Information Aggregation in Dynamic Markets with Strategic Risk Averse Traders

Vasileios Nastoulis & Timos Athanasiou

Twenty-fourth Conference on Research on Economic Theory and Econometrics

Ios | July, 2026

Athens University of Economics and Business

Presented by Vasileios Nastoulis

Roadmap

1. Motivation and benchmark
2. From quantity trading to public price revisions
3. Information aggregation and separability
4. Results and proof sketch
5. Speed of price discovery

Motivation

Do prices aggregate information when informed traders are risk-averse?

Prices are public signals: they summarize what the market has learned from past trades.

In actual trading environments, informed traders choose quantities, hence they affect public beliefs via prices:

$$\text{orders move prices} \Longrightarrow \text{prices update beliefs} \Longrightarrow \text{updated beliefs generate new trading incentives.}$$

But larger positions also expose traders to more risk.

Does Ostrovsky's **separability** remain the key informational condition, for information aggregation in a market where:

1. traders are strategic and partially informed;
2. the market maker commits to a linear price schedule;
3. information is revealed through quantity orders;
4. risk aversion and price impact jointly generate inertia.

Main result

Public prices still converge to the true payoff, albeit at a slower pace.

Benchmark: Ostrovsky's aggregation mechanism

Ostrovsky studies a dynamic public-price environment – similar to ours – with finitely many strategic traders and dispersed private information.

The benchmark result is:

$$\text{separability} \quad \Longrightarrow \quad \text{information aggregation in every equilibrium.}$$

In that framework, the public price history is the market's learning device. Traders do not need to observe each other's private signals directly.

They observe the sequence of public prices and revise their beliefs accordingly.

We keep

The same finite-state space, private partitions, common prior, learning from public price history, and separability condition.

We change

Price revisions are not direct announcements. They are generated by quantity trades through a market maker's linear price schedule.

Trading environment

In each round, one active trader submits an order q_k for a risky security with terminal payoff X.

The market maker executes the order through a known linear price schedule:

$$y_k = y_{k-1} + \Lambda q_k, \qquad \Lambda > 0.$$

Hence quantities and public prices are linked one-to-one:

$$q_k = \frac{y_k - y_{k-1}}{\Lambda}.$$

The trader's per-round payoff is

$$u_i(q_k, m_k) = X q_k - \frac{\delta}{2} q_k^2 + m_k. \qquad (*)$$

The parameter δ captures the trader's reluctance to take large risky positions.

From quantities to prices

A quantity choice can be rewritten as a public price revision.

Information and public learning

There is a finite state space Ω, a common prior μ, and a security payoff

$$X : \Omega \to \mathcal{Y} = [\underline{y}, \overline{y}].$$

Each trader observes only a private partition element:

$$\Pi_i(\omega) \in \Pi_i.$$

Private signals are never publicly revealed.

The market observes only the public price history:

$$H_{k-1} = (y_0, y_1, \ldots, y_{k-1}).$$

In round k, trader $i(k)$ is active. Given his private information and the public price history, he forms

$$Y_k = \mathbb{E}_\mu \left[X \mid \Pi_{i(k)}(\omega), H_{k-1} \right].$$

Private signals are not observed $\implies$ the market learns from the sequence of public prices $\implies$ each price revision reflects the posterior mean.

Connecting strategic trading and information aggregation

Kyle-style strategic trading

- Traders understand that their orders affect prices.
- Large orders may exploit information more aggressively, but they also move prices against the trader.
- In Kyle, price impact is an equilibrium object.
- Here, Λ is a known feature of the market maker's linear pricing rule.

Ostrovsky's public-price framework

- Public price history summarizes what the market learns.
- Strategic behavior concerns how much information is incorporated into the next price.
- In Ostrovsky, the scoring rule is assumed at the outset.
- Here, the semi-score is derived from the quantity-trading problem.

Kyle-style quantity trading + Ostrovsky-style public learning $\Longrightarrow$ semi-score price revision.

Lemma 1: from quantity choice to score maximization

In any trading round k, let the starting (previous) price to be $z = y_{k-1}$. Given the mover's information $(\Pi_{i(k)}, H_{k-1})$, his posterior mean is $Y_k = \mathbb{E}_\mu\big[X \mid \Pi_{i(k)}, H_{k-1}\big]$.

(i) the $k-$round trading problem with linear impact $y_k = z + \Lambda q_k$ and preferences as in $(*)$ solved by i is

$$\max_{q_k \in \mathcal{I}_z^{-1}(\mathcal{Y})} \Big\{ \big(\mathbb{E}_\mu\big[X \mid \Pi_{i(k)}, H_{k-1}\big] - z\big) q_k - \frac{\delta + \Lambda}{2} q_k^2 \Big\} \tag{2}$$

(ii) because quantities and public prices are linked one-to-one, **equivalently** we can write the problem as if the trader chooses the next public price y_k evaluated by

$$s^z(y_k, X) = -(1-\alpha)(X - y_k)^2 - \alpha(y_k - z)^2 \quad \alpha := \frac{\delta}{\delta + \Lambda} \tag{3}$$

Then the trader chooses y_k to maximize the expected semi-score,

$$\max_{y_k \in \mathcal{Y}} \mathbb{E}_\mu\big[s^z(y_k, X) \mid \Pi_{i(k)}, H_{k-1}\big] \tag{4}$$

(iii) Both problems, (2) and (4), lead to the same optimal public price that is unique:

$$y_k^* = \underbrace{\alpha}_{\text{inertia}} z + \underbrace{(1-\alpha)}_{\textit{responsiveness}} \mathbb{E}_\mu\big[X \mid \Pi_{i(k)}, H_{k-1}\big]. \tag{B.Rs.}$$

The semi-score is not an ad hoc payoff. It is the price-revision representation of the trader's quantity problem.

Separability: the key informational condition

Definition 1

The security X is **non-separable** under partition structure $\Pi = (\Pi_1, \ldots, \Pi_n)$ if there exist distribution μ' on the underlying state space Ω and a value $v \in \mathbb{R}$ such that:

1. $\mu'(\omega) > 0$ on at least one state ω in which $X(\omega) \neq v$;
2. for every trader i and every ω with $\mu'(\omega) > 0$,

$$\mathbb{E}_{\mu'}[X \mid \Pi_i(\omega)] = v.$$

Otherwise, X is **separable**.

Intuition

1. Pooled information may identify the state. However, this is not sufficient for the market to aggregate the information.
2. Separability ensures that some trader's posterior mean can reveal payoff-relevant information.

Non-Separable vs Separable security: examples

Non-separable security

$$X = (0, 1, 2, 3), \qquad \mu = \left(\frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8}\right) \quad y_0 = \frac{3}{2}, \quad \alpha = \frac{1}{2}.$$

$$\Pi_1 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}, \quad \Pi_2 = \{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}\}.$$

The join is singleton: pooled signals identify the state.

$$\Pi_1 \vee \Pi_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4\}\}.$$

But for every state ω,

$$\mathbb{E}_\mu[X \mid \Pi_1(\omega)] = \mathbb{E}_\mu[X \mid \Pi_2(\omega)] = \frac{3}{2}.$$

If $\omega = \omega_1$, then $X(\omega_1) = 0$.

$$Y_1 = Y_2 = \frac{3}{2}, \qquad y_1 = y_2 = \frac{3}{2}.$$

Every trader's conditional mean collapses to the same value v although the payoff is still uncertain.

Separable security

$$X = (0, 3, 1, 4), \quad \mu = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right), \quad y_0 = 2, \quad \alpha = \frac{1}{2}.$$

$$\Pi_1 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}, \quad \Pi_2 = \{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}\}.$$

Pooled signals identify the state (again!)

$$\Pi_1 \vee \Pi_2 = \{\{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_4\}\}.$$

If ω_1 is realized, Trader 1 observes $\{\omega_1, \omega_3\}$, so

$$Y_1 = \mathbb{E}_\mu[X \mid \{\omega_1, \omega_3\}] = \frac{1}{2}, \qquad y_1 = \frac{1}{2}y_0 + \frac{1}{2}Y_1 = \frac{5}{4}.$$

Trader 2 observes and combines y_1 with $\{\omega_1, \omega_4\}$, identifies ω_1, and

$$Y_2 = 0, \qquad y_2 = \frac{1}{2}y_1 = \frac{5}{8}.$$

After this, $Y_k = 0$ for all later rounds, so the public price path is

$$2 \to \frac{5}{4} \to \frac{5}{8} \to \frac{5}{16} \to \frac{5}{32} \to \cdots \to 0.$$

The difference is whether payoff-relevant information can pass through individual posterior means and public prices.

Main result: information aggregation in every equilibrium

Theorem (1)

Consider the game $\Gamma_{\mathrm{SSR}}(\Omega, \Pi, X, \mathcal{Y}, \alpha, y_0, s^z, \beta)$ *with initial price* $y_0 \in \mathcal{Y}$*, a discount factor* $\beta \in (0, 1]$

- *If* X *is separable under* Π*, then for every prior* μ *in any Nash equilibrium of* Γ_{SSR} *information gets aggregated.*
- *If* X *is non-separable under* Π*, then there exist a prior* $\widehat{\mu} \in \Delta(\Omega)$ *and a Nash equilibrium where the price of the security does not aggregate information.*

Theorem 1 Application: Equilibrium price path

Fix $\alpha = \frac{1}{2}$, $n = 3$, uniform prior $\mu(\omega) = 1/6$, and $X = \left(1, \frac{1}{2}, 1, \frac{1}{2}, 1, 0\right)$. The traders' partitions are

$$\Pi_1 = \big\{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}, \{\omega_5, \omega_6\}\big\}, \quad \Pi_2 = \big\{\{\omega_1, \omega_2, \omega_3\}, \{\omega_4, \omega_5, \omega_6\}\big\},$$

$$\Pi_3 = \big\{\{\omega_1, \omega_3, \omega_5\}, \{\omega_2, \omega_4, \omega_6\}\big\}.$$

Suppose ω_1 is realized, $y_0 = \mathbb{E}_\mu[X] = \frac{2}{3}$ and price updates: $y_k = \alpha y_{k-1} + (1-\alpha)\mathbb{E}_\mu[X \mid \Pi_{i(k)}, \mathcal{F}_{k-1}]$.

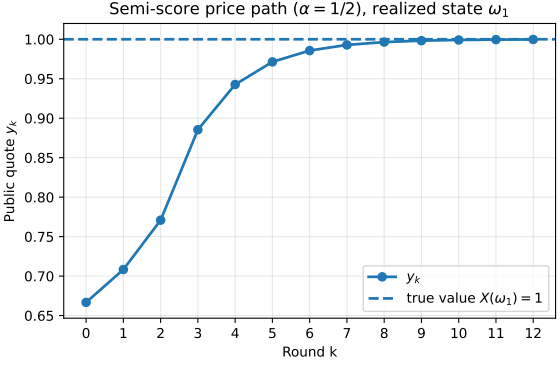

Evolution of beliefs along the equilibrium path

Starting from the uniform prior, price observations progressively eliminate states. After round 3, the public belief is

$$Q_3 = \left(\frac{1}{2}, 0, \frac{1}{2}, 0, 0, 0\right).$$

Though, the public has not learned whether the true state is ω_1 or ω_3, both states give the same payoff,

$$X(\omega_1) = X(\omega_3) = 1.$$

The public has learned the payoff-relevant information.

This is why the theorem is about convergence of prices to X, not necessarily full state revelation.

Proof sketch I: failed aggregation generates revision opportunities

Fix any Nash equilibrium and let Q_k be the public posterior belief after round k.

1. **Public belief convergence.** The public posterior process Q_k is a bounded martingale, so
$$Q_k \to Q_\infty.$$
2. **Two possible failures.** If aggregation fails, then either:
   - **Payoff-relevant uncertainty remains:**
   $$\exists \omega_b, \omega_c \in \Omega : \quad Q_\infty(\omega_b) > 0, \quad Q_\infty(\omega_c) > 0, \quad X(\omega_b) \neq X(\omega_c).$$
   - **The payoff is learned, but prices do not converge:**
   $$Q_\infty\big(\{\omega' \in \Omega : X(\omega') = X(\omega)\}\big) = 1, \quad \text{but} \quad y_k \overset{\mathbb{P}^\mu}{\not\longrightarrow} X.$$
3. **Separability implies a positive revision opportunity.** In both cases there exists some trader that has a strictly positive one-step revision opportunity.
4. **Therefore:**
   failure of aggregation $\Longrightarrow$ recurring positive one-step revision opportunities.

Proof sketch II: Recurring positive revision opportunities are impossible in equilibrium

Let $\bar{G}_k$ be the expected score gain induced by the equilibrium price revision in round k, and

$$\bar{\ell}_k = \mathbb{E}_\mu[(X - y_k)^2] \text{ be expected squared pricing error.}$$

Equilibrium score gains are bounded above by reductions in expected squared pricing error.

$$\bar{G}_k \leq (1-\alpha)(\bar{\ell}_{k-1} - \bar{\ell}_k).$$

1. **Discounted case $\beta < 1$.** The aggregate discounted continuation payoffs are uniformly bounded:

$$\sum_{k=1}^{K} U_k = \sum_{k=1}^{K} \sum_{m=0}^{\infty} \beta^m \bar{G}_{k+m} \leq \frac{1-\alpha}{1-\beta} d_Y^2.$$

   Recurring opportunities $> \eta^*$ would make these partial sums diverge.

2. **Undiscounted case $\beta = 1$.** Late continuation payoffs become arbitrarily small. If opportunities $> \eta^*$ arise arbitrarily late, a trader can choose a date at which his remaining equilibrium payoff is below $\eta^*/2$, take the one-step opportunity, and then stop revising.

3. **Conclusion.** Recurring positive revision opportunities are impossible in equilibrium in both cases.

Therefore, aggregation cannot fail: $y_k \xrightarrow{\mathbb{P}_\mu} X$.

Proposition 1: speed of public price discovery

Given that $0 < \delta < \infty$ and $0 < \Lambda < \infty$, Proposition 1 asks how quickly the public price moves toward the true payoff over a finite window

$$\{K+1, \ldots, K+m\}.$$

Public prices follow the optimal updating rule of public prices: $y_k = \alpha y_{k-1} + (1-\alpha) Y_k,\ \alpha = \frac{\delta}{\delta+\Lambda}$.
For every starting round K and horizon $m \geq 1$,

$$|y_{K+m} - X| \leq \underbrace{\alpha^m d_{\mathcal{Y}}}_{\text{inherited price inertia}} + \underbrace{\max_{1 \leq j \leq m} |Y_{K+j} - X|}_{\text{target inaccuracy}}.$$

Thus finite-window price error has two sources:

$$\text{past-price persistence} \quad + \quad \text{posterior-target inaccuracy}.$$

Define

$$m_\varepsilon(\alpha) := \min\{m \geq 1 : \alpha^m d_{\mathcal{Y}} \leq \varepsilon\}.$$

A larger α means a larger learning horizon $m_\varepsilon(\alpha)$.

Based on definition of α, higher risk aversion δ slows price discovery, while higher price impact Λ speeds up public price adjustment.

Comparative statics

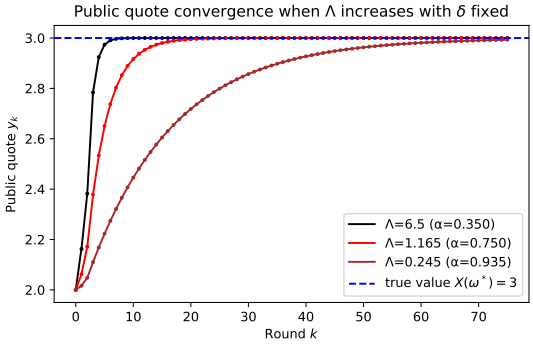

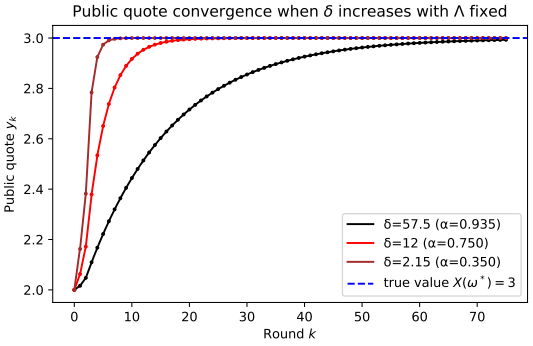

Figure: Variation in price impact and risk aversion. The left panel varies Λ (holding δ fixed), and the right panel varies δ (holding Λ fixed).

The parameters do not change the limiting value of the public price. They change the speed at which the public price approaches that value.

Limits of separability: hidden responsiveness

In the baseline model, the inertia parameter $\boxed{\alpha := \delta/(\delta + \Lambda)}$ and responsiveness parameter $\boxed{\rho := 1 - \alpha}$ are both common knowledge among the agents.

Because α is known, the public can recover the mover's posterior target by reverting: $y_k = \alpha y_{k-1} + (1 - \alpha) Y_k$, i.e.

$$Y_k = \frac{y_k - \alpha y_{k-1}}{1 - \alpha} = \frac{y_k - \alpha y_{k-1}}{\rho}.$$

If trader i has known heterogeneity, $\alpha_i = \frac{\delta_i}{\delta_i + \Lambda}$, $\rho_i := 1 - \alpha_i = \frac{\Lambda}{\delta_i + \Lambda}$, then, when trader i moves,

$$y_k = \alpha_i y_{k-1} + \rho_i Y_k, \qquad Y_k = \frac{y_k - \alpha_i y_{k-1}}{\rho_i}.$$

The difficulty begins when α_i — equivalently ρ_i — is privately known. Then the same public price revision may reflect

$$\text{information about } X \quad + \quad \text{information about hidden responsiveness.}$$

Separability remains a property of X and Π, but public price history is no longer a clean signal of private information.

Open question: existence of Nash equilibrium

The aggregation theorem is conditional on equilibrium existence:

$$X \text{ separable} \quad \Longrightarrow \quad y_k \xrightarrow{\mathbb{P}_\mu} X \quad \text{in every Nash equilibrium.}$$

The per-round problem is well behaved: the semi-score objective is strictly concave and the best reply is unique,

$$y_k^* = \alpha y_{k-1} + (1-\alpha)\mathbb{E}_\mu[X \mid \Pi_{i(k)}, H_{k-1}].$$

The difficulty is dynamic: strategies depend on continuous-valued public price histories. Hence small changes in histories may induce large changes in future actions, so continuity of continuation payoffs is not automatic.

Next step. Establish Nash existence for the finite-horizon semi-score game, or identify a restricted strategy class under which compactness and continuity are recovered.

Thank you

Vasileios Nastoulis

Athens University of Economics and Business

References I

Aumann, R. (1976). Agreeing to disagree.

DeMarzo, P., and Skiadas, C. (1998, 1999). Rational expectations equilibrium and separably oriented securities.

Dimitrov, D., and Sami, R. (2008). Strategic behavior in prediction markets.

Galanis, S., Ioannou, K., and Kotronis, V. (2024). Dynamic information aggregation under ambiguity.

Grossman, S. (1976). On the efficiency of competitive stock markets where traders have diverse information.

Hanson, R. (2003). Combinatorial information market design.

Hayek, F. A. (1945). The use of knowledge in society.

Holden, C., and Subrahmanyam, A. (1994). Risk aversion, imperfect competition, and long-lived information.

Kyle, A. (1985, 1989). Strategic trading models.

Offerman, T., Sonnemans, J., van de Kuilen, G., and Wakker, P. (2009). A truth serum for non-Bayesians: Correcting proper scoring rules for risk attitudes.

References II

Ostrovsky, M. (2012). Information aggregation in dynamic markets with strategic traders.

Peysakhovich, A., and Plagborg-Møller, M. (2012). Correcting proper scoring rules for risk attitudes.

Rostek, M., and Weretka, M. (2012, 2015). Price inference and market making.