Information Aggregation in Prediction Markets

Michael Ostrovsky

Do markets aggregate and reveal information?

- Hayek (1945), Grossman (1976), Radner (1979)
  - Yes; no strategic foundations.

- Hellwig (1982), Dubey, Geanakoplos, and Shubik (1987), Wolinsky (1990), Golosov, Lorenzoni, and Tsyvinski (2008)
  - Mostly yes; dynamic trading with myopic traders.

- Wilson (1977), Milgrom (1981), Pesendorfer and Swinkels (1997), Kremer (2002), Reny and Perry (2006)
  - Mostly yes; in auctions with many traders, a suitably chosen number of objects, and particular functional forms.

This talk

For a broad class of securities and information structures, information in dynamic markets with partially informed strategic traders always gets aggregated: in every equilibrium, as time approaches the end of the trading interval, the market price of a security converges in probability to its expected value conditional on the traders' pooled information.

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
- $\Pi = (\Pi_1, \ldots, \Pi_n)$ is the *partition structure*
- The join of partitions $\Pi_1, \ldots, \Pi_n$ consists of singleton sets
- Players have a common prior distribution P over states in Ω

Trading

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

(Note: the paper “Information aggregation in dynamic markets with strategic traders” also contains a model of trading based on Kyle (1985)).

Detour: Proper Scoring Rules

Suppose a risk-neutral individual has some information about the value of random variable X, is asked to make a prediction, y, about the value, and subsequently receives $s(y, x^*)$ when the realization x^* of X is observed.

Function $s(\cdot, \cdot)$ is a *proper scoring rule* if it is optimal for the individual to report his expectation of X truthfully.

Examples

- quadratic scoring rule (Brier, 1950):

$$s(y, x^*) = -(x^* - y)^2$$

- logarithmic scoring rule (Good, 1952):

$$s(y, x^*) = (x^* - a) \ln(y - a) + (b - x^*) \ln(b - y).$$

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 an equilibrium of 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 $\epsilon > 0$, 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 ϵ.

Example of non-aggregation (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 exists distribution P on Ω such that for some v:*

1. *for every player i, state ω with $P(\omega) > 0$, $E[X|\Pi_i(\omega)] = v$;*
2. *for some ω with $P(\omega) > 0$, $X(\omega) \neq v$.*

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

Theorem. *Consider n, Ω, X, and Π.*

1. *If X is separable under Π, then for any prior P, 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 equilibrium of game Γ^{MSR} information gets aggregated.*

2. *If X is non-separable under Π, then for some prior P, 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

Pick any equilibrium and consider the following stochastic process Q. Nature draws ω according to P. Each player i observes $\Pi_i(\omega)$.

- $Q_0 = (q_0^1, \ldots, q_0^{|\Omega|})$, where $q_0^w = P(\omega_w)$.

Next, player 1 makes forecast y_1. Based on y_1, the strategy of player 1, and the prior P, a Bayesian outside observer forms posterior beliefs about the probability q_1^w of each state ω_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.

Key observation: $(Q_k)_{k=1,\ldots}$ is a $|\Omega|$-dimensional martingale.

By the martingale convergence theorem, it has to converge to a random variable, $Q_\infty = (q_\infty^1, \ldots, q_\infty^{|\Omega|})$.

Rest of the proof: 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.,

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

Lemma. *Take any separable security X and distribution r such that $Var_r[X] > 0$. 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: 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: there is a 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 $\epsilon > 0$ such that after state ω is drawn by nature, for any K, there exists $k > K$ such that $Prob(|y_k - X(\omega)| > \epsilon) > \epsilon$. 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 "$\beta = 1$"

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^*.

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) $\Psi_k \geq$ the expected instant opportunity of the player who makes the forecast at t_k.

Consider now $\lim_{K\to\infty} \sum_{k=1}^{K} \Psi_k$. This limit is infinite, because each term Ψ_k is non-negative, and an infinite number of them are greater than η^*. At the same time, $\sum_{k=1}^{K} \Psi_k =$

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

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

Open Questions

- Within the current model
  - For non-separable securities, under a generic prior, price converges to a "common knowledge/belief" equilibrium of Dutta-Morris and DeMarzo-Skiadas—but to which one?
- Beyond the current model
  - Other dynamic microstructures (+1 in the paper . . .)
  - Risk-averse traders (with different utility functions)
  - Costly trading or information acquisition
  - Dynamic uncertainty

Key idea of the proof

Consider a $|\Omega|$-dimensional stochastic process Q tracking the evolution of an outside observer's beliefs about the probabilities of the true state ω.

Q is a bounded martingale, and therefore converges to some random variable Q_∞.

If there is a positive chance that Q_∞ does not place all weight on the states with the correct value of the security, then by separability, at least one trader has a non-vanishing arbitrage opportunity in the limit — and thus in any sufficiently late period.

This, in turn, implies that this trader never actually takes advantage of this opportunity, contradicting the assumption of profit-maximizing behavior.