Information Aggregation in Dynamic Markets with Strategic Traders

Michael Ostrovsky

Setup

- n risk-neutral players, $i = 1, \ldots, n$
- Finite set of states of the world Ω
- Random variable ("security") $X : \Omega \to \mathbb{R}$
- Each player i receives information about the true state $\omega \in \Omega$ according to partition Π_i of Ω
- The join of partitions $\Pi_1, \ldots, \Pi_n$ consists of singleton sets
- $\Pi = (\Pi_1, \ldots, \Pi_n)$ is the *partition structure*.
- Players have a common prior distribution P over states in Ω.

Trading – two models

- model of trading based on Kyle (1985);
- model of trading based on the Market Scoring Rule (MSR) of Hanson (2003).

Model of trading based on Kyle (1985)

- 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. There is also demand u_1 from noise traders, drawn randomly from $N(0, t_1 - t_0 = \frac{1}{2})$.

  - Competitive market makers observe aggregate demand $\sum_i d_1^i + u_1$, form their posterior beliefs about the state of the world, and set market price y_1 equal to the expected value of the security conditional on their beliefs.

  - The market clears, and all traders observe price y_1 and aggregate volume $\sum_i d_1^i + u_1$.

- At time $t_2 = \frac{3}{4}$, the next auction takes place, with each strategic player i choosing demand d_2^i and demand from noise traders u_2 drawn randomly from $N(0, t_2 - t_1 = \frac{1}{4})$.

- Subsequently, auctions are held at times $t_k = 1 - \frac{1}{2^k}$ with demand from noise traders drawn from $N(0, \frac{1}{2^k})$.

- The true value of the security, $x^* = X(\omega^*)$, is revealed at some time $t^* > 1$.

- Player i's payoff is equal to $\sum_{k=1}^{\infty} d_k^i \, (x^* - y_k)$.

Market Scoring Rules (Hanson, 2003; Dimitrov and Sami, 2008)

- Start with prediction y_0 offered by the market sponsor.
- Players take turns making predictions $y_k \in [\underline{y}, \overline{y}]$ at times t_1, t_2, t_3, ... in $(0, 1)$; sequence t_k converges to 1.
- At time $t^* > 1$, the value x^* of security X is revealed.
- For each revision of the prediction from y_{k-1} to y_k, player i is paid $s(y_k, x^*) - s(y_{k-1}, x^*)$.
- Discounted MSR: for each revision of the prediction from y_{k-1} to y_k, player i is paid $\beta^k(s(y_k, x^*) - s(y_{k-1}, x^*))$.
- The total payoff of each player is the sum of all payments for revisions.

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

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

Example (based on Geanakoplos and Polemarchakis, 1982)

- Two players, 1 and 2

- $\Omega = \{A, B, C, D\}$

- $X(A) = X(D) = 1$ and $X(B) = X(C) = -1$

- $\Pi_1 = \{\{A, B\}, \{C, D\}\}$ and $\Pi_2 = \{\{A, C\}, \{B, D\}\}$

If the players' common prior P assigns probability $\frac{1}{4}$ to every state, then for every ω it is common knowledge that each player's posterior belief about the value of the security is 0.

Dutta-Morris (1997) and DeMarzo-Skiadas (1998, 1999) use similar examples to illustrate the generic existence of not fully informative REE. D-S also define "separable orientation":

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

1. *$P(\omega)$ is positive on at least one state ω in which $X(\omega) \neq v$;*

2. *For every player i and every state ω with $P(\omega) > 0$,*

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

Otherwise, security X is separable.

Separability

- If $n = 1$, every security is separable
- Arrow-Debreu securities are separable
- Securities with additive payoffs are separable
- Securities that are order statistics (min, max, median, etc.) of players' signals are separable
- Monotone transformations of additive and multiplicative securities (e.g., call options on those securities) are separable
- Securities with payoffs increasing in signals are separable if $n = 2$ and may be non-separable if $n > 2$

Theorem. *Consider n, Ω, X, and Π.*

1. *If X is separable under Π, then for any prior P:*
   - *in any PBE of the corresponding game Γ^K information gets aggregated;*
   - *for any strictly proper scoring rule s, initial value y_0, bounds $\underline{y}$ and $\overline{y}$, and discount factor $\beta \in (0, 1]$, in any PBE of game Γ^{MSR} information gets aggregated.*

2. *If X is non-separable Π, then there exists prior P such that:*
   - *there exists a PBE of the corresponding game Γ^K in which information does not get aggregated;*
   - *for any s, y_0, $\underline{y}$, $\overline{y}$, and β, there exists a PBE of game Γ^{MSR} in which information does not get aggregated*

Proof of Statement 1 for game Γ^{MSR}

Pick any PBE and consider the following stochastic process Q. $Q_0 = (q_0^1, \ldots, q_0^{|\Omega|})$, where $q_0^w = P(\omega_w)$. Nature draws state ω according to distribution P and each player i observes $\Pi_i(\omega)$. Then, player 1 plays according to his equilibrium strategy and makes forecast y_1. Based on y_1, the 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 ω, forms posterior beliefs about the probability of each state ω_w. Denote this probability by q_1^w. $Q_1 = (q_1^1, \ldots, q_1^{|\Omega|})$. The rest of the process is constructed analogously: $Q_k = (q_k^1, \ldots, q_k^{|\Omega|})$, where q_k^w is the posterior belief of the observer about the probability of state ω_w after time t_k.

The key idea of the proof is that this process is a martingale. By the martingale convergence theorem, it has to converge to a random variable, $Q_\infty = (q_\infty^1, \ldots, q_\infty^{|\Omega|})$. We will show that Q_∞ has to place all weight on the states with the correct value of the security, and y_k has to converge to that value as well.

Let $r = (r^1, r^2, \ldots, r^{|\Omega|})$ be any probability distribution over the states and let z be any real number. Define *instant opportunity* of player i given r and z as his highest possible expected payoff from making only one change to the forecast, if the state is drawn according to r and the initial forecast is z, i.e.,

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

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

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

Now, suppose the statement of the theorem does not hold for this equilibrium. Consider Q_∞ and two possible cases.

Case 1

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

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

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

Case 2

Now suppose there is zero probability that Q_∞ assigns positive likelihoods to two states ω_a and ω_b with $X(\omega_a) \neq X(\omega_b)$. Then, for every realization ω of the 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.

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

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

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

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

Part "$\beta < 1$"

Let $\Psi_k = (S_k - S_{k-1}) + \beta(S_{k+1} - S_k) + \beta^2(S_{k+2} - S_{k+1}) + \ldots$. Then (i) $\Psi_k \geq 0$ and (ii) it is greater than or equal to the expected instant opportunity of the player who makes the forecast at t_k.

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

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

$= \sum_{k=0}^{\infty} \beta^k (S_{k+K} - S_k) < \frac{2M}{1-\beta}$ for some M.

Part "$\beta = 1$"

Take any player i. His expected payoff is equal to

$$\sum_{j=1}^{\infty} (S_{i+nj} - S_{i+nj-1}).$$

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

Open Questions

- Within the current model
  - Existence; slight variations (discretization, finite games, etc.) – next version of the paper (I hope)
  - For non-separable securities, under a generic prior, one can show that price converges to a “common knowledge/common belief” equilibrium of Dutta-Morris and DeMarzo-Skiadas. Can we find out to which one?
    * Are there multiple equilibria, under which information convergences to different points?
    * Is this question easier to answer in a continuous-time model?

- Beyond the current model: open-ended questions
  - Other dynamic microstructures
  - Multiple securities
  - Risk-averse traders (with different utility functions)
  - Costly information
  - Information changes over time (e.g., “Insider Trading with a Random Deadline” by Caldentey and Stacchetti)