Information Aggregation in Dynamic Markets with Strategic Traders

Michael Ostrovsky*
Stanford GSB and Hoover Institution
ostrovsky@gsb.stanford.edu

December 29, 2008

Abstract

This paper studies information aggregation in dynamic markets with partially informed strategic traders. A natural condition on traded securities and information structure, "separability," is introduced. If a security is separable, information about its value always gets aggregated, for any prior distribution over the states of the world. If the security is non-separable, then there exists a prior such that information does not get aggregated. Special cases satisfying separability include Arrow-Debreu securities, whose value is equal to one in one state of the world and to zero in all others, and "additive" securities, whose value is equal to the sum of traders' signals.

*I am grateful to Yossi Feinberg, Emir Kamenica, Konstantin Makarychev, Yury Makarychev, Sergey Norin, Rahul Sami, Michael Schwarz, and Robert Wilson for helpful comments, suggestions, and discussions.

1 Introduction

The idea that financial markets have the ability to aggregate and reveal dispersed information is an important part of economic thinking. The intuition behind this idea is arbitrage: if the price of a security is wrong, an informed trader will have an incentive to buy or sell this security, thus brining 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, thus revealing all his information. But what happens when there is a small 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 structure is more complicated?

This paper answers these questions for a wide class of securities and information structures, which I call "separable." These securities are traded in a dynamic market where each trader has infinitely many opportunities to trade. In equilibrium, in the limit, information about these securities is aggregated with probability 1.

In the model, there is a finite number of possible states of the world, and the security takes a particular value in each state. There is also a finite number of players, whose information about the states of the world is represented by Aumann's (1976) partition structure. The players share a common prior belief about the probability of each state. The security is called separable under the partition structure if, for *any* prior, either (i) the value of the security is the same in all states of the world that have positive probability under the prior or (ii) there exists a trader and two states of the world that occur with positive probability under the prior such that the trader's expectation of the value of the security conditional on the information that he observes in one state is not equal to his expectation conditional on the information that he observes in the other state. Arrow-Debreu securities (which have payoff 1 in one state of the world and 0 in all other states of the world), additive securities (for which the value of the security is equal to the sum of traders' signals), and monotone transformations of additive securities (e.g., call options on additive securities) are all separable. A simple example of a non-separable security is the following. Suppose there are two players, each tossing a coin and observing only his own toss. If both coins turn up heads or both coins turn up tails, the value of the security is 1. Otherwise, the value of the security is -1. This security is non-separable, because if the coins are fair and the tosses are independent, then both conditions (i) and (ii) above are violated.

There are many ways to model a dynamic market, with or without market makers, with or without noise traders, and so on. In each of them, from the no-trade theorem (Milgrom and Stokey, 1982; Sebenius and Geanakoplos, 1983), we know that if every trader and market maker is fully strategic, shares the same prior, and has no liquidity need for trading, then no profitable trading will take place. There needs to be a "source" of profits, which loses money in expectation, thus stimulating trade and price discovery. In the model used in this paper, this source is an automated market maker. There are no noise traders, and no strategic market makers: the only players are the partially informed traders. Specifically, the models of trading considered in the paper are the market scoring rule of Hanson (2003) and the discounted market scoring rule of Dimitrov and Sami (2008).

Under the market scoring rule, the game starts with a number written on the board—the market maker's initial prediction about the value of the security. The first player can erase that number and write his own prediction. 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 whole cycle repeats indefinitely. The players observe the full history of revisions. After the game is over, the true value of the security is revealed, and each prediction is evaluated according to a strictly proper scoring rule s. The payoff of a player from each revision is the difference between the score of his prediction and the score of the previous player's prediction—in essence, the player "buys out" the previous player's forecast. The total payoff of a player from the game is the sum of payoffs from all revisions, and players are risk-neutral. The discounted market scoring rule is very similar, except that the total payoff of a player is equal to the discounted sum of payoffs from all revisions, where the payoff from a revision made in period t is multiplied by β^t for some $\beta < 1$. I will refer to these games collectively as MSR games.

Note three things about these games. First, the losses of the market maker are deterministically bounded. Second, each player can always guarantee himself the payoff of zero from a revision by not modifying the number on the board. Third, if each player behaves myopically in each period, the number he will put on the board is his posterior belief about the 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).

MSR games look very different from a typical trading process in which securities are bought and sold, but in fact the former can be reinterpreted as the latter (see Hanson, 2003; Pennock, 2006; Chen and Pennock, 2007; and Abramowicz, 2007), with an automated market maker who continuously adjusts the price of the securities as traders buy them from him. Moreover, market scoring rules are increasingly widely used in practice to organize prediction markets (companies like Inkling Markets, Consensus Point, and Xpree, among others, run MSR-based internal prediction markets for a number of large corporations). Finally, the intuition behind the main result is quite general, and is likely to hold in a variety of other market microstructures. The main advantage of using MSR games as a modeling device is that it brings the key issue studied in this paper—aggregation of strategic traders' information—to the forefront. At the same time, it eliminates the need to consider noise traders, strategic market makers, and other aspects of more typical trading games that would divert attention from the key issue.

The main result of the current paper says that if a security is separable, then in the limit as time goes to infinity, the number written on the board will converge to the best estimate of the security's value, given all traders' pooled information. The intuition behind the proof is as follows. 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 the forecasts made by the traders. Consider the stochastic process that corresponds to this observer's vector of posterior beliefs about the likelihoods of the states of the world after each forecast revision. By construction, this process is a bounded martingale, and therefore, by the martingale convergence theorem, converges to some vector-valued random variable Q_∞ with probability 1. If Q_∞ puts positive probabilities 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 number on the board 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 in fact the correct one with probability 1. Now, if the random variable does put all mass on the states with the correct value of the security, but the number on the board does not converge to the same value, then even the uninformed outside observer could have made a profitable revision in infinitely many periods, and thus any informed player could have made such revisions as well, again contradicting the assumption of profit-maximizing behavior. Therefore, the outsider's posterior beliefs, in the limit, have to put all weight on the states with the correct value of the security, and the number on the board has to converge to the same value.

Two recent papers have studied equilibrium behavior of traders in MSR games. Chen et al. (2007) consider undiscounted games based on the particular scoring rule—logarithmic (see Section 2). In their model, the security can take one of two different values, and the number of revisions is finite. They find that if traders' signals are independent conditional on the value of the security, then it is an equilibrium for all traders to behave myopically, i.e., to set the forecast to their posterior belief. They also provide an example of a market in which signals are not conditionally independent and one of the traders has an incentive to behave non-myopically. Dimitrov and Sami (2008) also consider games based on the logarithmic scoring rule. In their models, in contrast to Chen et al., traders observe independent signals. Each realization of the vector of signals corresponds to a particular value of the security. The number of periods is infinite. Dimitrov and Sami find that in that case, in the MSR game with no discounting, myopic behavior is generically not an equilibrium and, moreover, there is no equilibrium in which all uncertainty is guaranteed to get resolved after a finite number of periods. They then introduce a 2-player 2-signal MSR game with discounting, and prove that in that game, information gets aggregated in the limit, under the additional assumption that the "complementarity bound" of the security is positive. They report that based on their experimentation, the bound is not always zero, but do not provide any sufficient conditions for it to be positive.

The rest of the paper is organized as follows. Section 2 presents the formal model. Section 3 states and proves the main result. Section 4 discusses classes of separable and non-separable securities and gives a convenient criterion for separability. Section 5 concludes.

2 Model

There are n players, $i = 1, \ldots, n$. There is a finite set of states of the world, Ω, and a random variable ("security") $X : \Omega \to \mathbb{R}$. 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 of the world is ω, player i observes $\Pi_i(\omega)$). For convenience, I assume that the join (the coarsest common refinement) of partitions $\Pi_1, \ldots, \Pi_n$ consists of singleton sets; i.e., for any two states $\omega_1 \neq \omega_2$ there exists player i such that $\Pi_i(\omega_1) \neq \Pi_i(\omega_2)$. $\Pi = (\Pi_1, \ldots, \Pi_n)$ is the *partition structure*. Players have a common prior distribution P over states in Ω.

At time $t = 0$, nature takes a random draw and selects the state of the world, ω^*, according to P. The uninformed market maker makes the initial prediction $y_0 \in \mathbb{R}$ about the value of X (a natural initial value for y_0 is the unconditional expected value of X under P, but it could also be equal to any other real number). At time $t = 1$, player 1 makes a "revised prediction," y_1. At time $t = 2$, 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 $t = \infty$, with players taking turns revising predictions. All predictions y_t are observed by all players. The action space is bounded, but the bounds are wide enough to allow for any prediction consistent with random variable X, i.e., each y_t is a number in an interval $[\underline{y}, \overline{y}]$, where $\underline{y} \leq \min_{\omega \in \Omega} X(\omega) \leq \max_{\omega \in \Omega} X(\omega) \leq \overline{y}$.

At $t = \infty$, the game ends and 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 the realization of a random variable and y is the 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 our setting), the logarithmic scoring rule ($s(y, x^*) = (x^* - a)\ln(y - a) + (b - x^*)\ln(b - y)$, for some $a < \underline{y}$ and $b > \overline{y}$), due to Good (1952). I assume that $s(y, x^*)$ is continuous and bounded on $[\underline{y}, \overline{y}] \times [\min_{\omega \in \Omega} X(\omega), \max_{\omega \in \Omega} X(\omega)]$.

Under the basic market scoring rule (introduced by Hanson, 2003, though the idea of repeatedly using a proper scoring rule to help forecasters aggregate information goes back to McKelvey and Page, 1990), players get multiple chances to make predictions, and are paid for each revision. Specifically, for each revision of the prediction from y_{t-1} to y_t, player i is paid $s(y_t, x^*) - s(y_{t-1}, x^*)$. Of course, this number can be negative, but each player can guarantee himself a zero payment for a revision by simply setting $y_t = y_{t-1}$, i.e., by not revising the forecast. A slight modification of this payoff rule, introduced by Dimitrov and Sami (2008), is a discounted market scoring rule: it is the same as the basic market scoring rule, except that the payment for the revision from y_{t-1} to y_t is equal to $\beta^t(s(y_t, x^*) - s(y_{t-1}, x^*))$, $0 < \beta \leq 1$. When $\beta = 1$, this rule coincides with the basic market scoring rule. The total payoff of each player is the sum of all payments for revisions. The players are risk-neutral. The resulting game is denoted $\Gamma(\Omega, \Pi, X, y_0, \underline{y}, \overline{y}, s, \beta)$.

Definition 1 *In a Perfect Bayesian Equilibrium of game* Γ, information gets aggregated *if for any* $\epsilon > 0$, *there exists* T *such that for any* $t > T$, *for any realization of the nature's draw* $\omega^* \in \Omega$, *the probability that* $|y_t - X(\omega^*)| > \epsilon$ *on the equilibrium path is less than* ϵ.

2.1 Separability

Consider the following example, which is a formal version of the coin-tossing example in the introduction (both are based on an example from Geanakoplos and Polemarchakis, 1982). There are two players, 1 and 2. $\Omega = \{A, B, C, D\}$; $P(\omega) = \frac{1}{4}$ for every $\omega \in \Omega$; $X(A) = X(D) = 1$ and $X(B) = X(C) = -1$. Partitions are $\Pi_1 = \{\{A, B\}, \{C, D\}\}$ and $\Pi_2 = \{\{A, C\}, \{B, D\}\}$. Other parameters of the game are arbitrary.

In this example, consider the following profiles of strategies and beliefs of the two players: $y_t = 0$ for any $t \geq 1$ and each player's beliefs put equal probabilities on the two possible observations of the other player, after any history. This is a Perfect Bayesian Equilibrium, yet information in this equilibrium does not get aggregated. This happens because even when a player announces his posterior belief truthfully, the other player cannot infer any information from this announcement. I call securities for which this situation is possible non-separable. Formally,

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:*

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

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

Otherwise, security X is separable. Note that non-separable securities are not degenerate—e.g., for any security with payoffs in ϵ-neighborhoods of the ones in the above example, there is a distribution P that would satisfy the conditions in Definition 2, and thus all such securities are non-separable.

3 Main result

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

Theorem 1 *Consider state space Ω, security X, and partition structure Π. If security X is separable under Π, then for any prior distribution P, proper scoring rule s, initial value y_0, bounds $\underline{y}$ and $\overline{y}$, and discount factor $\beta \in (0, 1]$, information in any Perfect Bayesian Equilibrium of the corresponding game Γ gets aggregated. If security X is non-separable under Π, then there exists prior P such that for any s, y_0, $\underline{y}$, $\overline{y}$, and β, there exists a Perfect Bayesian Equilibrium of the corresponding game Γ in which information does not get aggregated.*

Proof. See Appendix. ■

4 Applications

In light of Theorem 1, it is important to understand which securities are separable and which are not, beyond the basic definition. That is the focus of this section.

4.1 Arrow-Debreu securities

Consider first Arrow-Debreu securities, i.e., random variables X that are equal to 1 in one state, say ω_1, and to 0 in all other states. Any such security is separable. To see this, note first that for any Arrow-Debreu security X and any state ω' with $X(\omega') = 0$, there exists at least one player i who knows (in the sense of Aumann, 1976) that the value of X is 0, i.e., for any $\omega'' \in \Pi_i(\omega')$, $X(\omega'') = 0$. Suppose the security is non-separable and there exist distribution P and value v satisfying the bullet points of Definition 2. P has to be positive on at least two different states (otherwise the two bullet points cannot hold together). At least one of these two states, say ω_0, is such that $X(\omega_0) = 0$. Then, at state ω_0 there is a player who knows that the value of security X is 0, and so v has to be equal to 0. But that is only possible if $P(\omega_1) = 0$, contradicting the first bullet point.

Thus, information about Arrow-Debreu securities always gets aggregated. While the analysis of markets with multiple securities is beyond the scope of this paper, this result suggests that in complete markets, information always gets aggregated.

4.2 Additive payoffs

Another important class of separable securities are those with additive payoffs. Specifically, suppose each player's information $\Pi_i(\omega)$ can be interpreted as a signal $x_i(\omega) = x_i(\Pi_i(\omega)) \in \mathbb{R}$, so that the value of security $X(\omega)$ is equal to the sum of players' signals, $X(\omega) = \sum_i x_i$ (this, of course, also includes a seemingly more general case of $X(\omega)$ being a linear function of x_i's, such as the average, because x_i's can be rescaled, or even a stochastically monotone function (Nielsen et al., 1990), for the same reason).

Let us prove that X is separable.[1] Suppose it is not. Then there exist prior P and value v satisfying the bullet points of Definition 2. Consider the unconditional expectation $E[(X(\omega) - v)^2]$ under P.

On the one hand, from the first bullet point of Definition 2, we know that $E[(X(\omega) - v)^2]$ is positive.

On the other hand, since $E[X(\omega)] = v$ and thus $E[(X(\omega) - v)v] = 0$, we have

$$\begin{aligned}
E[(X(\omega)-v)^2] &= E[(X(\omega)-v)X(\omega)] \\
&= E[(X(\omega)-v)\sum_{i=1,\ldots,n} x_i(\omega)] \\
&= \sum_i E[(X(\omega)-v)x_i(\omega)] \\
&= \sum_i \sum_\omega P(\omega)(X(\omega)-v)x_i(\omega) \\
&= \sum_i \sum_{\pi\in\Pi_i} \sum_{\omega\in\pi} P(\omega)(X(\omega)-v)x_i(\omega) \\
&= \sum_i \sum_{\pi\in\Pi_i} \sum_{\omega\in\pi} P(\omega)(X(\omega)-v)x_i(\pi) \\
&= \sum_i \sum_{\pi\in\Pi_i} x_i(\pi)(\sum_{\omega\in\pi} P(\omega)(X(\omega)-v)) = 0.
\end{aligned}$$

The last equality follows from the second bullet point of Definition 2: it implies that for any $\pi \in \Pi_i$, $\sum_{\omega\in\pi} P(\omega)(X(\omega) - v) = 0$. We have arrived at a contradiction.

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

4.3 Increasing payoffs

A generalization of securities with additive payoffs are securities with "increasing" payoffs. Specifically, suppose each player's information $\Pi_i(\omega)$ can be interpreted as a signal $x_i(\omega) = x_i(\Pi_i(\omega)) \in \mathbb{R}$ so that $X(\omega_1) \geq X(\omega_2)$ whenever $x_i(\omega_1) \geq x_i(\omega_2)$ for all players i. If there are two traders in the market, then any such security is separable. To see this, suppose the security is non-separable and take the P and v that verify its non-separability. Take the lowest signal x_1 of player 1 and the corresponding element π_1 of partition Π_1 such that there exists $\omega \in \pi_1$ with $P(\omega) > 0$ and $X(\omega) \neq v$. Among ω in π_1 for which

[1] The proof is similar to the proof of Theorem 3 of Nielsen et al., 1990.

$P(\omega) > 0$), let ω' be the one with the largest corresponding signal of player 2. Then $X(\omega') > v$ and, by construction and monotonicity, all other states in $\Pi_2(\omega')$ that occur with positive probability under P have associated values of X greater than or equal to v, which contradicts the second bullet point of Definition 2.

With three or more players, it is no longer true that any security with increasing payoffs is separable. Indeed, consider the following example. There are three players, 1, 2, and 3. Each player's signal is equal to 1, 2, or 3. If the sum of signals is less than 6, then the value of security X is equal to -1. If the sum is greater than 6, it is equal to 1. If the sum of signals is equal to 6, then the value of security X is equal to -1 if the vector of signals is $(1,2,3)$, $(2,3,1)$, or $(3,1,2)$, to 1 if the vector of signals is $(3,2,1)$, $(2,1,3)$, or $(1,3,2)$, and to 0 if the vector of signals is $(2,2,2)$. It is straightforward to check that value $v = 0$ and prior probability P that places probability $\frac{1}{6}$ on every permutation of $(1,2,3)$ and 0 on all other states satisfy the requirements for non-separability of security X.

4.4 Characterization of separable securities

It is possible to give an alternative characterization of separable securities using duality theory.[2] This characterization is useful in applications, as Corollary 1 illustrates.

Theorem 2 *Security X is separable under partition structure Π if and only if for every $v \in \mathbb{R}$, there exist numbers λ_π corresponding to the elements π of partitions Π_i of all players, such that for every state ω with $X(\omega) \neq v$, the sign of $X(\omega) - v$ is equal to the sign of $\sum_i \lambda_{\Pi_i(\omega)}$, i.e., $(X(\omega) - v)(\sum_i \lambda_{\Pi_i(\omega)}) > 0$.*

Proof. See Appendix. ■

Corollary 1 *Suppose security X can be expressed as $X(\omega) = f\left(x_1(\Pi_1(\omega)) + \cdots + x_n(\Pi_n(\omega))\right)$, where $x_i(\Pi_i(\omega))$ is the "signal" observed by player i in state of the world ω and f is a monotone function. Then X is separable.*

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

Note that the results of Sections 4.1 and 4.2 follow directly from Corollary 1.

5 Conclusion

This paper leaves several important open questions. First, do the games studied in this paper have any Perfect Bayesian Equilibria, in both cases $\beta < 1$ and $\beta = 1$? In the former case, the answer is likely to be "yes" and the proof is likely to be mostly technical. In the latter, the issue is less clear.

A more substantive question is what happens when a security is non-separable, but the prior is generic. E.g., suppose the security and the partition structure are as in the example in Section 2, 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.

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

What happens in the strategic trading game? Does there exist an equilibrium in which information gets aggregated with probability 1? Does there exist an equilibrium in which with positive probability information does not get aggregated, and instead as time goes to infinity, players get "stuck" with not fully informative posterior beliefs that do not give any of them a chance to make a profitable deviation? Are the answers the same for all non-separable securities and various parameters of the game?

There is also a number of interesting questions that go beyond the current paper's model. First, while the intuition behind the proof of the main result seems quite general, it is important to consider other dynamic market microstructures and to check in which of them similar conclusions hold and in which they do not. Second, how important is the assumption that traders are risk-neutral? What happens if they are not? What if their coefficients of risk aversion are different? Finally, what happens if information is costly? How well is information extracted and aggregated in that case?

References

[1] Abramowicz, Michael (2007). "The Hidden Beauty of the Quadratic Market Scoring Rule: A Uniform Liquidity Market Maker, with Variations," *Journal of Prediction Markets*, 1(2), pp. 111–125.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

A Proof of Theorem 1

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

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

We first make an auxiliary observation. For $\delta \in [0, \frac{1}{2}]$, let $R(\delta)$ be the set of probability distributions r such that there are states ω_j and ω_k with $r(\omega_j) \geq \delta$, $r(\omega_k) \geq \delta$, and $X(\omega_j) \neq X(\omega_k)$.

Lemma 1 *If security X is separable, then $\forall \delta > 0, r \in R(\delta)\ \exists \phi > 0, i \in \{1, 2, \dots, n\}$ such that for any $z \in [\underline{y}, \overline{y}]$, the instant opportunity of player i given r and z is greater than ϕ.*

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

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

Of course, $Q = \{(q_t^1, \dots, q_t^{|\Omega|})\}_{t=0,1,\dots}$ is also a stochastic process in $\mathbb{R}^{|\Omega|}$. The key idea of the proof is that this process is a martingale, by the law of iterated expectations. And by the martingale convergence theorem, it has to converge to a random variable, $Q_\infty = (q_\infty^1, \dots, q_\infty^{|\Omega|})$.[3]

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

Case 1

Suppose there is a positive probability that Q_∞ assigns positive likelihoods to two states ω_j and ω_k with $X(\omega_j) \neq X(\omega_k)$. This implies that there is a vector of posterior probabilities $r = (r^1, \dots, r^{|\Omega|})$ and a $\delta > 0$ such that $r^j > \delta$, $r^k > \delta$, and for any $\epsilon > 0$, the probability that Q_∞ is in an ϵ-neighborhood of r is positive. Since Q_t converges to Q_∞, for any $\epsilon > 0$, there exists time T and $\zeta > 0$ such that for any $t > T$, the probability that Q_t is in the ϵ-neighborhood of r is greater than ζ.

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

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

Therefore, for some player i, time T, and $\eta > 0$, the expected (over all realizations of stochastic process Q) instant opportunity of player i in any period $nk + i > T$ is greater than η.

Case 2

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

Suppose now that with some positive probability, process y_t does not converge to the true value of the security. That is, there exist state ω and $\epsilon > 0$ such that after state ω is drawn by nature, for any T, there exists $t > T$ such that $Prob(|y_t - X(\omega)| > \epsilon) > \epsilon$. This, together with the fact that even for the uninformed outside observer 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 T, there exists period $nk + i > T$ in 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 periods $nk + i^*$ in which the expected instant opportunity of player i^* is greater than η^*. Fix i^* and η^*.

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

The rest of the proof is split into two parts, depending on the value of parameter β: $\beta < 1$ and $\beta = 1$.

Part "$\beta < 1$"

Take any period t. Let Ψ_t be the sum of all players' expected payoffs from the revision made in periods t and later, divided by β^t: $\Psi_t = (S_t - S_{t-1}) + \beta(S_{t+1} - S_t) + \beta^2(S_{t+2} - S_{t+1}) + \dots$. We can make two observations about Ψ_t. First, it is non-negative, because each player can guarantee himself a payoff of zero. Second, for a similar reason, it is greater than or equal to the expected instant opportunity of the player who makes the forecast at time t. Consider now $\lim_{t\to\infty} \sum_{t=1}^{T} \Psi_t$.

On the one hand, under both Case 1 and Case 2, this limit has to be infinite, because each term Ψ_t is non-negative, and an infinite number of them are greater than η^*.

On the other hand, for any T, $\sum_{t=1}^{T} \Psi_t =$

$$
\begin{array}{lcccccccc}
 & (S_1 - S_0) & + & \beta(S_2 - S_1) & + & \beta^2(S_3 - S_2) & + & \dots \\
+ & & & & & & & \\
 & (S_2 - S_1) & + & \beta(S_3 - S_2) & + & \beta^2(S_4 - S_3) & + & \dots \\
+ & & & & & & & \\
 & (S_3 - S_2) & + & \beta(S_4 - S_3) & + & \beta^2(S_5 - S_4) & + & \dots \\
+ & & & & & & & \\
 & \vdots & & & & & & \\
+ & & & & & & & \\
 & (S_T - S_{T-1}) & + & \beta(S_{T+1} - S_T) & + & \beta^2(S_{T+2} - S_{T+1}) & + & \dots
\end{array}
$$

$= \sum_{t=0}^{\infty} \beta^t (S_{t+T} - S_t) < \frac{2M}{1-\beta}$, where M is some number greater than the highest possible absolute value of the score function, s, given the possible range of outcomes $[\min_{\omega\in\Omega} X(\omega), \max_{\omega\in\Omega} X(\omega)]$ and the allowed range of predictions $[\underline{y}, \overline{y}]$. Thus, both Cases 1 and 2 are impossible, and thus y_t has to converge to the true value of security X.

Part "$\beta = 1$"

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

B Proof of Theorem 2

Notice first that strengthening the first bullet point in Definition 2 to "$P(\omega) = 0$ for every state ω such that $X(\omega) = v$" results in an equivalent definition. To see this, take a non-separable security X, take P and v satisfying the existing definition, set $P(\omega)$ to zero for all ω such that $X(\omega) = v$, and rescale the probabilities of remaining states so that they add up to one. The second bullet point is unaffected, so the new definition is also satisfied.

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

By Gordan's Transposition Theorem[4], exactly one of the following two systems of equations and inequalities has a solution:

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

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

Thus, either

- for some v, system (1) has a solution, and so the security is non-separable (Take solution x of (1); rescale it so that its elements add up to 1; and use rescaled probabilities as the common prior P. Then $AP = 0$, which implies the second bullet point of the definition of non-separability.)

  or

- for every v, system (2) has a solution.

Also, tautologically, every security is either separable or non-separable. Therefore, every security is separable if and only if for every v, system (2) has a solution. That is exactly the statement of Theorem 2.

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