pure common values. A natural direction for future research is to consider settings with risk-averse traders and/or traders who have private or interdependent components in their valuations, and see how far the principle stated above can be pushed. A recent paper by [Iyer, Johari, and Moallemi \(2010\)](#) suggests that it can indeed be useful in the richer markets discussed above. [Iyer, Johari, and Moallemi \(2010\)](#) extended the techniques developed in the current paper to study a market with an automated market maker, risk-averse strategic traders, and multiple securities. They made a stronger assumption than separability on the information structure of the traders, but assumed that as a group,

$\Omega' \subset \Omega$, $|\Omega'| > 1$. There must be a strategic player whose information is finer than Ω' (because the join of informed traders' partitions consists of singletons), and the genericity assumption implies that, under different possible information realizations of this player, the expected values of the security must be different. Thus, he has a persistent arbitrage opportunity, which by the same argument as in the proof of Theorem 1 is impossible in equilibrium.

¹⁷Note that for settings with multiple securities, if the market is complete (i.e., if the set of traded securities is rich enough to span all Arrow–Debreu securities for states in Ω), the results in Section 6 suggest that information would get aggregated, although of course the precise statement would depend on the specific assumptions about market microstructure.

these traders are still not fully informed: there is some residual uncertainty even after pooling all information. In that setting, not only does information get aggregated (in the sense that all traders learn all available information), but also the portfolios of securities they end up with by the end of trading represent efficient risk sharing among them.

Finally, this paper assumes that traders already possess the information at the beginning of trading, 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. This view gives rise to more questions: 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 A: PROOF OF THEOREM 1

This appendix consists of three sections. The first section contains the main parts of the proof of Theorem 1 for game Γ^{MSR} . The second section formally presents some technical details of the setup of game Γ^{MSR} and of the proof of the information aggregation result for that game. The last section contains the proof of Theorem 1 for game Γ^K .

A.1. Proof of Theorem 1 for Game Γ^{MSR}

The proof of this theorem consists of four steps.

STEP 1: In this step, we state and prove an auxiliary lemma on the existence of a uniform lower bound on the expected profits at least one strategic trader can make from improving a forecast about the value of a separable security under uncertainty.

Let r be some probability distribution over the states in Ω and let z be any real number. Define *instant opportunity* of player i as the highest expected payoff that he 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 (of course, since s is a strictly proper scoring rule, player i would achieve this payoff by changing the forecast to the expected value of the security given his information). Formally, the instant opportunity of player i given r and z is equal to

$$(1) \quad \sum_{\omega \in \Omega} r(\omega) (s(E_r[X|\Pi_i](\omega)), X(\omega)) - s(z, X(\omega))).$$

Now, let Δ be the set of probability distributions r such that there are states a and b with $r(a) > 0$, $r(b) > 0$, and $X(a) \neq X(b)$. In other words, for distributions in Δ , there is some uncertainty about the value of security X .

LEMMA 1: *If security X is separable, then, for every $r \in \Delta$, there exist $\chi > 0$ and $i \in \{1, 2, \dots, n\}$ such that, for every $z \in \mathbb{R}$, the instant opportunity of player i given r and z is greater than χ .*

Before moving on to the proof of the lemma, note that the separability assumption is crucial: without it, the statement would be false. Also, the order of quantifiers in the statement of the lemma is important: in one of the subsequent steps of the proof of Theorem 1, it will allow us to place a lower bound on the expected continuation profits of a trader without having to worry about the current market forecast.

PROOF: Let $x_r = E_r[X]$ —the expected value of security X under r . Since X is separable and there is uncertainty about its value under r , by the definition of separability, for at least one player i and at least one state ω , $r(\omega) > 0$ and $E_r[X|\Pi_i(\omega)] \neq x_r$. Together with $x_r = E_r[X]$, this implies that there exist two states, a and b , such that $r(a) > 0$, $r(b) > 0$, and $E_r[X|\Pi_i(a)] < x_r < E_r[X|\Pi_i(b)]$.

Let $\pi_a = \Pi_i(a)$ —the element of partition Π_i that contains state a . Let $\pi_b = \Pi_i(b)$. Since s is a proper scoring rule, for any element π of partition Π_i , we have

$$(2) \quad \sum_{\omega \in \pi} r(\omega)(s(E_r[X|\pi], X(\omega)) - s(z, X(\omega))) \geq 0.$$

Let x_a denote $E_r[X|\pi_a]$, and let $x_b = E_r[X|\pi_b]$. Since we can rewrite “ $\sum_{\omega \in \Omega}$ ” in equation (1) as “ $\sum_{\pi \in \Pi_i} \sum_{\omega \in \pi}$,” equation (2) implies that the instant opportunity of player i is at least as large as

$$(3) \quad \max \left[\sum_{\omega \in \pi_a} r(\omega)(s(x_a, X(\omega)) - s(z, X(\omega))), \right. \\ \left. \sum_{\omega \in \pi_b} r(\omega)(s(x_b, X(\omega)) - s(z, X(\omega))) \right].$$

Now, a strictly proper scoring rule s has to be “order-sensitive,” that is, the further away the forecast is from the true expected value, the lower is the expectation of the score (Lambert (2011, Proposition 1)). Thus, if z is greater than or equal to x_r , and thus (weakly) further away from x_a , we have

$$\sum_{\omega \in \pi_a} r(\omega)(s(x_a, X(\omega)) - s(z, X(\omega))) \\ \geq \sum_{\omega \in \pi_a} r(\omega)(s(x_a, X(\omega)) - s(x_r, X(\omega))),$$

and if $z \leq x_r$, we have

$$\begin{aligned} & \sum_{\omega \in \pi_b} r(\omega)(s(x_b, X(\omega)) - s(z, X(\omega))) \\ & \geq \sum_{\omega \in \pi_b} r(\omega)(s(x_b, X(\omega)) - s(x_r, X(\omega))). \end{aligned}$$

These inequalities, in turn, imply that, for all z , the expression in equation (3) is greater than or equal to

$$(4) \quad \min \left[\sum_{\omega \in \pi_a} r(\omega)(s(x_a, X(\omega)) - s(x_r, X(\omega))), \sum_{\omega \in \pi_b} r(\omega)(s(x_b, X(\omega)) - s(x_r, X(\omega))) \right].$$

The expression in equation (4) does not depend on z , and is strictly greater than zero: scoring rule s is strictly proper, and $x_a < x_r < x_b$. This completes the proof of Lemma 1. *Q.E.D.*

STEP 2: In this step, we construct a stochastic process corresponding to the beliefs of an outside observer about the realized state ω and establish its martingale properties. (Some additional technical details, such as the formal description of the underlying probability space for this process, are presented in Section A.2.)

Let $h = 1, \dots, H$ index the states in Ω . Let $q_0^h = P(h)$, that is, the (common) prior probability of state h . Take a Nash equilibrium of game Γ^{MSR} and consider the following stochastic process Q in \mathbb{R}^H . Q_0 is deterministic and is equal to $(q_0^1, q_0^2, \dots, q_0^H)$. 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 (possibly mixed) 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 h . Denote this probability by q_1^h . Q_1 is then equal to $(q_1^1, q_1^2, \dots, q_1^H)$. The rest of the process is constructed analogously: $Q_k = (q_k^1, q_k^2, \dots, q_k^H)$, where y_k is the forecast made at time t_k and q_k^h is the posterior belief of the Bayesian outside observer about the probability of state h , given his prior P , equilibrium strategies of players, and their history of forecasts up to and including time t_k .

Note that process Q is a martingale, by the law of iterated expectations. It is also bounded, since all its realizations are between 0 and 1. Thus, by the

martingale convergence theorem, it has to converge to some random variable, $Q_\infty = (q_\infty^1, \dots, q_\infty^H)$.

STEP 3: In this step, we show that if the statement of Theorem 1 does not hold for this equilibrium, then we can identify a “non-vanishing arbitrage opportunity”: there is a player, i^* , and a positive number, η^* , such that the expected instant opportunity of player i^* exceeds η^* at infinitely many trading times t_k .

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

Step 3, Case 1: Suppose there is a positive probability that Q_∞ assigns positive likelihoods to two states, a and b , with $X(a) \neq X(b)$. Then there exists a probability distribution $r = (r^1, \dots, r^H)$ such that $r^a > 0$, $r^b > 0$, and for any $\varepsilon > 0$, the probability that Q_∞ is in the ε -neighborhood of r is positive.¹⁸ Since Q_k converges to Q_∞ , this implies that, for any $\varepsilon > 0$, there exist 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 $\chi > 0$, the instant opportunity of player i is greater than χ given r and any $z \in \mathbb{R}$. By continuity,¹⁹ this implies that, for some $\varepsilon > 0$, the instant opportunity of player i is greater than χ for any z and any vector of probabilities r' in the ε -neighborhood of r .

Therefore, for some i , $\chi > 0$, t_K , and $\zeta > 0$, the instant opportunity of player i at any time $t_{n_{k+i}} > t_K$ is greater than χ with probability at least ζ , and thus for i , t_K , and $\eta = \chi\zeta > 0$, the expected instant opportunity of player i at any time $t_{n_{k+i}} > t_K$ is greater than η .

Step 3, Case 2: Now suppose there is zero probability that Q_∞ assigns positive likelihoods to two states a and b with $X(a) \neq X(b)$, that is, in the limit, the outside observer believes with certainty that the value of the security is equal to some x . Almost surely (i.e., with probability 1), Q_∞ has to assign some positive likelihood to the true state h ,²⁰ and thus for every realization h of nature’s draw, with probability 1, Q_∞ will place likelihood 1 on the value of the security being equal to $X(h)$. In other words, 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).

¹⁸See Section A.2.3 for the proof of this statement.

¹⁹A proper scoring rule does not need to be a continuous function. However, $E_r[s(E_r[X], X)]$ has to be convex (and thus continuous) in r for any proper scoring rule s : if $r = \alpha r' + (1 - \alpha)r''$ for $\alpha \in (0, 1)$, then $E_r[s(E_r[X], X)] = \alpha E_{r'}[s(E_r[X], X)] + (1 - \alpha)E_{r''}[s(E_r[X], X)] \leq \alpha E_{r'}[s(E_{r'}[X], X)] + (1 - \alpha)E_{r''}[s(E_{r''}[X], X)]$ (this argument is due to Savage (1971)). This, in turn, implies that the expression in equation (1) is continuous in r .

²⁰See Section A.2.4 for the proof of this statement.

Suppose now that process y_k does not converge in probability to the true value of the security. That is, there exist state $h \in \Omega$ and numbers $\varepsilon > 0$ and $\delta > 0$ such that, after state h is drawn by nature, for any K , there exists $k > K$ such that the probability that $|y_k - X(h)| > \varepsilon$ is greater than δ . 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 $t_{nk+i} > t_K$ 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_{nk+i^*} in which the expected instant opportunity of player i^* is greater than η^* . Fix i^* and η^* .

STEP 4: This step concludes the proof, by showing that the presence of a “non-vanishing arbitrage opportunity” is impossible in equilibrium.

Let \bar{s}_k be the expected score of prediction y_k (where the expectation is over all draws of nature and realizations of 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 $\beta^k(\bar{s}_k - \bar{s}_{k-1})$.

The rest of the proof proceeds separately for cases $\beta = 1$ and $\beta < 1$.

Step 4, Case “ $\beta = 1$ ”: Take any player i . His expected payoff is equal to $\sum_{\kappa=1}^{\infty} (\bar{s}_{n\kappa+i} - \bar{s}_{n\kappa+i-1})$. In equilibrium, the players’ expected payoffs exist and are finite, so the infinite sum has to converge. Therefore, for any $\varepsilon > 0$, there exists K such that, for any $\kappa > K$, $|\sum_{\kappa'=\kappa}^{\infty} (\bar{s}_{n\kappa'+i} - \bar{s}_{n\kappa'+i-1})| < \varepsilon$. But in both Case 1 and Case 2 of Step 3, that contradicts the assumption that players are profit-maximizing after any history. To see that, it is enough to consider player i^* identified in Step 3 and some period $t_{n\kappa+i^*}$ such that the expected instant opportunity of i^* is greater than η^* and $|\sum_{\kappa'=\kappa}^{\infty} (\bar{s}_{n\kappa'+i^*} - \bar{s}_{n\kappa'+i^*-1})|$ is less than η^* .

Step 4, Case “ $\beta < 1$ ”: If we knew that expected scores \bar{s}_k converged to a limit as k increased, the proof of this case would be essentially identical to that of the previous one. However, we do not know a priori that these scores do in fact converge, and so the proof needs to be somewhat more involved.

Take any period t_k . Let Ψ_k be the sum of all players’ expected continuation payoffs from the forecast revisions made in periods t_k and later, divided by β^k : $\Psi_k = (\bar{s}_k - \bar{s}_{k-1}) + \beta(\bar{s}_{k+1} - \bar{s}_k) + \beta^2(\bar{s}_{k+2} - \bar{s}_{k+1}) + \dots$. We can make two observations about Ψ_k . First, it is nonnegative, 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 $\lim_{K \rightarrow \infty} \sum_{k=1}^K \Psi_k$. On one hand, under both Case 1 and Case 2 of Step 3, this limit has to be infinite, because each term Ψ_k is nonnegative, and

an infinite number of them are greater than η^* . On the other hand, for any K ,

$$\begin{aligned} \sum_{k=1}^K \Psi_k &= (\bar{s}_1 - \bar{s}_0) + \beta(\bar{s}_2 - \bar{s}_1) + \beta^2(\bar{s}_3 - \bar{s}_2) + \dots \\ &+ (\bar{s}_2 - \bar{s}_1) + \beta(\bar{s}_3 - \bar{s}_2) + \beta^2(\bar{s}_4 - \bar{s}_3) + \dots \\ &+ \vdots \\ &+ (\bar{s}_K - \bar{s}_{K-1}) + \beta(\bar{s}_{K+1} - \bar{s}_K) + \beta^2(\bar{s}_{K+2} - \bar{s}_{K+1}) + \dots \\ &= (\bar{s}_K - \bar{s}_0) + \beta(\bar{s}_{K+1} - \bar{s}_1) + \beta^2(\bar{s}_{K+2} - \bar{s}_2) + \dots \\ &\leq 2M/(1 - \beta), \end{aligned}$$

where $M = \max_{\{y \in [\underline{y}, \bar{y}], \omega \in \Omega\}} |s(y, X(\omega))|$. Hence, both Cases 1 and 2 are impossible, and so y_k must converge in probability to the true value of security X .

A.2. Additional Technical Details for Game Γ^{MSR}

This section contains some additional mathematical formalism for the setup of game Γ^{MSR} as well as formal proofs of three technical statements used in the proof of Theorem 1 for Γ^{MSR} .

A.2.1. Setup

The initial uncertainty about the value of the security and about the information possessed by the players is captured by state of nature ω . However, once the game starts, players can use mixed strategies, thus introducing additional uncertainty into the economy. We formalize this additional uncertainty as follows. Assume that in each period t_k , the player who moves in that period observes a realization of a random variable ν_k , drawn from the uniform distribution on $[0, 1]$. These draws are independent of each other and of the initial state ω . The “full” state of the world (including initial uncertainty and randomization by players in the trading game) is described by state $\varphi = (\omega, \nu_1, \nu_2, \dots)$. The space of possible states is thus the product space $\Phi = \Omega \times [0, 1]^{\mathbb{N}}$. We denote the corresponding product probability space by $(\Phi, \mathcal{F}, \bar{P})$.

For player i and time t_k such that $k = i + n\kappa$ for some κ (i.e., player i makes a forecast at time t_k), the *strategy of player i at time t_k* is a measurable function $\sigma_{i,k} : \Pi_i \times [\underline{y}, \bar{y}]^{k-1} \times [0, 1] \rightarrow [\underline{y}, \bar{y}]$, denoting the forecast y_k made by player i given the element of partition Π_i he observed prior to trading, the history of forecasts (y_1, \dots, y_{k-1}) up to time t_k , and the realization of the random variable ν_k . The *strategy of player i* , denoted σ_i , is a set of strategies of player i at all times when it is his turn to make a forecast.

Given a profile of strategies $\sigma = (\sigma_1, \dots, \sigma_n)$, any realization of uncertainty (i.e., state φ) determines a sequence of forecasts: $y_1(\sigma, \varphi), y_2(\sigma, \varphi), \dots$. For

$\beta < 1$ (discounted MSR), we say that profile σ is a *Nash equilibrium* of game Γ^{MSR} if, for every player i and every alternative strategy σ'_i (and the corresponding profile of strategies $\sigma' = (\sigma_1, \dots, \sigma'_i, \dots, \sigma_n)$), we have

$$(5) \quad \mathbb{E} \left[\sum_{\kappa=0}^{\infty} \beta^{i+n\kappa} (s(y_{i+n\kappa}(\sigma, \varphi), X(\varphi)) - s(y_{i+n\kappa-1}(\sigma, \varphi), X(\varphi))) \right] \\ \geq \mathbb{E} \left[\sum_{\kappa=0}^{\infty} \beta^{i+n\kappa} (s(y_{i+n\kappa}(\sigma', \varphi), X(\varphi)) - s(y_{i+n\kappa-1}(\sigma', \varphi), X(\varphi))) \right],$$

where expectations are taken with respect to probability measure \bar{P} , and $X(\varphi)$ is the value of the security when the “full” state is φ , determined by the first component of φ : $X(\varphi) = X(\omega(\varphi))$.

Since forecasts are limited to the interval $[\underline{y}, \bar{y}]$, the scores (and thus differences in scores) are uniformly bounded, and therefore the infinite sums and their expectations in (5) are guaranteed to be finite and well-defined when $\beta < 1$.

For the case of undiscounted MSR ($\beta = 1$), we say that profile σ is a *Nash equilibrium* of game Γ^{MSR} if, for every player i , the limit

$$(6) \quad \lim_{K \rightarrow \infty} \mathbb{E} \left[\sum_{\kappa=0}^K (s(y_{i+n\kappa}(\sigma, \varphi), X(\varphi)) - s(y_{i+n\kappa-1}(\sigma, \varphi), X(\varphi))) \right]$$

exists, and for every alternative strategy σ'_i (and the corresponding profile of strategies $\sigma' = (\sigma_1, \dots, \sigma'_i, \dots, \sigma_n)$) such that the limit

$$(7) \quad \lim_{K \rightarrow \infty} \mathbb{E} \left[\sum_{\kappa=0}^K (s(y_{i+n\kappa}(\sigma', \varphi), X(\varphi)) - s(y_{i+n\kappa-1}(\sigma', \varphi), X(\varphi))) \right]$$

exists, we have (6) \geq (7).

A.2.2. *Formal Statement of Step 2 of the Proof of Theorem 1 for Γ^{MSR}*

Fix equilibrium strategy profile σ . Index states in Ω by $h = 1, \dots, H$. Define H -dimensional stochastic process Q on $(\Phi, \mathcal{F}, \bar{P})$ as follows. Q_0 is deterministic: it is equal to $(q_0^1, q_0^2, \dots, q_0^H)$, where $q_0^h = P(h)$. For every $k \geq 1$ and realization of uncertainty φ^* , set $Q_k(\varphi^*) = (q_k^1(\varphi^*), \dots, q_k^H(\varphi^*))$, where

$$q_k^h(\varphi^*) = \bar{P}(\omega(\varphi) = h | y_1(\sigma, \varphi) = y_1(\sigma, \varphi^*), y_2(\sigma, \varphi) = y_2(\sigma, \varphi^*), \dots, \\ y_k(\sigma, \varphi) = y_k(\sigma, \varphi^*)).$$

By the law of iterated expectations, stochastic process Q is a martingale. It is also bounded, since all q_k^h are between 0 and 1. Thus, it is uniformly integrable,

and so by the martingale convergence theorem, there exists a random variable Q_∞ on $(\Phi, \mathcal{F}, \bar{P})$ such that $Q_\infty(\varphi) = \lim_{k \rightarrow \infty} Q_k(\varphi)$ almost everywhere (\bar{P}) and

$$\int_{\Phi} \|Q_k - Q_\infty\| d\bar{P} \rightarrow 0 \quad \text{as } k \rightarrow \infty.$$

(For concreteness, we will use the Euclidean distance $\|\cdot\|$, although many other distance functions would work just as well.)

A.2.3. *Additional Argument in Step 3.1*

Suppose that, for states a and b in Ω , the probability (\bar{P}) that $q_\infty^a(\varphi)$ and $q_\infty^b(\varphi)$ are both positive is greater than some $\delta > 0$. Then it has to be the case that, for some positive integer L , the probability that $q_\infty^a(\varphi)$ and $q_\infty^b(\varphi)$ are both greater than $\frac{1}{L}$ is greater than δ (because $\{\varphi | q_\infty^a(\varphi) > 0, q_\infty^b(\varphi) > 0\} = \bigcup_{L \in \mathbb{N}} \{\varphi | q_\infty^a(\varphi) > \frac{1}{L}, q_\infty^b(\varphi) > \frac{1}{L}\}$). Fix that L .

Consider the H -dimensional box $B_0 = [0, 1] \times [0, 1] \times \dots \times [\frac{1}{L}, 1] \times \dots \times [0, 1] \times \dots \times [\frac{1}{L}, 1] \times \dots \times [0, 1]$ (where intervals $[\frac{1}{L}, 1]$ are at dimensions a and b , and intervals $[0, 1]$ are everywhere else). By construction, the probability \bar{P} that Q_∞ is in that box is greater than δ . By taking midpoints of the H intervals, splitting each interval into two subintervals, and considering various combinations of these subintervals, we can represent the box as a union of 2^H equally sized closed boxes. Clearly, for at least one of these boxes, Q_∞ belongs to that box with probability at least $\frac{\delta}{2^H}$. Denote this box by B_1 .

Proceeding by analogy, we construct a sequence of boxes B_l such that the size of each box in the sequence is equal to half the size of the previous one in the sequence (along each dimension), and the probability that Q_∞ is in box B_l is greater than $\frac{\delta}{2^{lH}}$. Let $r = \bigcap_{l \in \mathbb{N}} B_l$ (this intersection has to be a singleton). For any $\varepsilon > 0$, there exists l large enough that the entire box B_l belongs to the ε -neighborhood of r , and thus with probability greater than $\frac{\delta}{2^{lH}} > 0$, vector Q_∞ belongs to that ε -neighborhood. Note that it has to be the case that r is a distribution, that is, the sum of its coordinates is equal to 1 (otherwise, for a sufficiently small ε , no point in the ε -neighborhood of r is a distribution, and thus that neighborhood cannot ever contain distribution Q_∞). Note also that by construction, it has to be the case that r^a and r^b are greater than or equal to $\frac{1}{L} > 0$.

A.2.4. *Additional Argument in Step 3.2*

Consider state $h \in \Omega$. Let $\Phi^* = \{\varphi | \omega(\varphi) = h, q_\infty^h(\varphi) = 0\}$ and $\Xi^* = \{\varphi | q_\infty^h(\varphi) = 0\}$. In words, Φ^* is the set of “full” states of the world φ such that the security is in state h ($\omega(\varphi) = h$), but the outside observer, in the limit, assigns likelihood zero to that event ($q_\infty^h(\varphi) = 0$). Ξ^* is the set of φ in which the outside observer, in the limit, assigns likelihood zero to state h . Clearly, $\Phi^* \subset \Xi^*$.

By Lévy’s zero-one law, $q_\infty^h(\varphi) = \bar{P}(\omega = h | y_1(\varphi), y_2(\varphi), \dots)$, that is, q_∞^h is the probability of being in state h conditional on the observed infinite sequence of forecasts y_1, y_2, \dots . Therefore, by the definition of conditional probability,

$$\bar{P}(\Phi^*) = \int_{\Xi^*} q_\infty^h(\varphi) d\bar{P}(\varphi).$$

Thus, $\bar{P}(\Phi^*) = 0$, since by construction, $q_\infty^h(\varphi) = 0$ for all $\varphi \in \Xi^*$.

A.3. Proof of Theorem 1 for Game Γ^K

STEP 0: Fix an equilibrium of game Γ^K .

STEP 1: In this step, we obtain a bound on the expected losses of noise traders as time approaches the end of the trading interval.

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 $(-\sqrt{2}^k)$, that is,

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

where the expectation is over the draws of state $\omega \in \Omega$, the realizations of noise traders’ demands, and (if the equilibrium is in mixed strategies) the randomizations of strategic traders’ actions (the formal construction of the probability space is analogous to that in Section A.2, and is therefore omitted).

Since each strategic trader’s expected continuation payoff after any period t_k is nonnegative (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 $\mathbb{E}[y_k] = \mathbb{E}[x^*]$, we have $\mathbb{E}[(x^* - y_k)u_{k'}] = 0$ and so

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

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

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

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

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 the 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 correct 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 exists distribution r over states in Ω such that r places positive weights on two states in which the value of security X is different, and for any $\varepsilon > 0$ 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 .²¹ Fix distribution r for the rest of the proof and let $x_r = E_r[X]$.

STEP 4: In this step, we begin to identify “mispricings” and show that they occur with a non-vanishing positive probability. This is the step in which we use the assumption that security X is separable.

Since security X is separable, there exist trader i and elements π_a and π_b of his information partition such that $r(\pi_a) > 0$, $r(\pi_b) > 0$, and $x_a \equiv E_r[X|\pi_a] < x_r < x_b \equiv E_r[X|\pi_b]$. Define $\tau = (x_b - x_r)/5$ and $\rho = r(\pi_b)/2$. Pick $\nu > 0$ such that, for any r' in the (2ν) -neighborhood of r , differences $|x_r - E_{r'}[X]|$ and $|x_b - E_{r'}[X|\pi_b]|$ are less than τ and probability $r'(\pi_b)$ is greater than ρ . By the choice of r in Step 3, there exist $\zeta > 0$ and K_1 such that, for any $k > K_1$, the probability that Q_k is in the ν -neighborhood of r is greater than ζ . (Note: this is not a typo; we want Q_k to be in the ν -neighborhood of r and then, conditional on that, will ensure that Q_{k+1} is in the (2ν) -neighborhood with probability close to 1.) Fix ν for the rest of the proof.

STEP 5: The remaining steps show how trader i can take advantage of the “mispricings” identified in the previous step.

From this point on, it is convenient to introduce more formal notation, along the lines of Section A.2 (fully formal definitions are similar to those in that section, and are omitted). As in that section, the value of the security is determined by state $\omega \in \Omega$, but the description of the “full” state of the world, $\varphi \in \Phi$, also includes more information. First, it includes the realized demands u_k from noise traders in each period k . Second, it includes randomizations by strategic traders in each period, which we assume are driven by independent draws ι_k^i (for trader i in period k) from the uniform distribution on $[0, 1]$. Denote the corresponding product probability space by $(\Phi, \mathcal{F}, \bar{P})$.

Let φ_k^i denote all information about the full state φ available to trader i in period k , and let $\hat{\varphi}_k^i$ denote that information and in addition the draw ι_{k+1}^i that drives the randomization of player i in period $k + 1$ (for the analysis of pure-strategy equilibria, this would be unnecessary, but it is needed for the analysis of mixed-strategy ones). In other words, $\hat{\varphi}_k^i$ uniquely determines the action of player i in period $k + 1$. Note that it also uniquely determines Q_k ; that is, before making his move in period $k + 1$, player i knows the beliefs of an outside observer after period k . Hence, we can talk about functions like $Q_k(\hat{\varphi}_k^i)$ without ambiguity. Also, let φ_k denote all public information up to period k ; likewise, we can talk about $Q_k(\varphi_k)$.

²¹For concreteness, in this proof we will use the L^∞ norm to define distances between distributions on Ω , and so the ε -neighborhood of r is the set of r' such that, for every $\omega \in \Omega$, $|r(\omega) - r'(\omega)| < \varepsilon$.

STEP 6: Let $\bar{x}(\varphi_k^i) = \bar{x}(\hat{\varphi}_k^i)$ denote the expected value of security X conditional on the information available to trader i up to period k . Consider all public histories φ_k such that $Q_k(\varphi_k)$ is in the ν -neighborhood of r (also, denote that neighborhood as $B_\nu(r)$ and the (2ν) -neighborhood of r as $B_{2\nu}(r)$). We know that for $k > K_1$, the probability of such a history is greater than ζ .

We also know (by construction in Step 4) that conditional on every such history, the probability that the true state ω is in π_b is greater than ρ . Thus, with probability greater than $\zeta\rho$, the history $\hat{\varphi}_k^i$ observed by player i up to period k is such that he knows that $\omega \in \pi_b$ and $Q_k(\hat{\varphi}_k^i) \in B_\nu(r)$. Let us call this set of histories \mathcal{A}_k .

Note that

$$(8) \quad E_{\bar{P}}[\bar{x}(\hat{\varphi}_k^i)|\hat{\varphi}_k^i \in \mathcal{A}_k] = E_{\bar{P}}[X|(\omega \in \pi_b) \text{ and } (Q_k(\varphi_k) \in B_\nu(r))] > x_b - \tau,$$

where the inequality follows by construction from Step 4: for every $r' \in B_\nu(r) \subset B_{2\nu}(r)$, we have $E_{r'}[X|\omega \in \pi_b] > x_b - \tau$, and thus $E_{\bar{P}}[X|(\omega \in \pi_b) \text{ and } (Q_k(\varphi_k) = r')] = E_{r'}[X|\omega \in \pi_b] > x_b - \tau$, which in turn implies $E_{\bar{P}}[X|(\omega \in \pi_b) \text{ and } (Q_k(\varphi_k) \in B_\nu(r))] > x_b - \tau$.

Let \mathcal{B}_k be the subset of \mathcal{A}_k consisting of histories $\hat{\varphi}_k^i$ such that $\bar{x}(\hat{\varphi}_k^i) \geq x_b - 2\tau$. Since security X is bounded, the inequality in (8) implies that, for some $\xi > 0$, for any sufficiently large k , $\bar{P}(\mathcal{B}_k|\mathcal{A}_k) > \xi$, and thus $\bar{P}(\mathcal{B}_k) > \zeta\rho\xi$.

STEP 7: In this step, we use the bound on the losses of noise traders (which is also a bound on the profits of any strategic trader) from Step 1 to show that in the states identified above, a strategic trader cannot be buying or selling too many units of the security, since otherwise either his immediate profits or his subsequent continuation profits would violate the bound.

By the arguments in Step 1, 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 as k increases. Let us show that this implies that, for any $\lambda > 0$, conditional on event \mathcal{B}_k , the probability that player i buys more than $\lambda/\sqrt{2^{k+1}}$ units of the security at time t_{k+1} (i.e., more than λ times the standard deviation of noise traders' demand in that period) must converge to 0 as t_k goes to 1.

Indeed, suppose it does not, and instead has some limit point $\alpha > 0$. Take a large k for which the probability of player i buying more than $\lambda/\sqrt{2^{k+1}}$ units conditional on \mathcal{B}_k exceeds $\alpha/2$. Let \mathcal{C}_k be the subset of histories in \mathcal{B}_k following which player i buys more than $\lambda/\sqrt{2^{k+1}}$ units. Since $\bar{P}(\mathcal{C}_k|\mathcal{B}_k) > \alpha/2$, we have $\bar{P}(\mathcal{C}_k) > \zeta\rho\xi\alpha/2$.

Let us place a lower bound on the *unconditional* expected continuation payoff of player i following period k .

Since sequence Q_k converges, for any $\varepsilon > 0$, for any sufficiently large k , the unconditional probability that $\|Q_{k+1} - Q_k\| \geq \nu$ is less than ε . Hence, the probability that $\hat{\varphi}_k^i \in C_k$ and then Q_{k+1} is inside the (2ν) -neighborhood of r ($B_{2\nu}(r)$) is at least $\zeta\rho\xi\alpha/2 - \varepsilon$, which, by choosing any sufficiently small ε (i.e., any sufficiently large k), can be made arbitrarily close to $\zeta\rho\xi\alpha/2$. Take any $\delta > 0$, and let $\mathcal{D}_k(\delta)$ denote the subset of C_k consisting of histories $\hat{\varphi}_k^i$ for which the probability of Q_{k+1} being in $B_{2\nu}(r)$, following $\hat{\varphi}_k^i$, is at least $1 - \delta$.

Note that for any δ , as k goes to infinity, conditional probability $\bar{P}(\mathcal{D}_k(\delta)|C_k)$ converges to 1, and thus for any δ , for any sufficiently large k , $\bar{P}(\mathcal{D}_k(\delta))$ is greater than or close to $\zeta\rho\xi\alpha/2$. Also, since the value of security X is bounded, we can pick such δ_1 that knowing that the probability of $Q_{k+1} \in B_{2\nu}(r)$ is greater than $1 - \delta_1$ guarantees that the expected price in period $k + 1$ is at most $x_r + 2\tau$ (since $Q_{k+1} \in B_{2\nu}(r)$ implies $y_{k+1} < x_r + \tau$). Combining all of the above, we find that for any sufficiently large k , following every history $\hat{\varphi}_k^i \in \mathcal{D}_k(\delta_1)$, the expected continuation profit of player i is at least $\lambda/\sqrt{2}^{k+1} ((x_b - 2\tau) - (x_r + 2\tau)) = \tau\lambda/\sqrt{2}^{k+1}$, and since following any history his expected continuation profit is nonnegative, his unconditional expected continuation profit following any sufficiently large period k is greater than or close to $(\zeta\rho\xi\alpha/2) \cdot (\tau\lambda/\sqrt{2}^{k+1})$. Since Ψ_k converges to 0 as k increases, this expression is greater than $\Psi_k/\sqrt{2}^k$ for any sufficiently large k , contradicting prior assumptions.

Showing that the probability that player i sells more than $\lambda/\sqrt{2}^{k+1}$ units of the security at time t_{k+1} must converge to 0 as t_k goes to 1 follows a related logic: If this probability does not converge to zero, then, for a sufficiently large k , the expected losses incurred by player i from these large sales are of the order $\Omega(1/\sqrt{2}^k)$ and cannot be offset by any subsequent continuation profits after period t_{k+2} , which by Step 1 must be of the order $o(1/\sqrt{2}^k)$. Thus, such a strategy would result in negative expected continuation payoffs, which cannot happen in equilibrium.

STEP 8: In this step, we prove an auxiliary lemma that implies, roughly, that by changing his demand by a number of shares (at most) proportional to the standard deviation of noise traders' demand, a strategic trader does not move the beliefs of market makers (and thus prices) too much.

LEMMA 2: *For any $A > 0$, $\gamma > 0$, $\sigma > 0$, there exists $\delta > 0$ such that, for any $z \in [-A, A]$ and any measurable set $S \subset \mathbb{R}$ whose probability under $N(0, \sigma^2)$ is less than or equal to δ , the probability of set $(S - \sigma z)$ under $N(0, \sigma^2)$ is less than or equal to γ , that is, knowing that the probability of set S is at most δ guarantees that the probability of the "shifted" set $(S - \sigma z)$ is at most γ .*

PROOF: By rescaling, it is enough to prove the lemma for $\sigma = 1$. Let $\Phi(\cdot)$ denote the c.d.f. of $N(0, 1)$ and, without loss of generality, assume that $\gamma \leq 1$.

Let u be the solution of the equation $\Phi(u - A) = 1 - \frac{\gamma}{4}$. Let $\delta_0 = \frac{\gamma}{2}e^{-u^2/2}$. 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 δ_0 . Take any $z \in [-A, A]$. Then

$$\begin{aligned} \Phi(S - z) &= \Phi((S - z) \cap [-u + A, u - A]) \\ &\quad + \Phi((S - z) \cap ((-\infty, -u + A) \cup (u - A, \infty))). \end{aligned}$$

By the choice of u , for the second term of this sum we have $\Phi((S - z) \cap ((-\infty, -u + A) \cup (u - A, \infty))) \leq \Phi((-\infty, -u + A) \cup (u - A, \infty)) = \frac{\gamma}{2}$.

Also, $\Phi(S \cap [-u, u]) \leq \Phi(S) \leq \delta_0$. The density of the normal distribution is at least $\frac{1}{\sqrt{2\pi}}e^{-u^2/2}$ everywhere on $[-u, u]$, and so the Lebesgue measure of $S \cap [-u, u]$ is at most $\frac{\delta_0}{(1/\sqrt{2\pi})e^{-u^2/2}}$. Hence, the Lebesgue measure of $((S - z) \cap [-u + A, u - A]) \subset ((S - z) \cap [-u - z, u - z]) = ((S \cap [-u, u]) - z)$ is also at most $\frac{\delta_0}{(1/\sqrt{2\pi})e^{-u^2/2}}$. But then, since the density of the normal distribution is at most $\frac{1}{\sqrt{2\pi}}$ everywhere, the probability of $(S - z) \cap [-u + A, u - A]$ under the normal distribution is at most

$$\frac{1}{\sqrt{2\pi}} \frac{\delta_0}{\frac{1}{\sqrt{2\pi}}e^{-u^2/2}} = \frac{\delta_0}{e^{-u^2/2}} = \frac{\gamma}{2}.$$

Therefore, $\Phi(S - z) = \Phi((S - z) \cap [-u + A, u - A]) + \Phi((S - z) \cap ((-\infty, -u + A) \cup (u - A, \infty))) \leq \frac{\gamma}{2} + \frac{\gamma}{2} = \gamma$, and we can set $\delta = \delta_0$ to finish the proof. *Q.E.D.*

STEP 9: This step concludes the proof by using Lemma 2 to show that, by trading an amount proportional to the standard deviation of noise traders’ demand, a strategic trader can take advantage of mispricings identified in earlier steps and make expected profits that exceed the bound established in Step 1.

Take any $\lambda > 0$. Take a small $\varepsilon > 0$ and a large k (how small ε should be and how large k should be will be clear later in the step).

Recall from Step 6 that \mathcal{A}_k is the set of histories $\hat{\varphi}_k^i$ such that $Q_k(\hat{\varphi}_k^i) \in B_\nu(r)$ and $\omega \in \pi_b$, and \mathcal{B}_k is the subset of \mathcal{A}_k consisting of histories $\hat{\varphi}_k^i$ such that $\bar{x}(\hat{\varphi}_k^i) \geq x_b - 2\tau$. Also recall that for any sufficiently large k , $\bar{P}(\mathcal{B}_k) > \zeta\rho\xi > 0$.

Let \mathcal{G}_k be the set of histories in \mathcal{B}_k following which player i buys or sells at most $\lambda/\sqrt{2}^{k+1}$ units of the security in period $k + 1$. By Step 7, as $k \rightarrow \infty$, $\bar{P}(\mathcal{G}_k)/\bar{P}(\mathcal{B}_k) \rightarrow 1$.

Take any $\delta > 0$ and let $\mathcal{H}_k(\delta)$ be the set of histories in \mathcal{G}_k following which the probability that Q_{k+1} is in $B_{2\nu}(r)$ is at least $1 - \delta$. By the same argument

as for sets $\mathcal{D}_k(\delta)$ and \mathcal{C}_k in Step 7, we know that, for any $\delta > 0$, as $k \rightarrow \infty$, $\overline{P}(\mathcal{H}_k(\delta))/\overline{P}(\mathcal{G}_k) \rightarrow 1$.

Fix some small $\delta > 0$ and a large k , and take some history $\hat{\varphi}_k^i \in \mathcal{H}_k(\delta)$. Suppose, following this history, in period $k + 1$, player i buys $\Delta/\sqrt{2}^{k+1}$ units of the security (recall that $\hat{\varphi}_k^i$ includes player i 's randomization in period $k + 1$, and thus uniquely determines his action in that period). Let F denote the distribution of the sum of other strategic traders' demands in period $k + 1$ following history $\hat{\varphi}_k^i$ (denote it by d_k^{-i}), and recall that the distribution of noise traders' demands (u_k) is independent of all prior history and is given by $N(0, 1/2^{k+1})$. Let \mathcal{Z} denote the realizations of d_k^{-i} after which the probability of Q_{k+1} being outside $B_{2\nu}(r)$ is greater than $\sqrt{\delta}$. Clearly, the probability of \mathcal{Z} under F is less than $\sqrt{\delta}$ (otherwise the probability of Q_{k+1} being outside of $B_{2\nu}(r)$ following $\hat{\varphi}_k^i$ would exceed $\sqrt{\delta} \cdot \sqrt{\delta} = \delta$). Following the remaining realizations, $\neg\mathcal{Z}$, with probability at least $1 - \sqrt{\delta}$, Q_{k+1} will be inside $B_{2\nu}(r)$ (and thus price y_{k+1} will be within τ of x_r).

Consider the following deviation by player i following history $\hat{\varphi}_k^i \in \mathcal{H}_k(\delta)$ chosen above: instead of buying $\Delta/\sqrt{2}^{k+1}$ units, buy $\lambda/\sqrt{2}^{k+1}$ units. Let us see how that impacts price y_{k+1} . With probability at most $\sqrt{\delta}$, the realization of d_k^{-i} is in \mathcal{Z} , in which case what we know about the price is that its absolute value is always at most $\max_{\omega} |X(\omega)|$. With probability at least $1 - \sqrt{\delta}$, the realization of d_k^{-i} is in $\neg\mathcal{Z}$. Fix any such realization d^* . Let \mathcal{U} be the set of realizations of noise traders' demands such that, added to $\Delta/\sqrt{2}^{k+1} + d^*$, they result in Q_{k+1} being in $B_{2\nu}(r)$. By construction, the probability of \mathcal{U} under the normal distribution $N(0, 1/2^{k+1})$ is at least $1 - \sqrt{\delta}$. Now, let $\mathcal{U}' = \mathcal{U} - (\lambda/\sqrt{2}^{k+1} - \Delta/\sqrt{2}^{k+1})$. Note that \mathcal{U}' is exactly the set of realizations of noise traders' demands such that, added to the $\lambda/\sqrt{2}^{k+1} + d^*$, they result in Q_{k+1} being in $B_{2\nu}(r)$. And by Lemma 2, for any $\gamma > 0$, by choosing $\sqrt{\delta}$ (and thus δ) appropriately small, we can ensure that the probability of \mathcal{U}' under $N(0, 1/2^{k+1})$ is at least $1 - \gamma$. Thus, under the alternative strategy, for this realization $d_k^{-i} \in \neg\mathcal{Z}$, with probability $1 - \gamma$, price y_{k+1} will be less than $x_r + \tau$.

Thus, following any history $\hat{\varphi}_k^i \in \mathcal{H}_k(\delta)$, the expected profit from the alternative strategy in period $k + 1$ is at least

$$\frac{\lambda}{\sqrt{2}^{k+1}} \left((x_b - 2\tau) - \sqrt{\delta} \max_{\omega} |X(\omega)| - (1 - \sqrt{\delta}) \left((1 - \gamma)(x_r + \tau) + \gamma \max_{\omega} |X(\omega)| \right) \right),$$

which for a sufficiently small δ (and thus γ) is greater than

$$\frac{\lambda}{\sqrt{2}^{k+1}} \left((x_b - 2\tau) - (x_r + 2\tau) \right) = \frac{\lambda\tau}{\sqrt{2}^{k+1}}.$$

Recall also that for any δ , $\lim_{k \rightarrow \infty} \overline{P}(\mathcal{H}_k(\delta))/\overline{P}(\mathcal{B}_k) = 1$ and for any sufficiently large k , $\overline{P}(\mathcal{B}_k) > \zeta\rho\xi > 0$. Thus, for any sufficiently large k , $\overline{P}(\mathcal{H}_k(\delta)) > \zeta\rho\xi/2$, and the unconditional expected continuation profit of trader i following period k must be greater than $(\zeta\rho\xi/2) \cdot (\lambda\tau/\sqrt{2}^{k+1})$ (because a feasible strategy for trader i is to not trade in any period later than $k + 1$, and in period $k + 1$ only trade following histories in $\mathcal{H}_k(\delta)$, buying $\lambda/\sqrt{2}^{k+1}$ units). This is greater than $\varepsilon/\sqrt{2}^k$ for a sufficiently small ε , contradicting the assumption that the original strategy was optimal for trader i .

APPENDIX B: PROOFS OF RESULTS IN SECTION 4

B.1. Proof of Theorem 2

Recall that strategy S_{ik} of player i at time t_k assigns a probability to each of the $2M_k + 1$ possible actions as a function of element $\pi \in \Pi_i$ observed by the player at time t_0 , demands d_1^i, \dots, d_{k-1}^i he has previously submitted, and aggregate market demands v_1, \dots, v_{k-1} observed in prior periods. Other than the aggregate market demands, all other values above come from finite sets. The proof will proceed by embedding the set of possible profiles of strategies into an appropriate topological vector space, considering the best response correspondence, and showing that all the conditions of a generalization of the Kakutani fixed point theorem are satisfied.

We first introduce some notation. Take any m and let F^m denote the set of bounded functions $f: \mathbb{Z}^m \rightarrow \mathbb{R}$. For any function $f \in F^m$, define its norm $\|f\|$ as

$$\|f\| = \sum_{x=(x_1, \dots, x_m) \in \mathbb{Z}^m} |f(x)| \frac{1}{2^{|x_1| + \dots + |x_m|}}.$$

Define distance between functions and the corresponding topology accordingly.

For every period k , player i , index j (denoting one of $2M_k + 1$ possible actions of player i in period k), element $\pi \in \Pi_i$, and possible prior demands d_1^i, \dots, d_{k-1}^i , let $H_{k,i,j,\pi,d_1^i, \dots, d_{k-1}^i} = F^{k-1}$. Let H be the Cartesian product of $H_{k,i,j,\pi,d_1^i, \dots, d_{k-1}^i}$ for all possible combinations of indices. Note that this Cartesian product is finite. For any $h \in H$, define its norm $\|h\|$ as the sum of the norms of all its individual elements, and define the induced distance and topology accordingly. Note that H is thus a normed vector space, and is therefore Hausdorff and locally convex (these properties will be used later in the proof, to invoke the appropriate fixed point theorem). Note also that set Ξ of possible strategy profiles is a nonempty, compact, and convex subset of space H : it is the set of profiles $S \in H$ such that, for any k, i, j, π , and d_1^i, \dots, d_{k-1}^i , element $S_{kij}(\pi, d_1^i, \dots, d_{k-1}^i)$ is a nonnegative function, and, for any k, i, π ,

and d_1^i, \dots, d_{k-1}^i , the sum $\sum_{j=1}^{2M_k+1} S_{kij}(\pi, d_1^i, \dots, d_{k-1}^i)$ is a constant function equal to 1 everywhere. The compactness of set S follows from the facts that any uniformly bounded infinite sequence of functions $f^m \in F^m$ has a subsequence converging pointwise to some limit function f^* , and that under the topology induced by the norm above, pointwise convergence of such a sequence of functions implies convergence in that topology.

Now, take any strategy profile $S = (S_1, \dots, S_n)$. We say that strategy Z_i of trader i is a best response to S if in every continuation game, after every possible history (including the ones where trader i has played actions inconsistent with S), for any other strategy Z'_i of trader i , the expected payoff from playing Z_i is at least as high as that from playing Z'_i , provided that prices are updated using pricing rule Y corresponding to S and other traders follow their strategies prescribed by S . Let $BR_i(S)$ be the set of best responses of trader i to strategy profile S and let $BR(S) = BR_1 \times \dots \times BR_n$. We need to show that there exists strategy profile S such that $S \in BR(S)$.

The existence of a fixed point S of mapping BR follows from a generalization of the Kakutani fixed point theorem to locally convex Hausdorff spaces (Glicksberg (1952)). To use the theorem, we need to show that BR has a closed graph and that, for any S , set $BR(S)$ is convex and non-empty.

Take any trader i , any period k , any history h observed by trader i up to k (i.e., his own demands and aggregate market demands prior to period k), any strategy profile S , and any strategy Z_i of trader i . Let $W(i, k, h, Z_i, S)$ denote the expected continuation profit of trader i , starting with period k , following history h , when he plays according to strategy Z_i , other strategic traders play according to their strategies in S , and the pricing rule is consistent with S . Note that function W is continuous in both Z_i and S . This, together with the compactness of set Ξ , implies that for any S , $BR(S)$ is non-empty and that the graph of BR is closed.

To see that for any $S \in \Xi$, set $BR(S)$ is convex, let $Z', Z'' \in BR(S)$ and let $Z_i = \alpha Z'_i + (1 - \alpha)Z''_i$ for some $\alpha \in (0, 1)$. Let us show that Z_i is a best response to S after any history. The proof is by induction, going from the last period forward. Let $k = K$, and take any history up to k . Following that history, any action taken with positive probability under either Z' or Z'' has to result in the same expected payoff, and so after any history, the expected continuation payoff from Z', Z'' , and Z has to be the same. Suppose we have shown that after any history up to period $k = k' + 1$, expected continuation payoffs under Z', Z'' , and Z are the same. Take any history h up to period $k = k'$. By assumption, the continuation payoffs under Z' and Z'' are the same (and the best attainable). Suppose the continuation payoff under Z is less than that. Consider “strategy” \hat{Z} , under which, following history h , trader i plays according to Z' with probability α and plays according to Z'' with probability $1 - \alpha$ (this is not a strategy according to our definition, but we can still compute expected payoffs under that behavior rule). Clearly, the expected continuation payoff from \hat{Z} is equal to that from Z' and Z'' . Moreover, after any history in

period $k + 1$, the expected continuation payoff from \hat{Z} is equal to that from Z' and Z'' , and thus, by induction hypothesis, to that from Z . But in period k , the probability distributions of actions of trader i under \hat{Z} and Z are identical, and thus the expected payoff in period k under the two strategies is also the same. Therefore, the continuation payoff following history h under strategy Z is the same as under \hat{Z} , which in turn is equal to that under Z' and Z'' , and therefore Z is also a best response.

Thus, all the conditions of the fixed point theorem are satisfied, and mapping $\text{BR}(S)$ has a fixed point S^* .

B.2. Proof of Theorem 3

Consider sequence Γ_K of games obtained from Γ^∞ by removing all trading periods following period K . By Theorem 2, each of these games has an equilibrium, S_K^* . Consider sequence S_K^* . Take any period K' , and for $K \geq K'$, let $S_{K,K'}^*$ denote the sequence of strategy profiles S_K^* restricted to the first K' periods. By the compactness of the set of strategy profiles in game $\Gamma_{K'}$ (in the corresponding normed space $H_{K'}$, as in the proof of Theorem 2), sequence $S_{K,K'}^*$ has a limit point for any K' . We can now construct equilibrium S^* of game Γ as follows. Take the limit point of $S_{K,1}^*$ (say, Z_1), and take subsequence of equilibria Seq_1 converging to it. In period 1, under S^* , traders act according to Z_1 . Now, consider the restriction of Seq_1 to the first two periods. In that subsequence, strategies in the first period converge to Z_1 , by construction. Strategies in the second period also have a limit point (call it Z_2) and a subsequence Seq_2 of Seq_1 converging to it. In period 2, under S^* , let traders act according to Z_2 . Proceeding analogously, we add Z_3, Z_4 , etc., to S^* , ending up with a profile of strategies.

The fact that S^* is an equilibrium follows from the facts that (i) game Γ^∞ is continuous at infinity (i.e., for any $\varepsilon > 0$, one can find K' such that the payoffs of players are affected by their actions after period K' by less than ε), and (ii) as in the proof of Theorem 2, “counterfactual” continuation payoffs $W(i, k, h, Z_i, S)$ in any finite restriction of Γ are continuous in strategies.

B.3. Proof of Theorem 4

The proof follows the same steps as the proof of Theorem 1 for Γ^K . Moreover, Steps 0 and 2–8 are unchanged, and only Steps 1 and 9 need to be slightly adjusted, to address discretization, as described below.

STEP 1: Only the last equation in this step needs to be changed. In that equation, $\text{Var}(u_{k'})$ is now equal to $V_{\sigma,d}/2^{k'}$ instead of $1/2^{k'}$, where $V_{\sigma,d}$ is the variance of the discretized normal distribution with mean 0, variance parameter σ^2 , and unit of discretization d . This difference, however, does not change the conclusion (the last paragraph) of the step.

STEP 9: This step needs two minor modifications. First, we cannot pick an arbitrary λ , so we pick a specific one: let $\lambda = d$ (i.e., the smallest allowed positive unit of discretization).

Second, in the paragraph on the original and “shifted” sets of realizations of draws from the normal distribution (\mathcal{U} and \mathcal{U}'), instead of being interpreted as the actual realizations of noise traders’ demands (as in the proof of Theorem 1 for Γ^K) these sets should now be interpreted as sets of normal draws \tilde{u}_k that are subsequently rounded to determine the actual demands from noise traders.

APPENDIX C: PROOF OF THEOREM 5

STEP 0: Suppose that for some $\varepsilon > 0$, the statement is not true. Then one can find an increasing sequence K_m and a corresponding sequence of equilibria of games $\Gamma_{K_m}, (S_m^*, Y_m^*)$, such that in these equilibria, for each m , $E[|y_{K_m} - X(\omega)|] \geq \varepsilon$. Consider this sequence.

STEP 1: In this step, we select a “well-behaved” subsequence from the above sequence.

Pick any m and consider equilibrium (S_m^*, Y_m^*) . Take any $k \leq K_m$. Define residual variance $\Sigma(k, m)$ as the expected square of the difference between price y_k and the true value of the security, that is, $\Sigma(k, m) = E[(X(\omega) - y_k)^2]$, where the expectation is taken over all possible draws of the original state and noise traders’ demands, as well as the (possibly) mixed strategy of the strategic trader. By construction, y_k is an unbiased estimator of $X(\omega)$, and so we have $E[X(\omega)|y_k] = y_k$, $E[X(\omega) - y_k|y_k] = 0$, and hence $E[X(\omega) - y_k] = 0$ and $Cov(X(\omega) - y_k, y_k) = 0$. Thus, $Var(X(\omega)) = Var((X(\omega) - y_k) + y_k) = Var(X(\omega) - y_k) + Var(y_k) + 2Cov(X(\omega) - y_k, y_k) = E[(X(\omega) - y_k)^2] + Var(y_k) = \Sigma(k, m) + Var(y_k)$.

Let $V_m^1 = \Sigma(\lfloor \frac{1}{2}m \rfloor, m)$, that is, the residual variance after approximately half of the trading has occurred. Sequence V_m^1 is bounded (by $Var(X(\omega))$), and so it must have a converging subsequence. Consider any such subsequence, Seq_1 , and let V^1 be the limit of V_m^1 along this subsequence.

Next, let $V_m^2 = \Sigma(\lfloor \frac{3}{4}m \rfloor, m)$, that is, the residual variance after approximately three quarters of the trading has occurred. Consider sequence V_m^2 for m restricted to subsequence Seq_1 . It also must have a converging subsequence, Seq_2 . Let V^2 be the limit of V_m^2 along this subsequence.

Proceeding analogously, we obtain subsequences Seq_l and the corresponding limits V^l for all $l \geq 1$. Now, let m_1 be the first element of subsequence Seq_1 , $m_2 > m_1$ be the second element of subsequence Seq_2 , and so on. The resulting subsequence m_j has the property that, for every l , the limit of V_m^l along this subsequence is equal to V^l . Without loss of generality, to simplify the notation, assume that this subsequence is the original sequence $\{1, 2, 3, \dots\}$.

Note that by construction, for any l , $V^l \geq V^{l+1} \geq 0$.²² Thus, as l goes to infinity, V^l converges to some $V^\infty \geq 0$.

STEP 2: In this step, we obtain an upper bound on the expected losses of noise traders as time approaches the end of trading, for equilibria in the chosen subsequence.

Take any m and $k < K_m$. Let $\Psi(k, m)$ denote the unconditional expected total profit of noise traders arriving after period k , that is,

$$\Psi(k, m) = \mathbb{E} \left[\sum_{k'=k+1}^{K_m} (X(\omega) - y_{k'})u_{k'} \right].$$

Since the strategic trader’s expected continuation payoff after any period k is nonnegative (because he can always guarantee himself a payoff of zero by simply not trading), and the expected continuation payoff of market makers is zero (by assumption), the expected payoff of noise traders arriving after period k cannot be positive, and so $\Psi(k, m) \leq 0$. In the remainder of this step, we will obtain an upper bound on $|\Psi(k, m)|$.

Since $\mathbb{E}[u_{k'}] = 0$ and $u_{k'}$ is independent of anything that happened before period k' , we have $\mathbb{E}[X(\omega)u_{k'}] = \mathbb{E}[X(\omega)]\mathbb{E}[u_{k'}] = 0$, $\mathbb{E}[y_{k'-1}u_{k'}] = \mathbb{E}[y_{k'-1}]\mathbb{E}[u_{k'}] = 0$, and thus $\mathbb{E}[(X(\omega) - y_{k'})u_{k'}] = \mathbb{E}[(y_{k'-1} - y_{k'})u_{k'}]$. Therefore,

$$\Psi(k, m) = \mathbb{E} \left[\sum_{k'=k+1}^{K_m} (y_{k'-1} - y_{k'})u_{k'} \right].$$

Take any A and B such that $1 \leq A < B \leq K_m$ and consider $\mathbb{E}[\sum_{k'=A+1}^B (y_{k'-1} - y_{k'})u_{k'}]$. We have

$$\begin{aligned} & \left| \mathbb{E} \left[\sum_{k'=A+1}^B (y_{k'-1} - y_{k'})u_{k'} \right] \right| \\ &= \left| \sum_{k'=A+1}^B \mathbb{E}[(y_{k'-1} - y_{k'})u_{k'}] \right| \\ &\leq \sum_{k'=A+1}^B \left| \mathbb{E}[(y_{k'-1} - y_{k'})u_{k'}] \right| \end{aligned}$$

²²To see this, take any $m > 1$ and any k' and k'' such that $1 \leq k' < k'' \leq m$. We have $\text{Var}(X(\omega)) = \Sigma(k', m) + \text{Var}(y_{k'}) = \Sigma(k'', m) + \text{Var}(y_{k''})$, and since sequence y_k is a martingale, we also have $\text{Var}(y_{k'}) \leq \text{Var}(y_{k''})$. Thus, $\Sigma(k', m) \geq \Sigma(k'', m)$.

$$\begin{aligned} &\leq \sum_{k'=A+1}^B \sqrt{\mathbb{E}[(y_{k'-1} - y_{k'})^2] \mathbb{E}[(u_{k'})^2]} \\ &= \frac{\sigma}{\sqrt{K_m}} \sum_{k'=A+1}^B \sqrt{\mathbb{E}[(y_{k'-1} - y_{k'})^2]}, \end{aligned}$$

where the last inequality follows from the Cauchy–Schwarz inequality and the last equality follows from the original assumption about the distribution of $u_{k'}$.

Now, each y_k is equal to the expectation of $X(\omega)$ based on the publicly available information up to period k , and thus changes in y_k prior to period B are uncorrelated with each other and with $X(\omega) - y_B$. Thus, we have

$$\begin{aligned} \mathbb{E}[(X(\omega) - y_A)^2] &= \mathbb{E}\left[\left((X(\omega) - y_B) + \sum_{k'=A+1}^B (y_{k'} - y_{k'-1})\right)^2\right] \\ &= \mathbb{E}[(X(\omega) - y_B)^2] + \sum_{k'=A+1}^B \mathbb{E}[(y_{k'-1} - y_{k'})^2], \end{aligned}$$

and so

$$\sum_{k'=A+1}^B \mathbb{E}[(y_{k'-1} - y_{k'})^2] = \Sigma(A, m) - \Sigma(B, m).$$

Since square root is a concave function, Jensen’s inequality implies that

$$\begin{aligned} \sum_{k'=A+1}^B \sqrt{\mathbb{E}[(y_{k'-1} - y_{k'})^2]} &= (B - A) \frac{\sum_{k'=A+1}^B \sqrt{\mathbb{E}[(y_{k'-1} - y_{k'})^2]}}{B - A} \\ &\leq (B - A) \sqrt{\frac{\sum_{k'=A+1}^B \mathbb{E}[(y_{k'-1} - y_{k'})^2]}{B - A}} \\ &= \sqrt{B - A} \sqrt{\Sigma(A, m) - \Sigma(B, m)}. \end{aligned}$$

We can now obtain a bound on $|\Psi(k, m)|$. Pick any $\varphi > 0$. Recall the sequence V^l constructed in Step 1. It decreases monotonically and converges to V^∞ . Take the smallest a such that $V^a - V^\infty < (\varphi/4)^2$. Take the smallest $b > a$ such that $\frac{2^a}{2^b} V^b < (\varphi/4)^2$ (such b exists because sequence V^l is bounded). Now, for any large m (e.g., such that $K_m > 2^{b+1}$), let $k(\varphi, m)$ be the largest k

such that $\frac{k}{K_m} \leq 1 - \frac{1}{2^a}$ and let $g(\varphi, m)$ be the largest g such that $\frac{g}{K_m} \leq 1 - \frac{1}{2^b}$.²³ Then we have

$$\begin{aligned}
 & |\Psi(k(\varphi, m), m)| \\
 &= \left| \mathbb{E} \left[\sum_{k'=k(\varphi, m)+1}^{K_m} (y_{k'-1} - y_{k'}) u_{k'} \right] \right| \\
 &\leq \frac{\sigma}{\sqrt{K_m}} \sum_{k'=k(\varphi, m)+1}^{K_m} \sqrt{\mathbb{E}[(y_{k'-1} - y_{k'})^2]} \\
 &= \frac{\sigma}{\sqrt{K_m}} \left(\sum_{k'=k(\varphi, m)+1}^{g(\varphi, m)} \sqrt{\mathbb{E}[(y_{k'-1} - y_{k'})^2]} \right. \\
 &\quad \left. + \sum_{k'=g(\varphi, m)+1}^{K_m} \sqrt{\mathbb{E}[(y_{k'-1} - y_{k'})^2]} \right) \\
 &\leq \frac{\sigma}{\sqrt{K_m}} \left(\sqrt{g(\varphi, m) - k(\varphi, m)} \sqrt{\Sigma(k(\varphi, m), m) - \Sigma(g(\varphi, m), m)} \right. \\
 &\quad \left. + \sqrt{K_m - g(\varphi, m)} \sqrt{\Sigma(g(\varphi, m), m) - \Sigma(K_m, m)} \right) \\
 &< \sigma \left(\sqrt{\frac{K_m - k(\varphi, m)}{K_m}} \sqrt{\Sigma(k(\varphi, m), m) - \Sigma(g(\varphi, m), m)} \right. \\
 &\quad \left. + \sqrt{\frac{K_m - g(\varphi, m)}{K_m}} \sqrt{\Sigma(g(\varphi, m), m)} \right).
 \end{aligned}$$

From earlier arguments, and by construction, we know that as m goes to infinity,

- (i) $\frac{K_m - k(\varphi, m)}{K_m}$ converges to $\frac{1}{2^a}$ and $\frac{K_m - g(\varphi, m)}{K_m}$ converges to $\frac{1}{2^b}$;
- (ii) $\Sigma(g(\varphi, m), m)$ converges to V^b ; and
- (iii) $\Sigma(k(\varphi, m), m) - \Sigma(g(\varphi, m), m)$ converges to $V^a - V^b$, which is less than or equal to $V^a - V^\infty$.

Moreover, a and b were chosen in such a way that $V^a - V^\infty < (\frac{\sigma}{4})^2$ and $V^b \frac{1}{2^b} < (\frac{\sigma}{4})^2 \frac{1}{2^a}$. Combining all of the above, we find that for any $\xi > 1$, for a sufficiently

²³That is, $\frac{k(\varphi, m)}{K_m} \approx 1 - \frac{1}{2^a}$ and $\frac{g(\varphi, m)}{K_m} \approx 1 - \frac{1}{2^b}$, and so $\Sigma(k(\varphi, m), m) \approx V^a$ and $\Sigma(g(\varphi, m), m) \approx V^b$.

large m ,

$$|\Psi(k(\varphi, m), m)| < \xi\sigma \left(\sqrt{\frac{1}{2^a}} \frac{\varphi}{4} + \sqrt{\frac{1}{2^a}} \frac{\varphi}{4} \right),$$

and so, in particular, by setting $\xi = 2$ we find that, for a sufficiently large m ,

$$|\Psi(k(\varphi, m), m)| < \varphi\sigma\sqrt{\frac{1}{2^a}}.$$

To interpret this bound, recall that $\frac{k(\varphi, m)}{K_m}$ is, by construction, approximately equal to $1 - \frac{1}{2^a}$, and that a was chosen only based on φ . Thus, this bound tells us that for any arbitrarily small but positive φ , there exists time $t = 1 - \frac{1}{2^a} < 1$ such that the expected loss of noise traders in the continuation of the game from period t until the end of trading is less than $\varphi\sigma\sqrt{1-t}$, for any sufficiently large m .

STEP 3: In this step, we identify potential arbitrage opportunities.

Pick any m and consider the corresponding K_m and (S_m^*, Y_m^*) . Also, take any $k \leq K_m$. By assumption, $E[|y_{K_m} - X(\omega)|] \geq \varepsilon$, which implies $E[(y_{K_m} - X(\omega))^2] \geq (E[|y_{K_m} - X(\omega)|])^2 \geq \varepsilon^2$.

This, in turn, implies that for any $k \leq K_m$, $E[(y_k - X(\omega))^2] \geq \varepsilon^2$.

Let $Q(k, m)$ be the random variable (with values in the $|\Omega|$ -dimensional unit simplex) denoting the beliefs of a market maker (or any other outside observer) about the true state ω , after period k in game Γ_{K_m} under equilibrium (S_m^*, Y_m^*) .

By the law of iterated expectations, we can rewrite the last inequality as

$$E_{Q(k, m)}[E[(y_k - X(\omega))^2 | Q(k, m)]] \geq \varepsilon^2,$$

which, together with the fact that $y_k = E[X(\omega) | Q(k, m)]$, in turn implies that for some $\varepsilon_1 > 0$, depending only on the primitives of the model and on the chosen ε , but not on k or m , with probability at least ε_1 , the realization of $Q(k, m)$ is such that

$$\text{Var}(X(\omega) | Q(k, m)) > \varepsilon_1.$$

Let us call such realizations of $Q(k, m)$ “arbitrageable.”

Next, since set Ω is finite, we can find $\delta > 0$ such that, for any arbitrageable realization q , we can find two states, ω_1 and ω_2 , such that the probability of each of them under q is greater than δ and also $|X(\omega_1) - X(\omega_2)| > \delta$.²⁴

²⁴To see this, take any $\delta < 1/|\Omega|$ and any arbitrageable realization q . Take any state ω_0 such that $q(\omega_0) \geq 1/|\Omega| > \delta$. We have $\text{Var}(X(\omega) | q) \leq \sum_{\omega' \in \Omega} q(\omega')(X(\omega') - X(\omega_0))^2$. If for

Crucially, ε_1 and δ depend only on the ε chosen in Step 0 and on the primitives of the model, but not on m and k .

STEP 4: In this step, we show how the strategic trader can take advantage of the potential arbitrage opportunities identified in the previous step. More formally, we obtain a lower bound on the expected profits of the strategic trader following such opportunities.

Take any $m, k < K_m$, any arbitrageable realization $Q(k, m)$, and the states of nature ω_1 and ω_2 identified in Step 3. Consider any public history h_k (i.e., a sequence of aggregate demands v_1, v_2, \dots, v_k) leading to this realization. Consider a counterfactual trading policy for the strategic trader following this history: he completely stops trading. Denote by $\bar{y}(0)$ the expectation of $\frac{1}{K_m - k} \sum_{k'=k+1}^{K_m} y_{k'}$ (i.e., the expected average market price of the security after period k) under that policy following history h_k . Clearly, at least one of $X(\omega_1)$ or $X(\omega_2)$ differs from $\bar{y}(0)$ by at least $\frac{\delta}{2}$; assume for concreteness that it is $X(\omega_1)$. Also, note that conditioning on the true state of the world does not change the expectation of the average future price, since under this counterfactual (non)trading policy, the true state has no influence on the future evolution of prices. Without loss of generality, assume that $X(\omega_1) - \bar{y}(0) \geq \frac{\delta}{2}$ (the case when $X(\omega_1) - \bar{y}(0) \leq -\frac{\delta}{2}$ is completely symmetric: instead of buying the security, as below, the trader would sell it).

Consider now the following continuation trading strategy for the strategic trader after history h_k . If the true state of the world is not ω_1 , do not trade. If the true state of the world is ω_1 , buy $c\Delta$ units of security in every period, where c is a constant to be determined below and $\Delta = \frac{\sigma}{\sqrt{K_m(K_m - k)}}$. Note that if $c = 1$, then over the $K_m - k$ trading periods remaining after history h_k , the trader will end up buying a total of $\sigma\sqrt{\frac{K_m - k}{K_m}}$ units, which is equal to one standard deviation of the total demand from noise traders over that period.

We now need to carefully choose the constant c . Let $\gamma(v_{k+1}, v_{k+2}, \dots, v_{K_m})$ be the average price of the security during periods $k + 1$ through K_m when the public history of total demands is $h_k, v_{k+1}, \dots, v_{K_m}$. By definition,

$$\bar{y}(0) = \int_{\mathbb{R}^{K_m - k}} \gamma(u_{k+1}, u_{k+2}, \dots, u_{K_m}) f(\bar{u}) d\bar{u},$$

where

$$f(\bar{u}) = \left(\sqrt{\frac{K_m}{2\pi\sigma^2}} \right)^{K_m - k} e^{-(K_m/(2\sigma^2))(u_{k+1}^2 + \dots + u_{K_m}^2)}$$

every ω' , we have $q(\omega') \leq \delta$ or $|X(\omega') - X(\omega_0)| \leq \delta$, then this expression is less than $|\Omega|\delta(2\max_{\omega} \{X(\omega)\})^2 + |\Omega|\delta^2$, which is smaller than ε_1 for a sufficiently small δ .

is the density of the multivariate normal distribution of noise traders' demands in periods $k + 1$ and later.

Now, let $\bar{y}(z)$ denote the expected average price of the security following history h_k when the strategic trader buys z units in every period $k' > k$. We have

$$\begin{aligned} \bar{y}(z) &= \int_{\mathbb{R}^{K_m-k}} \gamma(u_{k+1} + z, u_{k+2} + z, \dots, u_{K_m} + z) \\ &\quad \times \left(\sqrt{\frac{K_m}{2\pi\sigma^2}} \right)^{K_m-k} e^{-(K_m/(2\sigma^2))(u_{k+1}^2 + \dots + u_{K_m}^2)} d\bar{u}. \end{aligned}$$

We will now pick $c \in (0, 1)$, independent of k and m , in such a way that $|\bar{y}(c\Delta) - \bar{y}(0)|$ is small. Let $\gamma^* = \max_{\omega \in \Omega} \{|X(\omega)|\}$. Clearly, $|\gamma(\cdot)|$ is never greater than γ^* . Also, let $M > 0$ be such that the probability that the absolute value of a random variable drawn from the standard normal distribution (with mean 0 and variance 1) is greater than M is equal to $\frac{\delta}{16\gamma^*}$. Both γ^* and M are independent of k and m .

Consider $|\bar{y}(c\Delta) - \bar{y}(0)|$. We have

$$\begin{aligned} |\bar{y}(c\Delta) - \bar{y}(0)| &= \left| \int_{\mathbb{R}^{K_m-k}} \gamma(\bar{u} + c\Delta) f(\bar{u}) d\bar{u} - \int_{\mathbb{R}^{K_m-k}} \gamma(\bar{u}) f(\bar{u}) d\bar{u} \right| \\ &= \left| \int_{\mathbb{R}^{K_m-k}} \gamma(\bar{u}) f(\bar{u} - c\Delta) d\bar{u} - \int_{\mathbb{R}^{K_m-k}} \gamma(\bar{u}) f(\bar{u}) d\bar{u} \right|, \end{aligned}$$

where the second equality is obtained by the change of variables under the first integral.

We will now break up the last expression into several parts. Let V_1 be the subset of \mathbb{R}^{K_m-k} in which $|\sum_{k'=k+1}^{K_m} u_{k'}| \leq (M + 1)\sigma\sqrt{\frac{K_m-k}{K_m}}$ and let $V_2 = \mathbb{R}^{K_m-k} \setminus V_1$. Then,

$$\begin{aligned} |\bar{y}(c\Delta) - \bar{y}(0)| &\leq \left| \int_{V_1} \gamma(\bar{u})(f(\bar{u} - c\Delta) - f(\bar{u})) d\bar{u} \right| \\ &\quad + \left| \int_{V_2} \gamma(\bar{u}) f(\bar{u} - c\Delta) d\bar{u} \right| \\ &\quad + \left| \int_{V_2} \gamma(\bar{u}) f(\bar{u}) d\bar{u} \right|. \end{aligned}$$

Since for any $\bar{u} \in V_2$, $|\gamma(\bar{u})| < \gamma^*$ and (because $c < 1$ and $(K_m - k)\Delta = \sigma\sqrt{\frac{K_m-k}{K_m}}$) both $|\sum u_{k'}|$ and $|\sum(u_{k'} - c\Delta)|$ are greater than $M\sigma\sqrt{\frac{K_m-k}{K_m}}$, and also because $\sum u_{k'}$ is distributed normally with mean zero and standard deviation $\sigma\sqrt{\frac{K_m-k}{K_m}}$, by the choice of M we know that each of the last two terms is

bounded by $\gamma^* \frac{\delta}{16\gamma^*}$, and thus we have

$$|\bar{y}(c\Delta) - \bar{y}(0)| \leq \left| \int_{V_1} \gamma(\bar{u})(f(\bar{u} - c\Delta) - f(\bar{u})) d\bar{u} \right| + \frac{\delta}{8}.$$

Now,

$$\left| \int_{V_1} \gamma(\bar{u})(f(\bar{u} - c\Delta) - f(\bar{u})) d\bar{u} \right| \leq \int_{V_1} \gamma^* \left| \frac{f(\bar{u} - c\Delta)}{f(\bar{u})} - 1 \right| f(\bar{u}) d\bar{u}.$$

Next,

$$\begin{aligned} \frac{f(\bar{u} - c\Delta)}{f(\bar{u})} &= \frac{e^{-(K_m/(2\sigma^2))((u_{k+1}-c\Delta)^2+\dots+(u_{K_m}-c\Delta)^2)}}{e^{-(K_m/(2\sigma^2))(u_{k+1}^2+\dots+u_{K_m}^2)}} \\ &= e^{-(K_m/(2\sigma^2))(\sum_{k'=k+1}^{K_m} ((u_{k'}-c\Delta)^2-u_{k'}^2))} \\ &= e^{-(K_m/(2\sigma^2))((K_m-k)c^2\Delta^2-2c\Delta\sum u_{k'})}. \end{aligned}$$

Since $\bar{u} \in V_1$, we can place a bound on $\ln \frac{f(\bar{u}-c\Delta)}{f(\bar{u})}$:

$$\begin{aligned} \left| \ln \frac{f(\bar{u} - c\Delta)}{f(\bar{u})} \right| &\leq \frac{K_m}{2\sigma^2} \left((K_m - k)c^2\Delta^2 + 2c\Delta(M + 1)\sigma\sqrt{\frac{K_m - k}{K_m}} \right) \\ &= \frac{K_m}{2\sigma^2} \left((K_m - k)c^2 \left(\frac{\sigma}{\sqrt{K_m(K_m - k)}} \right)^2 \right. \\ &\quad \left. + 2c \left(\frac{\sigma}{\sqrt{K_m(K_m - k)}} \right) (M + 1)\sigma\sqrt{\frac{K_m - k}{K_m}} \right) \\ &= \frac{c^2}{2} + c(M + 1). \end{aligned}$$

Crucially, this bound does not depend on m or k , and so we can pick constant $c > 0$ such that $\ln \frac{f(\bar{u}-c\Delta)}{f(\bar{u})}$ is arbitrarily close to 0. Hence, $\frac{f(\bar{u}-c\Delta)}{f(\bar{u})} - 1$ can also be made arbitrarily close to zero for an appropriate choice of $c > 0$; in particular, its absolute value can be made less than $\frac{\delta}{8\gamma^*}$. Choose such c . Then

$$|\bar{y}(c\Delta) - \bar{y}(0)| \leq \int_{V_1} \gamma^* \frac{\delta}{8\gamma^*} f(\bar{u}) d\bar{u} + \frac{\delta}{8} < \frac{\delta}{4}.$$

Since by construction, $X(\omega_1) - \bar{y}(0) \geq \frac{\delta}{2}$, we have $X(\omega_1) - \bar{y}(c\Delta) > \frac{\delta}{4}$. Thus, following public history h_k , when the true state of nature is ω_1 , by buying $c\Delta$ units in each period $k + 1, \dots, K_m$, the strategic trader can obtain the expected

profit of

$$(X(\omega_1) - \bar{y}(c\Delta))(K_m - k)c\Delta > \frac{\delta}{4}c\sigma\sqrt{\frac{K_m - k}{K_m}}.$$

By construction, conditional on having an arbitrageable realization of $Q(k, m)$, with probability at least $\delta > 0$, the true state of the world is such that the strategic trader can obtain this profit. The probability of having an arbitrageable realization (for any k and m) is at least ε_1 . Finally, in all other cases, the continuation value of the strategic trader is nonnegative. Combining these facts, we find that for $\lambda = \varepsilon_1\delta\frac{\delta}{4}c > 0$, for any k and m , the expected continuation value of the strategic trader in equilibrium (S_m^*, Y_M^*) of game Γ_{K_m} following period k is at least $\lambda\sigma\sqrt{1 - \frac{k}{K_m}}$.

Combined with the fact that the expected continuation profit of market makers following any period k is zero, this bound contradicts the finding in Step 2 that, for any $\varphi > 0$, for a sufficiently large m , the expected loss of noise traders following period $k(\varphi, m) \approx K_m(1 - \frac{1}{2^a})$ is less than $\varphi\sigma\sqrt{\frac{1}{2^a}}$.

APPENDIX D: PROOFS OF RESULTS IN SECTION 6

D.1. Proof of Theorem 7

The “if” direction of the theorem is, in essence, proved in Proposition 3 of DeMarzo and Skiadas (1999) using the following “adding-up” argument, which also clarifies the intuition behind the condition. Suppose X is non-separable and take P and v satisfying the requirements of non-separability. Take functions λ_i as in the statement of the theorem. Consider the unconditional expectation $E[(X(\omega) - v) \sum_i \lambda_i(\Pi_i(\omega))]$ under P . On one hand, by the choice of functions λ_i , the expectation is strictly positive. On the other hand,

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

where the last equality follows from requirement (ii) of Definition 2.

The “only if” direction follows from Proposition 4 of DeMarzo and Skiadas (1999), but also has the following short self-contained proof.²⁵ Suppose security

²⁵I am grateful to Yury Makarychev for this proof.

X is separable. Take any $v \in \mathbb{R}$. Ignore all states ω with $X(\omega) = v$ and let $h = 1, \dots, H$ denote the remaining states. Let $m = 1, \dots, M$ denote the elements π of all players' partitions; if the same element π belongs to partitions of two (or more) players, it should be indexed correspondingly many times.

Construct an $M \times H$ matrix A as follows. If state h belongs to element m of players' partitions, then the entry in row m and column h of the matrix is equal to $X(h) - v$. Otherwise, it is equal to zero.

By Gordan's transposition theorem,²⁶ exactly one of the following two systems of equations and inequalities has a solution:

1. $Ax = 0, x \geq 0, x \neq 0$ (where $x \in \mathbb{R}^H$);
2. $A^T y > 0$ (where $y \in \mathbb{R}^M$).

Note that if system (1) has a solution, then the security is non-separable. (Take solution x of system (1); rescale it so that its elements sum to 1; and use these 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. Take any solution y , and for every player i and element $m \in \Pi_i$, let $\lambda_i(m) = y_m$. Then for any $\omega \in \Omega$ such that $X(\omega) \neq v$, we have $(X(\omega) - v) \sum_i \lambda_i(\Pi_i(\omega)) > 0$.

D.2. Proof of Corollary 1

Let $x_{(j)}(\omega)$ denote the j th lowest of the n numbers $(x_i(\Pi_i(\omega)))_{i=1, \dots, n}$. Take any $v \in \mathbb{R}$. For every i and ω , set $\lambda_i(\Pi_i(\omega))$ equal to 1 if $x_i(\Pi_i(\omega)) \geq v$ and to $-\frac{n-j+1}{j-0.5}$ if $x_i(\Pi_i(\omega)) < v$. Then $X(\omega) > v \Leftrightarrow x_{(j)}(\omega) > v \Rightarrow \sum_i \lambda_i(\Pi_i(\omega)) \geq 1 \cdot (n - j + 1) - \frac{n-j+1}{j-0.5} \cdot (j - 1) = (n - j + 1)(1 - \frac{j-1}{j-0.5}) > 0$ and $X(\omega) < v \Leftrightarrow x_{(j)}(\omega) < v \Rightarrow \sum_i \lambda_i(\Pi_i(\omega)) \leq 1 \cdot (n - j) - \frac{n-j+1}{j-0.5} \cdot j < (n - j) - (n - j + 1) < 0$. Thus, $X(\omega) \neq v \Rightarrow (X(\omega) - v) \sum_i \lambda_i(\Pi_i(\omega)) > 0$, and since the initial choice of v was arbitrary, by Theorem 7 security X is separable.

D.3. Proof of Corollary 2

Assume that function f is increasing (the proof for the opposite case is analogous) and continuous and unbounded (since security X takes only a finite number of values, this is w.l.o.g.). Take any $v \in \mathbb{R}$. Take any z such that $f(z) = v$. Setting $\lambda_i(\Pi_i(\omega)) = x_i(\Pi_i(\omega)) - \frac{z}{n}$ for every player i and state ω , we get $X(\omega) \neq v \Rightarrow (X(\omega) - v) \sum_i \lambda_i(\Pi_i(\omega)) = (f(\sum_i x_i(\Pi_i(\omega))) - f(z)) \cdot (\sum_i x_i(\Pi_i(\omega)) - z) > 0$ (the inequality is strict, because $X(\omega) = f(\sum_i x_i(\Pi_i(\omega))) \neq v = f(z)$ implies $\sum_i x_i(\Pi_i(\omega)) \neq z$). Since the initial choice of v was arbitrary, Theorem 7 implies that security X is separable.

²⁶<http://com.springer.de/m/m130240.htm>.

REFERENCES

- AMIHUD, Y., AND H. MENDELSON (1980): "Dealership Market: Market-Making With Inventory," *Journal of Financial Economics*, 8 (1), 31–53. [2599]
- AUMANN, R. J. (1976): "Agreeing to Disagree," *The Annals of Statistics*, 4 (6), 1236–1239. [2601]
- BACK, K. (1992): "Insider Trading in Continuous Time," *Review of Financial Studies*, 5 (3), 387–409. [2610]
- BACK, K., AND S. BARUCH (2004): "Information in Securities Markets: Kyle Meets Glosten and Milgrom," *Econometrica*, 72 (2), 433–465. [2602,2606,2610]
- BACK, K., C. H. CAO, AND G. A. WILLARD (2000): "Imperfect Competition Among Informed Traders," *Journal of Finance*, 55 (5), 2117–2155. [2597,2610]
- BERGIN, J., AND A. BRANDENBURGER (1990): "A Simple Characterization of Stochastically Monotone Functions," *Econometrica*, 58 (5), 1241–1243. [2615]
- BLACK, F. (1971a): "Towards a Fully Automated Stock Exchange, Part I," *Financial Analyst Journal*, 27 (4), 28–35, 44. [2599]
- (1971b): "Towards a Fully Automated Stock Exchange, Part II," *Financial Analyst Journal*, 27 (6), 24–26, 86–87. [2599]
- BRIER, G. W. (1950): "Verification of Forecasts Expressed in Terms of Probability," *Monthly Weather Review*, 78 (1), 1–3. [2603]
- CALDENTY, R., AND E. STACCHETTI (2010): "Insider Trading With a Random Deadline," *Econometrica*, 78 (1), 245–283. [2602,2610]
- CHEN, Y., AND D. M. PENNOCK (2007): "A Utility Framework for Bounded-Loss Market Makers," in *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*. Corvallis, OR: AUAI Press, 49–56. [2599]
- CHEN, Y., S. DIMITROV, R. SAMI, D. M. REEVES, D. M. PENNOCK, R. D. HANSON, L. FORTNOW, AND R. GONEN (2010): "Gaming Prediction Markets: Equilibrium Strategies With a Market Maker," *Algorithmica*, 58 (4), 930–969. [2600]
- CHEN, Y., D. M. REEVES, D. M. PENNOCK, R. D. HANSON, L. FORTNOW, AND R. GONEN (2007): "Bluffing and Strategic Reticence in Prediction Markets," in *Proceedings of the 3rd International Workshop on Internet and Network Economics*. Berlin: Springer, 70–81. [2600]
- CONSENSUS POINT (2011a): "Customers," available at <http://www.consensuspoint.com/customers>, retrieved September 8, 2011. [2600]
- (2011b): "Prediction Markets," available at <http://www.consensuspoint.com/business-value/prediction-markets>, retrieved September 8, 2011. [2600]
- CRIPPS, M. W., G. J. MAILATH, AND L. SAMUELSON (2004): "Imperfect Monitoring and Impermanent Reputations," *Econometrica*, 72 (2), 407–432. [2607]
- (2007): "Disappearing Private Reputations in Long-Run Relationships," *Journal of Economic Theory*, 134 (1), 287–316. [2607]
- DEMARZO, P., AND C. SKIADAS (1998): "Aggregation, Determinacy, and Informational Efficiency for a Class of Economies With Asymmetric Information," *Journal of Economic Theory*, 80 (1), 123–152. [2601,2604,2613,2615]
- (1999): "On the Uniqueness of Fully Informative Rational Expectations Equilibria," *Economic Theory*, 13 (1), 1–24. [2601,2604,2613,2615,2643]
- DIMITROV, S., AND R. SAMI (2008): "Non-Myopic Strategies in Prediction Markets," in *Proceedings of the 9th ACM Conference on Electronic Commerce*. New York: ACM, 200–209. [2598,2600,2603]
- DUBEY, P., J. GEANAKOPOLOS, AND M. SHUBIK (1987): "The Revelation of Information in Strategic Market Games," *Journal of Mathematical Economics*, 16 (2), 105–137. [2596]
- DUTTA, J., AND S. MORRIS (1997): "The Revelation of Information and Self-Fulfilling Beliefs," *Journal of Economic Theory*, 73 (1), 231–244. [2604,2615]
- FOSTER, F. D., AND S. VISWANATHAN (1996): "Strategic Trading When Agents Forecast the Forecasts of Others," *Journal of Finance*, 51 (4), 1437–1478. [2597,2610]
- FUDENBERG, D., AND D. LEVINE (1983): "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory*, 31 (2), 251–268. [2607]

- GEANAKOPOLOS, J. D., AND H. M. POLEMARCHAKIS (1982): "We Can't Disagree Forever," *Journal of Economic Theory*, 28 (1), 192–200. [2596,2598,2603,2604,2615]
- GLICKSBERG, I. L. (1952): "A Further Generalization of the Kakutani Fixed Point Theorem, With Application to Nash Equilibrium Points," *Proceedings of the American Mathematical Society*, 3 (1), 170–174. [2633]
- GOLOSOV, M., G. LORENZONI, AND A. TSYVINSKI (2011): "Decentralized Trading With Private Information," Working Paper, MIT. Available at <http://econ-www.mit.edu/files/4445>. [2596]
- GOOD, I. J. (1952): "Rational Decisions," *Journal of the Royal Statistical Society, Ser. B*, 14 (1), 107–114. [2603]
- GROSSMAN, S. (1976): "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information," *Journal of Finance*, 31 (2), 573–585. [2596]
- HANSON, R. (2003): "Combinatorial Information Market Design," *Information Systems Frontiers*, 5 (1), 107–119. [2598,2599,2602,2603]
- (2007): "Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation," *Journal of Prediction Markets*, 1 (1), 3–15. [2598,2602,2603]
- (2009): "On Market Maker Functions," *Journal of Prediction Markets*, 3 (1), 61–63. [2600]
- HAYEK, F. A. (1945): "The Use of Knowledge in Society," *American Economic Review*, 35 (4), 519–530. [2596]
- HELLWIG, M. F. (1982): "Rational Expectations Equilibrium With Conditioning on Past Prices: A Mean-Variance Example," *Journal of Economic Theory*, 26 (2), 279–312. [2596]
- IYER, K., R. JOHARI, AND C. C. MOALLEMI (2010): "Information Aggregation and Allocative Efficiency in Smooth Markets," Working Paper, Stanford University. Available at <http://ssrn.com/abstract=1581519>. [2616]
- INKLING MARKETS (2011a): "Frequently Asked Questions," available at http://www.inklingmarkets.com/homes/frequently_asked_questions, retrieved September 8, 2011. [2600]
- (2011b): "Inkling Markets Homepage," available at <http://www.inklingmarkets.com>, retrieved September 8, 2011. [2600]
- KREMER, I. (2002): "Information Aggregation in Common Value Auctions," *Econometrica*, 70 (4), 1675–1682. [2596]
- KYLE, A. S. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53 (6), 1315–1336. [2596-2598,2600-2602,2610]
- LAFFONT, J.-J., AND E. S. MASKIN (1990): "The Efficient Market Hypothesis and Insider Trading on the Stock Market," *Journal of Political Economy*, 98 (1), 70–93. [2597,2598]
- LAMBERT, N. S. (2011): "Elicitation and Evaluation of Statistical Forecasts," Working Paper, Stanford University. Available at <http://www.stanford.edu/~nlambert/papers/elicitat.pdf>. [2618]
- MCKELVEY, R. D., AND T. PAGE (1986): "Common Knowledge, Consensus, and Aggregate Information," *Econometrica*, 54 (1), 109–127. [2596,2615]
- (1990): "Public and Private Information: An Experimental Study of Information Pooling," *Econometrica*, 58 (6), 1321–1339. [2603]
- MILGROM, P. R. (1981): "Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica*, 49 (4), 921–943. [2596]
- MILGROM, P., AND N. STOKEY (1982): "Information, Trade and Common Knowledge," *Journal of Economic Theory*, 26 (1), 17–27. [2598]
- NIELSEN, L. T., A. BRANDENBURGER, J. GEANAKOPOLOS, R. MCKELVEY, AND T. PAGE (1990): "Common Knowledge of an Aggregate of Expectations," *Econometrica*, 58 (5), 1235–1239. [2596]
- OSTROVSKY, M. (2012): "Supplement to 'Information Aggregation in Dynamic Markets With Strategic Traders'," *Econometrica Supplemental Material*, 80, http://www.econometricsociety.org/ecta/Supmat/8479_proofs.pdf. [2612]
- PENNOCK, D. M. (2006): "Implementing Hanson's Market Maker," available at <http://blog.oddhead.com/2006/10/30/>, retrieved December 25, 2008. [2599]

- PESENDORFER, W., AND J. M. SWINKELS (1997): "The Loser's Curse and Information Aggregation in Common Value Auctions," *Econometrica*, 65 (6), 1247–1282. [2596]
- RADNER, R. (1979): "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica*, 47 (3), 655–678. [2596,2604]
- RENY, P. J., AND M. PERRY (2006): "Toward a Strategic Foundation for Rational Expectations Equilibrium," *Econometrica*, 74 (5), 1231–1269. [2596]
- SAVAGE, L. J. (1971): "Elicitation of Personal Probabilities and Expectations," *Journal of the American Statistical Association*, 66 (336), 783–801. [2599,2620]
- SEBENIUS, J. K., AND J. GEANAKOPOLOS (1983): "Don't Bet on It: Contingent Agreements With Asymmetric Information," *Journal of the American Statistical Association*, 78 (382), 424–426. [2598]
- SMITH, L., AND P. SØRENSEN (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68 (2), 371–398. [2607]
- WILSON, R. (1977): "A Bidding Model of Perfect Competition," *Review of Economic Studies*, 44 (3), 511–518. [2596]
- WOLINSKY, A. (1990): "Information Revelation in a Market With Pairwise Meetings," *Econometrica*, 58 (1), 1–23. [2596]

Graduate School of Business, Stanford University, Stanford, CA 94305, U.S.A.;
ostrovsky@gsb.stanford.edu.

Manuscript received March, 2009; final revision received September, 2011.