

Department of Economics

Information Aggregation with Costly Information Acquisition

Spyros Galanis and Sergei Mikhalishchev

Working Paper No. 5, 2025

Department Economics

Durham University Business School

Mill Hill Lane

Durham DH1 3LB, UK

Tel: +44 (0)191 3345200

<https://www.durham.ac.uk/business/about/departments/economics/>

© Durham University, 2025

Information Aggregation with Costly Information Acquisition*

Spyros Galanis[†] Sergei Mikhalevich[‡]

This draft: November 24, 2025

Abstract

We study information aggregation in a dynamic trading model with partially informed traders. [Ostrovsky \[2012\]](#) showed that ‘separable’ securities aggregate information in all equilibria, however, determining whether a security is separable requires knowing the exact information structure of agents. To remedy this problem, we allow traders to acquire signals with cost κ , in every period. We show that ‘ κ separable securities’ characterize information aggregation and, as the cost decreases, almost all securities become κ separable, irrespective of the traders’ initial private information. Moreover, the switch to κ separability happens not gradually but discontinuously, hence even a small decrease in costs can result in a security aggregating information. We provide a complete classification of securities in terms of how well they aggregate information, which surprisingly depends only on their payoff structure.

JEL: C91, D82, D83, D84, G14, G41

Keywords: Information Aggregation, Information Acquisition, Financial Markets, Prediction Markets.

*We would like to thank Emir Kamenica, Anastasios Karantounias, Ehud Lehrer, Suraj Malladi, Marciano Siniscalchi and participants at Hitotsubashi, Northwestern, DIMACS 2024 Workshop on Forecasting at Rutgers, the Workshop on Information Aggregation in Durham, the 10th LGT Conference in Lancaster, Athens University of Economics and Business, the 7th World Congress of the Game Theory Society in Beijing, the EC24 Conference on Economics and Computation at Yale, Glasgow, Sussex, the CRETE Conference in Milos, Greece, the MIMA Workshop in Macroeconomic Theory at Warwick, the CEPR Workshop in Turin, the Durham York Workshop in Economic Theory, the Durham Economic Theory Conference, EWMES2023 in Manchester, and SAET2023 in Paris. This research is funded under ESRC grant ES/V004425/1.

[†]Department of Economics, Durham University, Mill Hill Lane, Durham, UK, DH1 3LB. Email: spyros.galanis@durham.ac.uk.

[‡]Department of Economics, Durham University, Mill Hill Lane, Durham, UK, DH1 3LB. Email: sergei.mikhalevich@durham.ac.uk.

1 Introduction

The question of whether financial markets reveal and aggregate the private information of traders has been studied at least since Hayek [1945]. Ostrovsky [2012] provides a strong result, that information gets aggregated in *all Nash equilibria* if the traded securities are *separable* and trading takes place for infinitely many periods. A security is a function from states to payoffs. It is non-separable given an information structure if there is a prior under which its payoffs are uncertain, yet all traders agree on its expected value given their private information. If no such prior exists, it is separable. However, determining whether a security is separable requires knowing the exact information structure of agents. As a result, a market designer who does not know who participates or what is their information structure, cannot be sure that the equilibrium price is a good predictor of the security's value.

In this paper, we examine whether the ability to acquire costly signals during trading can make markets more efficient at aggregating information, so that the price of a security always converges to its value. This question becomes more relevant as the continuous improvements in information technology have created an abundance of available information, which is now cheaper than ever to acquire, analyze, and act upon.¹

Our main result, Theorem 1 with Proposition 1, is that for a generic security, which pays differently across all states, and for sufficiently low cost of information acquisition, information will get aggregated, *irrespective* of who trades and what is the information structure. In other words, the availability of cheap information makes ‘most’ markets aggregate information under very general conditions.²

We use the dynamic trading model of Ostrovsky [2012] with infinitely many periods and payoffs given by the Market Scoring Rule (MSR) [Hanson, 2003, 2007]. Each Trader's private information is represented by a partition of the state space, and the conjunction of everyone's private information reveals the true value of the security. The model is enhanced by allowing traders to acquire a signal $\mathcal{R}$ in each period, before trading, at a cost of $cK(\mathcal{R})$, where K is a cost function of signals and c is a strictly positive marginal cost. Given $\kappa = (K, c)$, we say that a security

¹For example, recent advances in generative AI tools such as ChatGPT could add considerable value for investors with information processing constraints [Kim et al., 2023] and assist in picking stocks [Pelster and Val, 2023].

²This result is also supported by empirical evidence. Farboodi et al. [2022] show, using a structural model, that as the value of a firm's data grows, which is equivalent in our model to a decrease in the cost of information acquisition, the information content of the price of the firm's stock increases as well.

is κ *non-separable* given an information structure if it is non-separable at some prior and no trader finds it myopically optimal to acquire any signal at that prior. Otherwise, it is κ separable.

Theorem 1 shows that κ separable securities are necessary and sufficient for information aggregation, when the cost of information is κ , thus generalising Ostroy [2012]. As the marginal cost c decreases, the securities that eventually become κ separable, and therefore aggregate information in *all* Nash equilibria and for *all* information structures, have a very simple structure: they specify a different payoff at each state (Proposition 1).

Surprisingly, there is a second, small class of securities, which Proposition 2 characterises, that never become κ separable for *all* information structures, and therefore may fail information aggregation, even when c converges (but is not equal) to zero.³ Such a security specifies the same payoff d in two states, while in two other states it pays either higher or lower than d .

The third and final class consists of securities which are always separable, irrespective of who trades and what their information structure is. These securities aggregate information even if costly signals are unavailable. Unfortunately, this class is very small and uninformative. Proposition 3 shows that it only contains the Arrow-Debreu (A-D) security, which pays a at some state and b in all other states, and the security that specifies three payoffs: the largest is paid in one state of the world, the lowest in another, and the middle in all other states.⁴ For any other security, there is a market (information structure and common prior) at which there is no information aggregation so that the security's price does not converge to its true value. This means that information may not aggregate in 'most' markets if signals are too costly or unavailable.

These three classes provide a complete classification of securities in terms of information aggregation, which depends only on their payoff structure. First, the most informative class of κ separable securities for some κ , ensure that information gets aggregated as long as the cost of information acquisition is low. Second, the securities which are κ non-separable for all κ , may fail information aggregation even if the cost is low. Finally, the always separable securities aggregate information even if costly signals are unavailable.

If the marginal cost of information acquisition gets arbitrarily small, do we

³If the cost of information acquisition is zero, then information aggregates trivially.

⁴These securities are not very informative because the A-D can only predict whether one state has occurred or not, whereas the other security can only predict whether two states have occurred or not. Note that combining more than one A-D security to construct a composite and more informative security will not solve this issue, because it will be non-separable for some information structures.

even need markets to aggregate information through prices? Each Trader could buy the necessary signals and then bid very close to the true value. We argue that this intuition is incorrect, because markets become even more important in an environment with information acquisition. As the marginal cost c decreases, the security switches discontinuously from κ non-separable to κ separable, hence even a slight reduction can enable a market to aggregate information in all Nash equilibria. This is in contrast to a poll, defined as the average of the traders' opinions after receiving the information, because its predictive accuracy improves smoothly as costs decrease (see Section 5). Hence, the availability of cheap information leverages the value of the markets, enabling them to aggregate information long before the cost goes to zero. In fact, Proposition 4 shows that a security is κ separable if and only if the market is more accurate than a poll, for all priors.

1.1 Literature

Our paper contributes to two strands of the literature. The first studies the inefficiency of information acquisition and its effect on information aggregation in markets.⁵ Pavan et al. [2022] show that traders acquire and use information inefficiently. Moreover, as the cost of information declines, traders over-invest in information acquisition and trade too much on their private information. Several experimental studies support these results and find that traders tend to over-acquire information. In addition, while information acquisition is positively correlated with market efficiency, market prices do not aggregate all private information [Kraemer et al., 2006, Page and Siemroth, 2017, 2021, Corgnet et al., 2022]. Mele and Sangiorgi [2015] analyze costly information acquisition in asset markets with ambiguity-averse traders and show that when uncertainty is high enough, information acquisition decisions become strategic complements and lead to multiple equilibria. Our paper complements and differs from this literature. We find that, as the cost of information acquisition decreases, the number of securities (and therefore markets) that aggregate information increases. However, some securities are never able to aggregate information, even if the cost is almost zero.

The second strand looks at the information aggregation properties of financial and, in particular, prediction markets.⁶ DeMarzo and Skiadas [1998, 1999] first introduced the notion of separable securities. Ostrovsky [2012] and Chen et al. [2012] show that in a market with dynamically consistent traders, separable securities are both necessary and sufficient for information aggregation. Dimitrov

⁵See Lim and Brooks [2011] for a survey of the empirical literature on market efficiency.

⁶See Wolfers and Zitzewitz [2004] for an early overview of the literature.

and Sami [2008] and Chen et al. [2010] examine information aggregation by varying the assumptions regarding the traders’ information structure. Galanis et al. [2024] show theoretically and experimentally that ambiguity aversion can lead to no information aggregation with separable securities, but a new class - strongly separable securities - overcomes this limitation. Galanis and Kotronis [2021] show that information aggregation may fail if boundedly rational traders are unable to update their awareness when confronted with price announcements that they cannot rationalize.⁷ We contribute to this literature by allowing traders to acquire costly signals at every period and propose a new class of κ separable securities which characterize information aggregation.

We conclude by motivating our choice of the MSR model. First, in the MSR model, there are no noise traders and no strategic market makers, hence the issue of information aggregation is not intertwined with that of information revelation, as in Kyle [1985].⁸ Unlike the MSR, in Kyle [1985] it is not always the case that the price will converge to the true value of the security, even if there is only one Trader and therefore information aggregation is achieved by default. Second, a prediction market with the MSR can be reinterpreted as an inventory-based market with a market maker who continuously adjusts the price of the securities depending on the orders she receives.⁹ The advantage of the MSR over more well-known market mechanisms, such as the continuous double auction, is that an agent can make her prediction/trade without waiting for another agent to take the opposite side, or submit a limit order and wait for it to be filled. This feature makes it an attractive mechanism for markets with relatively few participants who do not trade daily, or in markets with automated market makers.¹⁰ MSR-based prediction markets have been used widely, for example, by firms such as Ford, Google, General Electric, and Chevron (see Ostrovsky [2012], Cowgill and Zitzewitz [2015]) as well as governments, for example, in the UK and the Czech Republic (The Economist [2021]).

The paper adheres to the following plan. Section 2 describes the model and Section 3 presents our main results. In Section 4 we provide a complete clas-

⁷Unawareness and ambiguity aversion generate dynamic inconsistency and negative value of information, which are partly responsible for no information aggregation. See Galanis [2011, 2013] for a model of unawareness and Galanis [2021] for a connection between dynamic inconsistency and the negative value of information.

⁸Note that Ostrovsky [2012] uses both the MSR and the Kyle [1985] model.

⁹See Ostrovsky [2012] and Galanis et al. [2024] for examples.

¹⁰Automated market makers are widely used in Decentralized Finance, see Schlegel et al. [2022] for an axiomatization of the logarithmic MSR. Frongillo et al. [2024] show the equivalence of prediction markets and constant function market makers that are used overwhelmingly when trading on the blockchain.

sification of securities. Section 5 compares a poll with a prediction market and shows that a low cost of information acquisition leverages the value of the market. Section 6 concludes.

2 The Model

2.1 Preliminaries

Uncertainty is described by a finite state space $\Omega = \{\omega_1, \dots, \omega_l\}$ and the set of traders is denoted $I = \{1, \dots, n\}$. Trader i 's initial private information is represented by partition Π_i of Ω . Let $\Pi_i(\omega)$ be a partition element of Π_i that contains ω , so that $\omega \in \Pi_i(\omega) \in \Pi_i$. When the true state is $\omega \in \Omega$, Trader i considers all states in $\Pi_i(\omega) \subseteq \Omega$ to be possible. We assume that the join (the coarsest common refinement) of partitions $\Pi = \{\Pi_1, \dots, \Pi_n\}$ consists of singleton sets so that $\bigcap_{i \in I} \Pi_i(\omega) = \omega$ for all $\omega \in \Omega$, which means that the traders' pooled information always reveals the true state.¹¹ This implies that, for any two states $\omega_1 \neq \omega_2$, there exists Trader i such that $\Pi_i(\omega_1) \neq \Pi_i(\omega_2)$. Let $\mathcal{P}$ be the collection of all information structures Π where Ω has at least three states and $\bigcap_{i \in I} \Pi_i(\omega) = \{\omega\}$ for all $\omega \in \Omega$. Traders have a full-support common prior μ_0 over Ω and they are risk-neutral.

2.2 Trading environment

Trading is organized as follows. There are infinitely many periods and finitely many I traders, with $|I| = n$, who trade a security $X : \Omega \rightarrow \mathbb{R}$. At time $t_0 = 0$, nature selects a state $\omega^* \in \Omega$ and the uninformed market maker makes a prediction y_0 about the value of security X . At time $t_1 > t_0$, Trader 1 makes a revised prediction y_1 , at $t_2 > t_1$ Trader 2 makes his prediction, and so on. At time $t_{n+1} > t_n$, Trader 1 makes another prediction y_{n+1} , and the whole process repeats until time $t_\infty \equiv \lim_{k \rightarrow \infty} t_k = 1$. All predictions are observed by all traders. Each prediction y_k is required to be within the set $[\underline{y}, \bar{y}] = [\min_{\omega \in \Omega} X(\omega), \max_{\omega \in \Omega} X(\omega)]$. At some time $t^* > 1$ the true value $x^* = X(\omega^*)$ of the security is revealed.

The traders' payoffs from each announcement are computed using a scoring rule, $s(y, x^*)$, where x^* is the true value of the security and y is a prediction. A scoring rule is *proper* if, for any probability measure μ and any random variable

¹¹This assumption is also made by [Ostrovsky \[2012\]](#) and it is without loss of generality because if the conjunction of the traders' private information does not reveal the state, we cannot expect that trading the security will reveal it.

X , the expectation of s is maximized at $y = E_\mu[X]$. It is *strictly proper* if y is unique. We focus on continuous strictly proper scoring rules. Examples are the quadratic, where $s(y, x) = -(x - y)^2$, and the logarithmic, where $s(y, x) = (x - a)\ln(y - a) + (b - x)\ln(b - y)$ with $a < \min_{\omega \in \Omega} X(\omega)$, $b > \max_{\omega \in \Omega} X(\omega)$.

Under the market scoring rule (MSR) (McKelvey and Page [1990], Hanson [2003, 2007]), a Trader is paid for each revision he makes. In particular, his payoff, from announcing y_n at t_n , is $s(y_n, x^*) - s(y_{n-1}, x^*)$, where y_{n-1} is the previous announcement and x^* is the true value of the security. For all proper scoring rules, as $E_q[X]$ converges to $X(\omega)$, $s(E_q[X], X(\omega))$ converges to 0. Moreover, if y_{n-1} is further away from $X(\omega)$ than $E_q(X)$ is from $X(\omega)$, then $s(E_q[X], X(\omega)) - s(y_{n-1}, X(\omega))$ is strictly positive. We then say that the Trader “buys out” the previous Trader’s prediction. If he repeats the previous announcement, his period payoff is zero. We say that prior μ is non-degenerate given security X if it does not assign probability 1 to a unique value of X .

2.3 Cost of Information

We enhance the model of Ostrovsky [2012] by allowing traders to acquire a costly signal in each time t where they make an announcement. A signal (or statistical experiment) $\mathcal{R}$ consists of a signal realization space $T_{\mathcal{R}} \subseteq T$ and a family of distributions $\{\mathcal{R}(\cdot|\omega)\}_{\omega \in \Omega}$ over $T_{\mathcal{R}}$. Let $\mathcal{E}$ be the compact set of available signals, in the standard product topology on $\bigcup_{\mathcal{R} \in \mathcal{E}} \Delta(T_{\mathcal{R}})^\Omega$.¹²

Given a signal $\mathcal{R}$ and a prior μ , each signal realization $\tau \in T_{\mathcal{R}}$ generates a posterior $\gamma \in \Delta(\Omega)$. Let $\mathcal{R}_\mu \in \Delta(\Delta(\Omega))$ be the generated random posterior, so that the probability of posterior γ is $\mathcal{R}_\mu(\gamma)$.

Definition 1. *The cost of signal $\mathcal{R}$ is $cK(\mathcal{R})$, where $c > 0$ is the marginal cost and $K : \mathcal{E} \rightarrow [0, +\infty]$ is a continuous function. Let $\mathcal{K}$ be the collection of all cost structures $\kappa = (c, K)$, where K satisfies the following assumptions:*

- (i) *If $\kappa = (c, K) \in \mathcal{E}$ then $\kappa' = (c', K) \in \mathcal{E}$, for all $c' > 0$,*
- (ii) *There is an uninformative signal $\mathcal{R}_0$, so that $\mathcal{R}_0(\cdot|\omega) = \mathcal{R}_0(\cdot|\omega')$ for all $\omega, \omega' \in \Omega$, with $K(\mathcal{R}_0) = 0$, and K is sufficiently smooth around $\mathcal{R}_0$,¹³*

¹²We use the Levy-Prokhorov metric d_W (Dekel et al. [2006]) to generate the topology of weak convergence on $\Delta(T)$. The distance between two signals is defined as $D(R, P) = \max_{\omega \in \Omega} d_W(\mathcal{R}(\cdot|\omega), \mathcal{P}(\cdot|\omega))$.

¹³That is, $\limsup_{\mathcal{R} \rightarrow \mathcal{R}_0} \frac{K(\mathcal{R})}{[D(\mathcal{R}, \mathcal{R}_0)]^2} < \infty$ for all $\mathcal{R} \in \mathcal{E}$.

(iii) If $\mathcal{R}'$ is a garbling of $\mathcal{R}$ then $K(\mathcal{R}) \geq K(\mathcal{R}')$.¹⁴

We separately specify the marginal cost so that, for any K , we can express a sufficiently low cost of information acquisition with $c \rightarrow 0$, and (i) specifies that all marginal costs are available. Condition (ii) says that an uninformative signal is available and free, and K is sufficiently smooth around zero information. Finally, (iii) says that more informative signals are costlier.

For the second part of Theorem 1 and Proposition 4, we assume that if a Trader can generate a random posterior Q with a sequence of signals, it is weakly less costly to generate it with just one signal. This assumption is similar to the Sequential Learning-Proof property in Bloedel and Zhong [2024]. In particular, consider a prior μ , a signal $\mathcal{R}$, and a collection of signals $\{\mathcal{R}^\tau\}_{\tau \in T_{\mathcal{R}}}$, one for each signal realization $\tau \in T_{\mathcal{R}}$. Let $Q = \{\mathcal{R} \otimes \{\mathcal{R}^\tau\}_{\tau \in T_{\mathcal{R}}}\}_{\mu}$ be the random posterior generated by the combination of these signals, according to the following procedure. First, the Trader conducts signal $\mathcal{R}$. After receiving each signal realization $\tau \in T_{\mathcal{R}}$ and updating his posterior to γ^τ , he conducts signal $\mathcal{R}^\tau$. After receiving a new signal realization $\tau' \in T_{\mathcal{R}^\tau}$, he updates his posterior to $\gamma_{\mathcal{R}^\tau}^{\tau'}$. The ex ante cost of the generated random posterior is $C = c \left(K(\mathcal{R}) + \sum_{\tau \in T_{\mathcal{R}}} \mathcal{R}_\mu(\gamma^\tau) K(\mathcal{R}^\tau) \right)$. We assume that the same random posterior $Q = \{\mathcal{R} \otimes \{\mathcal{R}^\tau\}_{\tau \in T_{\mathcal{R}}}\}_{\mu}$ can be generated by just one signal $\mathcal{R}'$, which costs weakly less than C .

Assumption 1. For any prior μ and signals $\mathcal{R}$, $\{\mathcal{R}^\tau\}_{\tau \in T_{\mathcal{R}}}$, there is signal $\mathcal{R}'$ such that $\{\mathcal{R} \otimes \{\mathcal{R}^\tau\}_{\tau \in T_{\mathcal{R}}}\}_{\mu} = \mathcal{R}'_{\mu}$ and $K(\mathcal{R}) + \sum_{\tau \in T_{\mathcal{R}}} \mathcal{R}_\mu(\gamma^\tau) K(\mathcal{R}^\tau) \geq K(\mathcal{R}')$.

2.4 The Game

A game $\Gamma(\Omega, \Pi, X, \mu, \mathcal{E}, c, K, y_0, s, \beta)$ is a tuple where Ω is the state space, Π is the collection of partitions which describes the initial private information of traders, X is the traded security, μ is the common prior, $\mathcal{E}$ is the collection of signals, c is the marginal cost and K is the cost function of signals, y_0 is the market maker's initial announcement at time t_0 , s is a strictly proper scoring rule, and β is the common discount rate. The range of possible announcements is $[\underline{y}, \bar{y}] = [\min_{\omega \in \Omega} X(\omega), \max_{\omega \in \Omega} X(\omega)]$.

At time t_k , one Trader chooses a signal and specifies an announcement in $[\underline{y}, \bar{y}]$, for each possible signal realization. Then, nature moves and picks a signal realization.

¹⁴Recall that signal $\mathcal{R}'$ is a garbling of signal $\mathcal{R}$ if there exists a map $g : T_{\mathcal{R}} \rightarrow \Delta(T_{\mathcal{R}'})$ such that for each $\tau' \in T_{\mathcal{R}'}$, $\mathcal{R}'(\tau'|\omega) = \sum_{\tau \in T_{\mathcal{R}}} g(\tau'|\tau) \mathcal{R}(\tau|\omega)$, for all $\omega \in \Omega$.

tion. Given the Trader's specification, the announcement $y_k \in [\underline{y}, \bar{y}]$ corresponding to that signal realization is made.

Then, nature moves, picks a signal realization and the corresponding announcement y_k is made. The announcement y_k becomes common knowledge but the choice of the signal and the signal realization are the private information of the Trader who announced.

Following [Dimitrov and Sami \[2008\]](#), the payoff of Trader i who makes an announcement y_k at time t_k is $\beta^k (s(y_{t_k}, x^*) - s(y_{t_{k-1}}, x^*) - cK(\mathcal{R}))$, where $0 < \beta \leq 1$ is the common discount factor and $\mathcal{R} \in \mathcal{E}$ is the signal that he acquired.¹⁵ Both the choice of signal $\mathcal{R}$ and the signal realization is private and not revealed to other traders. The total payoff of each Trader is the sum of all payments for revisions, minus the cost of the signals he acquired. We also consider the special case of myopic, or non-strategic, traders, so that when making an announcement they do not take into consideration the impact of the other traders' strategies on their future payoffs. Note that each Trader can guarantee a zero payoff at any round, by repeating the previous announcement and not buying any signal.

Let $(y_1, \dots, y_k) \in [\underline{y}, \bar{y}]^k$ be a history of announcements and $(\tau_1, \dots, \tau_k) \in (\bigcup_{\mathcal{R} \in \mathcal{E}} T_{\mathcal{R}})^k$ be a history of signal realisations, up to time t_k . Denote by T_i^k the collection of all histories about i 's signal realisations up to time t_k . For simplicity, we assume that a signal realisation $\tau \in T_{\mathcal{R}}$ is unique to a specific signal, therefore a history of realisations in T_i^k also identifies the chosen signals.

Let $[\underline{y}, \bar{y}]^{T_{\mathcal{R}}}$ be the set of all functions from signal realizations of $\mathcal{R}$ to announcements in $[\underline{y}, \bar{y}]$. Trader i 's mixed strategy at time t_k is a measurable function

$$\sigma_{i,k} : \Pi_i \times T_i^{k-1} \times [\underline{y}, \bar{y}]^{k-1} \times [0, 1] \longrightarrow \bigcup_{\mathcal{R} \in \mathcal{E}} [\underline{y}, \bar{y}]^{T_{\mathcal{R}}}.$$

It specifies a signal $\mathcal{R}$ and an announcement y_k for each realization that is drawn from that signal, given the element of his partition at the true state ω , the history of his past chosen signal realizations, the history of all announcements up to time t_{k-1} , and the realization of random variable $\iota_k \in [0, 1]$, which is drawn from the uniform distribution. One such draw takes place at each time t_k , it is observed only by the Trader who makes the announcement at t_k , and the draws are independent of each other and of the true state ω .¹⁶ The draws of the uniform distribution allow each Trader to randomize both in terms of the chosen signals

¹⁵For some examples, we allow myopic traders with $\beta = 0$, where each Trader makes the myopically best announcement in each period.

¹⁶This formulation, with draws on $[0, 1]$ denoting the Trader's randomization, is taken from [Ostrovsky \[2012\]](#).

and the announcements given the signal realisations. Note that, although all announcements are public, the history of past signals and realisations for each Trader i is private.

The full state $\phi = (\omega, \iota_1, \iota_2, \dots)$ resolves the initial uncertainty about Ω and the randomizations of the players according to ι . Let $X(\phi) \equiv X(\omega)$ be the true value of X at $\phi = (\omega, \iota_1, \iota_2, \dots)$ and $\Phi = \Omega \times [0, 1]^\mathbb{N}$ be the full state space. Denote by σ_i the collection $\sigma_{i,k}$ of i 's strategies, at all times t_k where it is his turn to make an announcement. Let $\sigma = (\sigma_1, \dots, \sigma_n)$ be a profile of strategies and Σ be the set of all strategy profiles. Given a strategy σ and state ϕ , let $y_{i+nk}(\sigma, \phi)$ be the announcement of Trader i and $cK(\mathcal{R}_{i+nk}(\sigma, \phi))$ the cost of acquiring signal $\mathcal{R}_{i+nk}(\sigma, \phi)$ at time t_{i+nk} .

Definition 2. *A strategy profile σ is a Nash equilibrium if, for every Trader i and every alternative strategy $\sigma' = (\sigma_{-i}, \sigma'_i)$, we have*

$$E_\mu \left[\sum_{k=0}^{\infty} \beta^{i+nk} \left(s(y_{i+nk}(\sigma, \phi), X(\phi)) - s(y_{i+nk-1}(\sigma, \phi), X(\phi)) - cK(\mathcal{R}_{i+nk}(\sigma, \phi)) \right) \right] \geq$$

$$E_\mu \left[\sum_{k=0}^{\infty} \beta^{i+nk} \left(s(y_{i+nk}(\sigma', \phi), X(\phi)) - s(y_{i+nk-1}(\sigma', \phi), X(\phi)) - cK(\mathcal{R}_{i+nk}(\sigma', \phi)) \right) \right],$$

where the expectation is taken with respect to the common prior μ .

2.5 Information Aggregation

We say that information is aggregated if the traders' predictions converge to the intrinsic value $X(\phi)$ of security X , for all $\phi \in \Phi$. For every $\phi \in \Phi$ and strategy profile σ , let $y_k(\sigma, \phi)$ be the announcement of the Trader who moves in period t_k . The announcement $y_k(\sigma, \phi)$ depends on $\phi = (\omega, \iota_1, \iota_2, \dots)$ because traders have different private information across Ω and randomize based on the realizations of random variable ι . Because $\{y_k\}_{k=1}^\infty$ is a sequence of random variables, we need a probabilistic version of convergence.

Definition 3. *Under a profile of strategies in Γ , information aggregates if sequence $\{y_k\}_{k=1}^\infty$ converges in probability to random variable X .*

2.6 Separable Securities

Consider the following example, which illustrates the notion of separability.

Example 1. The state space is $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, the security is $X = (0, 1, 2, 3)$, and there are two myopic traders with common prior $\mu = (\frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8})$. Trader 1's partition is $\{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$ and Trader 2's is $\{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}\}$.

As traders are myopic in this example, under the MSR they take turns in announcing their expected value of X . Suppose that the true state is ω_1 . In period 1, Trader 1 with conditional beliefs $(\frac{1}{4}, 0, \frac{3}{4}, 0)$ will announce $\frac{3}{2}$. The same announcement would be made by Trader 1 in all states, and thus no information is transmitted to Trader 2. In period 2, Trader 2 also makes the same announcement at ω_1 , and the same announcement would be made in all states. Hence, no information is transmitted to Trader 1. Because the two traders agree on the announcement, there is no information updating, and the process ends. We say that there is no information aggregation at ω_1 because the final announcement $\frac{3}{2}$ is not equal to the intrinsic value of X at ω_1 , which is 0.¹⁷

At all states, it is common knowledge that both traders agree that the expected value of X is $\frac{3}{2}$, hence there can be no more learning from further (myopic) announcements. However, it is also common knowledge that the intrinsic value of X is not $\frac{3}{2}$; it is either 0, 1, 2 or 3, which implies that there is no information aggregation. When both these conditions are satisfied for some prior and some partitions, the security is non-separable. If there is no such prior, the security is separable.

Definition 4. A security X is called non-separable under information structure Π if there exists probability distribution μ and value $v \in \mathbb{R}$ such that:

- (i) $X(\omega) \neq v$ for some $\omega \in \text{Supp}(\mu)$,
- (ii) $E_\mu[X|\Pi_i(\omega)] = v$ for all $i = 1, \dots, n$ and $\omega \in \text{Supp}(\mu)$.

We then say that security X is non-separable at μ . Otherwise, it is called separable.

A security X is non-separable if, for some prior μ and at all states in its support, all traders' expected value of X is v , yet there is uncertainty about the value of X . If for all priors at least one condition is violated, then the security is separable. Note that separability (and non-separability) is a property that depends on the information structure Π . For a different Π , a security may switch from separable to non-separable and vice versa. Theorem 1 in [Ostrovsky \[2012\]](#) shows that separability characterizes information aggregation, in an environment without information acquisition.

¹⁷Information aggregation fails in the first round in this example and no-one updates from μ . In general, traders could start from a different common prior μ' , and after several rounds of updating they could update to some other posterior μ'' , at which there is no information aggregation.

2.7 Information Acquisition

To illustrate how information gets aggregated, even when the security is non-separable, we enhance Example 1 by allowing traders to acquire extra information about the value of the security through noisy signals, before making their announcement. Suppose that Trader 1 can acquire signal $\mathcal{R}$ that generates a signal realization z , given each state, with probabilities $\mathcal{R}(z|\omega_1) = 1$ and $\mathcal{R}(z|\omega_2) = \mathcal{R}(z|\omega_3) = \mathcal{R}(z|\omega_4) = 0.5$.

At state ω_1 , Trader 1's prior belief is $(\frac{1}{4}, 0, \frac{3}{4}, 0)$ and, after receiving the signal realization z , his posterior belief is $(\frac{2}{5}, 0, \frac{3}{5}, 0)$. Then, he announces the expected value of X , which is $\frac{6}{5}$. Trader 2 considers states ω_1 and ω_4 to be possible. The public announcement of $\frac{6}{5}$ reveals to Trader 2 that the true state is ω_1 . The reason is that, regardless of whether Trader 1 received a signal z or not, his posteriors at ω_4 are $(0, \frac{3}{4}, 0, \frac{1}{4})$ and he would have announced $\frac{3}{2}$. As a result, Trader 2 announces 0 in the second round, and the game ends. Note that the final price is equal to the intrinsic value of X at ω_1 , hence information gets aggregated. In summary, the ability to acquire extra signals transforms X from non-separable to κ separable, enabling information aggregation.

We make two observations. First, we have abstracted from the cost of acquiring an experiment. Each Trader will acquire the signal only if the expected continuation value from making a better prediction outweighs the cost. Second, Trader 2 free-rides on Trader 1 buying the signal. By moving the price from $\frac{6}{5}$ to 0, he books a profit, without paying the cost of a signal.

2.8 κ Separable Securities

Example 1 showed that even if a security is non-separable at some prior μ , as long as traders can buy costly signals they can further move the announcement and achieve information aggregation. Hence, information aggregation may get 'stuck' only if no-one is willing to acquire any new signals. We say that there is no (myopic) information acquisition if no Trader is willing to acquire any signal, given the cost structure $\kappa = (K, c)$.

Definition 5. *There is no information acquisition for security X , given prior μ , cost structure $\kappa = (K, c)$ and information structure $\Pi \in \mathcal{P}$, if for all states $\omega \in \text{Supp}(\mu)$, all traders i , and all signals $\mathcal{R}$,*

$$\sum_{\gamma \in \text{Supp}(\mathcal{R}_{\mu_{\Pi_i}(\omega)})} \mathcal{R}_{\mu_{\Pi_i}(\omega)}(\gamma) \sum_{\omega' \in \Omega} \gamma(\omega') \left[s(E_\gamma[X], X(\omega')) - s(E_{\mu_{\Pi_i}(\omega)}[X], X(\omega')) \right] - cK(\mathcal{R}) \leq 0.$$

Otherwise, there is information acquisition.

To explain this inequality, suppose that at time t it is Trader i 's turn to make an announcement. Having observed all previous announcements and using the public information that is revealed, an outside observer updates the common prior μ_0 to a belief μ over Ω . If the true state is ω , then Trader i 's private information is $\Pi_i(\omega)$ and his posterior belief is the Bayesian update of μ , denoted $\mu_{\Pi_i(\omega)}$. In other words, he updates using both his private information and the public information revealed by previous announcements.

His myopic problem at ω consists of buying an experiment $\mathcal{R}$, at cost $cK(\mathcal{R})$, which generates a random posterior $\mathcal{R}_{\mu_{\Pi_i(\omega)}}$. For each posterior belief $\gamma \in \text{Supp}(\mathcal{R}_{\mu_{\Pi_i(\omega)}})$, his myopic best response is to announce $E_\gamma[X]$ because the scoring rule is proper. If his expected utility is weakly negative for all signals $\mathcal{R}$, then we say that he does not want to acquire any information, because buying no signal and repeating the previous announcement $E_{\mu_{\Pi_i(\omega)}}[X]$ guarantees a payoff of zero.¹⁸

We say that security X is κ non-separable given information structure Π and cost structure κ if there is a prior μ such that there is no myopic information acquisition and the security is non-separable at μ .

Definition 6. *Security X is κ non-separable given an information structure $\Pi \in \mathcal{P}$ and cost structure $\kappa \in \mathcal{K}$, if there exists prior μ such that*

- *Security X is non-separable at μ ,*
- *There is no information acquisition for X given μ .*

Otherwise, security X is κ separable.

3 Main Result

Our main result consists of two parts, Theorem 1 and Proposition 1. Collectively, they specify that for a security which pays differently across all states, and for sufficiently low cost of information acquisition, information will get aggregated, *irrespective* of who trades and what is the information structure Π .

The first part, Theorem 1, shows that κ separable securities characterize information aggregation when the cost structure is $\kappa = (K, c)$.

¹⁸Note that, as in [Ostrovsky \[2012\]](#), the previous announcement does not influence the choice of the myopically optimal choice of the announcement (and the signal). We have specified $E_{\mu_{\Pi_i(\omega)}}[X]$ to be the previous announcement so that we normalize at zero the utility from no information acquisition, because then the optimal myopic announcement is $E_{\mu_{\Pi_i(\omega)}}[X]$. Interestingly, the myopic best depends on the previous announcement when traders are ambiguity averse, as shown in [Galanis et al. \[2024\]](#).

Theorem 1. *Fix information structure Π and cost structure $\kappa = (K, c)$. Then:*

- *If security X is κ separable under Π , then information gets aggregated in any Nash equilibrium of game $\Gamma(\Omega, \Pi, X, \mu, \mathcal{E}, c, K, y_0, s, \beta)$,*
- *If security X is κ non-separable under Π , then there exists prior μ such that for all s , y_0 , and β , there exists a Nash equilibrium of the corresponding game Γ in which information does not get aggregated.*

See [proof](#) on page 28.

Although Theorem 1 ensures that information aggregates in all Nash equilibria, it does not guarantee that an equilibrium exists. This is true also in [Ostrovsky \[2012\]](#), and the reason is that the action spaces are infinite. In [Appendix A](#), we prove the existence of a Perfect Bayesian Equilibrium if the action spaces are finite. We first show that if we discretize the action spaces, so that traders can choose from a finite set of announcements and signals, then a Perfect Bayesian Equilibrium exists whenever the time horizon is finite. Second, we show that each game Γ is continuous at infinity. Finally, we adapt the proofs of [Fudenberg and Levine \[1983, 1986\]](#) to approximate the infinite horizon game with a sequence of finite horizon games and show that the sequence of Perfect Bayesian equilibria in the games with finite horizon converges to a Perfect Bayesian equilibrium in the game with infinite horizon.¹⁹ We also show that the sequence of announcements converges to a limit which is as close as possible to the correct value of the security. Hence, information aggregation improves as the action space becomes richer.

The second part of our main result says that if security X pays differently across all states, then it is κ separable for some marginal cost of information c , given *any* information structure Π and cost function K . In other words, if information is sufficiently cheap to acquire, there are easily describable securities where information aggregation is achieved at all Nash equilibria and independently of who participates in the market or what is their private information.

Proposition 1. *If $X(\omega) \neq X(\omega')$ for all $\omega, \omega' \in \Omega$, then, given any Π and K , X is κ separable for some $c > 0$.*

See [proof](#) on page 33.

[Proposition 1](#) echoes the following result within the context of perfectly competitive markets, as pointed out in [Laffont and Maskin \[1990\]](#). As long as there

¹⁹[Ostrovsky \[2012\]](#) describes such an approach but does not provide further details. Existence of equilibrium is also proven by [Galanis et al. \[2024\]](#), for the identical game without information acquisition but with ambiguity aversion.

are “enough prices”, so that trade can be made contingent on sufficiently many events, then the competitive equilibrium is generically separating and the function relating information to prices is invertible (Grossman [1976], Radner [1979], Allen [1981]). We could interpret Proposition 1 as a generalization of this result, because our setting applies to all Nash equilibria, not just competitive equilibria.

It would be natural to expect that, for a sufficiently low marginal cost c , any security eventually becomes κ separable. But this is not true. Consider the following example, which appears in similar form in Geanakoplos and Polemarchakis [1982] and in Example 1 of Ostrovsky [2012].

Example 2. *The state space is $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, the security is $X = (0, 1, 0, 1)$, and there are two traders with common prior $\mu = (1/4, 1/4, 1/4, 1/4)$. Trader 1’s partition is $\{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ and Trader 2’s is $\{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}\}$.*

With a uniform prior, all traders agree that the expected value of X is 0.5, at all states, hence the security is non-separable. Consider now prior $\mu = (\frac{1-2m}{2}, m, \frac{1-2m}{2}, m)$, where $0 < m < 0.5$. Given any m , each Trader’s expected value of X at all states is $2m$, so for m very close to 0.5, the consensus announcement is very close to 1. Given K and for any $c > 0$, we can always specify m that is sufficiently close to 0.5, so that no Trader is willing to buy a signal $\mathcal{R}$, because this will move their expected value of X from $1 - 2\epsilon$, which is the previous announcement, to $1 - \epsilon$ when they receive the signal realisation, for small enough $\epsilon > 0$. Their maximum benefit, when 1 is the true state, is $s(1 - \epsilon, 1) - s(1 - 2\epsilon, 1)$. Because the scoring rule is proper, this benefit is strictly positive but decreasing, as ϵ decreases. However, as the signal becomes ever more precise, its cost $cK(\mathcal{R})$ is weakly increasing. This means that for high enough $m < 0.5$, the expected benefit will be lower than the cost. Because there is no information acquisition and all traders agree on the expected value, security X is $\kappa = (K, c)$ non-separable for all $c > 0$.²⁰

The following Proposition generalizes this example, showing that the only securities that never become κ separable pay (a, d, b, d) , in four states, where either $a, b < d$ or $a, b > d$.

Proposition 2. *If Ω has at least four states, the following are equivalent:*

- X is κ non-separable given some Π and for all $\kappa \in \mathcal{K}$,

²⁰It is interesting to note that the incompatibility of the $(0, 1, 0, 1)$ vector with information aggregation is met in other settings as well. In the model of DeGroot [1974], agents update their beliefs naively, by looking at their immediate neighbours, according to a fixed network. If the network is represented by a periodic matrix, such as $A = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, then beliefs do not converge (Golub and Jackson [2010]).

- X pays (a, d, b, d) in four states, where either $a, b < d$ or $a, b > d$.

See [proof](#) on page 34.

To interpret this result, note that state space Ω is fixed independently of the securities and it describes all the higher order beliefs of traders about what others know and believe. The Proposition says that information aggregation may fail if there are two distinct states which specify the exact same payoff, yet they differ in terms of the higher order beliefs of traders and the only thing that matters to them are the payoffs.

4 Classification of Securities

Using the results from the previous section, we can provide a complete classification of securities in terms of how well they aggregate information, which surprisingly depends only on their payoff structure. There are three types of securities. First, the *always separable* securities, which we characterise in this section, are separable for *all* information structures and therefore aggregate information even if information acquisition was unavailable. Second, the κ non-separable for all κ and *some* information structure, such as the $(1, 0, 1, 0)$ security, which may fail information aggregation even if $c > 0$ is very small. Finally, the securities which are κ separable for some κ and *all* information structures, pay differently across all states and aggregate information if c is low enough. This classification does not depend on who is trading or what is their information structure. Moreover, it is very simple to classify each security, which is not true when determining whether a security is separable or not, given an information structure.

The always separable securities aggregate information even if costly signals are not available. Unfortunately, this class is very small and uninformative, as shown by the following Proposition. It consists of just three types. The first is the constant, paying the same at all states. The second is the Arrow-Debreu (A-D), which pays a at some state ω and b at all other states. The third pays a at some state ω , d at ω' , and b at all other states, where $a < b < d$.

Proposition 3. *The only non-constant securities that are separable for all information structures in $\mathcal{P}$ are the A-D and the security that is of the following form. There are values $a < b < d$ such that $X(\omega_a) = a$ and $X(\omega_d) = d$ for two states ω_a, ω_d , and $X(\omega) = b$ for all $\omega \neq \omega_a, \omega_d$.*

See [proof](#) on page 35.

All three types of securities are not very informative. Even when there is information aggregation and the price of X always converges to the true value $X(\omega)$ at state ω , the A-D only reveals whether ω has occurred or not, and the second type only reveals that either ω , ω' , or neither, have occurred. In other words, the price of X does not reveal information about most events in Ω . In contrast, the most informative security X' pays differently across all states. If there is information aggregation, then X' reveals whether any event in Ω has occurred or not. However, because X' is not always separable, we know that for some information structure Π , information does not aggregate.

We now formally state our classification. Let X be a security defined on Ω that has at least four states. If there is m such that $X(\omega') = X(\omega'') = m$ for at least two states $\omega', \omega'' \in \Omega$, let M be the set of states ω with $X(\omega) = m$. Otherwise, set $M = \emptyset$.

Case 0. $M = \Omega$. Security X is constant and therefore (trivially) always separable, for all information structures Π .

Case 1. There is either a unique $\omega \notin M$, so that the security is A-D, or there are exactly two states $\omega, \omega' \notin M$, such that $X(\omega) > m > X(\omega')$. From Proposition 3, X is always separable, for all information structures Π .

Case 2. There exist $\omega, \omega' \notin M$ with either $X(\omega), X(\omega') < m$, or $X(\omega), X(\omega') > m$. From Proposition 2, X is κ non-separable for all κ , for some information structure Π .

Case 3. $M = \emptyset$, hence $X(\omega) \neq X(\omega')$, for all $\omega, \omega' \in \Omega$. From Proposition 1, X is κ separable for some κ , for all information structures Π .

We conclude the section with a few remarks on the properties of κ separable securities. First, separability implies κ separability, for all κ , because there does not exist a prior at which the security is non-separable. Second, non-separability implies κ non-separability for some κ , because for a high enough marginal cost c , no Trader will acquire any information.

Remark 1. *If security X is separable given an information structure $\Pi \in \mathcal{P}$, then it is also κ separable for any cost structure $\kappa \in \mathcal{K}$.*

Remark 2. *If security X is non-separable given an information structure $\Pi \in \mathcal{P}$, then for any cost function K it is $\kappa = (K, c)$ non-separable for some marginal cost of information c .*

Note that if X is $\kappa = (K, c)$ non-separable, then there is a μ for which there is no information acquisition and X is non-separable at μ . If we increase the marginal cost to $c' > c$, then there is still no information acquisition for μ , hence

X is $\kappa' = (K, c')$ non-separable. Conversely, if X is $\kappa = (K, c)$ separable, it is separable for all μ for which there is no information acquisition. If we decrease the marginal cost to $c' < c$, then the set of priors for which there is no information acquisition will shrink and therefore X will be $\kappa' = (K, c')$ separable for all $c' < c$. We, therefore, have the following remark.

Remark 3. *If X is $\kappa = (K, c)$ non-separable (separable), then it is $\kappa' = (K, c')$ non-separable (separable) for all $c' > c$ ($c' < c$).*

5 The Value of the Market

If the cost of information drops significantly, do we even need markets to aggregate information through prices? If each Trader can independently afford all relevant signals, then the average of their predictions would be very close to the outcome of the market, rendering it obsolete. In this section, we argue with an example that this intuition is not correct, because markets can aggregate information long before it becomes economically viable for each Trader to acquire the required information on their own. Hence, markets become even more important in an environment with information acquisition.

To illustrate this point, we compare the market's prediction accuracy to that of a poll where traders simultaneously make a single announcement. The poll's prediction is then calculated as the average of their individual predictions.²¹

Fix a cost function K and a prior μ for which the security is non-separable and then gradually decrease c , the marginal cost of information. For high enough c , the market and the poll are equally accurate. At a threshold $\bar{c}$, the security becomes κ separable and the accuracy of the market increases discontinuously to 1, whereas the accuracy of the poll increases gradually as c decreases further.

²¹Note that there are many ways of improving the accuracy of a poll by aggregating announcements differently [Baron et al., 2014]. In our framework, markets will always be more accurate than polls because more information is disseminated through multiple rounds of announcements, and the value of information is positive. Several papers have examined the two settings in experiments and real-life settings, and the results are mixed. Snowberg et al. [2013] argue that prediction markets are better. Berg et al. [2008] show that the Iowa Electronic Markets were more accurate than 964 polls in predicting the outcomes of five presidential elections between 1988 and 2004. Cowgill and Zitzewitz [2015] show that internal prediction markets in Google and Ford were more accurate than the predictions of professional forecasters. On the other hand, Atanasov et al. [2017], Dana et al. [2019] argue that while prediction markets are more accurate than the simple mean of forecasts from polls, the latter outperform prediction markets when forecasts are aggregated with transformation algorithms or made in teams. Camerer et al. [2016] show that markets are equally accurate with a survey in predicting the replicability of economic experiments.

Recall Example 1 with state space $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, security $X = (0, 1, 2, 3)$, and common prior $\mu_0 = (\frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8})$. Trader 1's partition is $\{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$ and Trader 2's is $\{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}\}$. Security X is non-separable at μ_0 because the expected value of X for both traders is $v = \frac{3}{2}$ at all states, yet there is uncertainty about the value of the security. Traders can acquire signals where the cost is proportional to the expected reduction in entropy relative to the fixed posterior $\mu = 0.5$ for every ω [Shannon, 1948]. We assume that traders are myopic and proper scoring rules are used in both settings, hence each announcement is the expected value of the security given the acquired information.

Figure 5 shows the expected accuracy of markets and polls for prior μ . When the marginal cost of information is high ($c > 4.5$), no Trader acquires any information and the accuracy of the market is equal to that of the poll. As the marginal cost decreases below 4.5, Trader 2 starts acquiring information if the state is either 1 or 4. This implies that his announcement differs across partitions, revealing whether event $\{\omega_1, \omega_4\}$ or $\{\omega_2, \omega_3\}$ is true. Trader 1 combines this with his private information and learns which state is true, thus announcing $v = X(\omega)$. Therefore, a small change in the marginal cost of information allows the market to aggregate information, with a prediction accuracy of 1. In contrast, the poll's accuracy improves gradually as information gets cheaper. This means that, as the cost of information decreases, the prediction accuracy of the market suddenly jumps to 1, whereas the accuracy of the poll gradually increases.




Figure 1: Prediction accuracy ($A(\mu) = E_\mu A(\omega, \mu)$) for markets and polls.

We now provide the formal definitions of a poll and of prediction accuracy. We then state Proposition 4 which characterizes κ separable securities in terms of the market being strictly more accurate than a poll for all priors. Recall that the full state $\phi = (\omega, \iota_1, \iota_2, \dots)$ describes the initial uncertainty $\omega \in \Omega$ and the

randomizations of the players, as well as the signal realisations.

Definition 7. *In a poll there is only one round of announcements and traders are myopic. For each $\phi \in \Phi$ and common prior μ , the prediction of the poll is the average of the myopic predictions, y_i , where each Trader i optimally obtains a signal $\mathcal{R}_i$ and then they all announce simultaneously:*

$$p^p(\phi, \mu) = \frac{\sum_{i=1}^n y_i}{n}.$$

The accuracy of the poll at ϕ and μ is $A^p(\phi, \mu) = 1 - |p^p(\phi, \mu) - X(\phi)|$.

The prediction of the market depends on the equilibrium that is played. Let $\mathcal{L}(\mu)$ be the set of Nash equilibria given μ . We know that in every Nash equilibrium $l \in \mathcal{L}$ and state ϕ , the announcements converge to some price $p^m(\phi, \mu, l)$, which we denote as the prediction of the market. We define the accuracy of the market given ϕ and μ as $A^m(\phi, \mu) = \inf_{l \in \mathcal{L}(\mu)} \{1 - |p^m(\phi, \mu, l) - X(\phi)|\}$, the worst accuracy among all equilibria. The highest possible accuracy is 1, when $p^m(\phi, \mu, l) = X(\phi)$ at all Nash equilibria. The expected accuracy of the market given μ is $A^m(\mu) = E_\mu A^m(\phi, \mu)$, whereas for the poll it is $A^p(\mu) = E_\mu A^p(\phi, \mu)$.

For Proposition 4, we assume that if the support of $\mathcal{R}(\cdot|\omega)$ is not the same for all $\omega \in \Omega$, then the cost of $\mathcal{R}$ is infinite. This implies that it is impossible for a Trader to exclude a state ω from being possible, just by observing a signal realization. This assumption makes the analysis non-trivial; otherwise, traders would be able to exclude states from being possible by acquiring signals with sufficiently low c , so obtaining information from trading behavior would be irrelevant.

Assumption 2. *If $\text{Supp}(\mathcal{R}(\cdot|\omega)) \neq \text{Supp}(\mathcal{R}(\cdot|\omega'))$ for some $\omega, \omega' \in \Omega$, then $K(\mathcal{R}) = +\infty$.*

Recall that a non-degenerate prior μ given X does not assign probability 1 to a unique value of security X . If the security is $\kappa = (K, c)$ separable, then information gets aggregated for all marginal costs $c' \leq c$, which is not true for polls. We, therefore, have the following remark for each $\Pi \in \mathcal{P}$ and $\kappa = (K, c)$.

Remark 4. *Under Assumption 2, if security X is κ separable, then for any non-degenerate μ given X and for all $0 < c' \leq c$, information gets aggregated by the market with cost $\kappa' = (K, c')$ at all Nash equilibria, so that $A^m(\mu) = 1$, but it is not aggregated by the poll, so that $A^p(\mu) < 1$.*

To see this, note first that X is κ' separable, as shown in Remark 3, hence the first part follows directly from Theorem 1, which shows that information gets aggregated in all Nash equilibria. Assumption 2 implies that individual traders cannot buy a signal that reveals the true state with probability 1 in the first round. Because μ is non-degenerate given X , we have that $p^p \neq X(\phi)$.

We can further generalise this result, showing that κ separability is equivalent to the market being strictly more accurate than the poll, for all non-degenerate priors given X .

Proposition 4. *Fix cost structure κ . Under Assumptions 1 and 2, security X is κ separable given $\Pi \in \mathcal{P}$ if and only if $A^m(\mu, l) > A^p(\mu)$ for all non-degenerate priors μ given X .*

See proof on page 38.

To interpret this result, let Δ_X be the set of all non-degenerate priors given X . Suppose we define the value of the market (with security X) to be $V(X) = \min_{\mu \in \Delta_X} [A^m(\mu) - A^p(\mu)]$, the minimum improvement in accuracy by switching from the poll to the market. Then, Proposition 4 implies that X is κ separable if and only if $V(X) > 0$.

6 Concluding Remarks

The paper provides a thorough examination of the interplay between information aggregation and information acquisition. We show that κ separability characterizes information aggregation when the cost of information acquisition is κ and we group securities into three classes. First, the ‘always separable’ securities aggregate information irrespective of who trades or what is their information structure. Second, the securities which are κ separable for some κ and all information structures, aggregate information if the cost is sufficiently low. Finally, there is a small class of securities such that for any κ , each is κ non-separable for some information structure. This means that even if the cost is very close to zero, information may not aggregate. Surprisingly, these three classes are easily distinguishable just by looking at the payoff structure of each security.

An interesting question is whether information aggregates if the security is non-separable but the prior is generic. A security may be non-separable at a particular common prior, hence at that prior the information may not aggregate. However, if we slightly perturb this prior then information may aggregate. By perturbing the common prior slightly, a non-separable security could become separable, and

therefore information could aggregate. [Ostrovsky \[2012\]](#) shows that information gets aggregated in all pure-strategy equilibria with a generic prior, even for non-separable securities. However, in his setting, it is an open question what happens with mixed-strategy equilibria. Intuitively, even if we perturb the initial prior, and given that agents have infinite action spaces, it is not clear whether they will converge to some belief at which the security is separable.

The current paper provides a novel perspective on this issue of genericity, by endogenizing the perturbation of beliefs. If information is not aggregated at time t and all traders agree on the price, a Trader could buy an additional signal, if the cost is not too high, and profit from changing his belief. By examining what happens when the cost is very low, we effectively allow for arbitrarily small perturbations of beliefs to be feasible for traders. A “generic” security pays differently across states and [Proposition 1](#) shows that if the cost is sufficiently low, such a security is κ separable, for all information structures. Therefore, using [Theorem 1](#) we can say that generically (for almost all securities) information gets aggregated, in all mixed-strategy equilibria, thus providing an answer to the open question of [Ostrovsky \[2012\]](#).

A Existence of a Perfect Bayesian Equilibrium

Perfect Bayesian equilibria do not always exist in games with an infinite horizon and infinite actions. In this section, we discretize the action space and show that, because the game is continuous at infinity, a Perfect Bayesian equilibrium always exists. Continuity at infinity is achieved by shortening the time period, t_k , as $k \rightarrow \infty$, so that the discount factor decreases. This is similar to the approach of [Ostrovsky \[2012\]](#) and [Galanis et al. \[2024\]](#).

We first establish that a Perfect Bayesian equilibrium always exists in games with finitely many actions and periods. Then, by adapting the proofs of [Fudenberg and Levine \[1983, 1986\]](#), we approximate the game of infinitely many periods with a sequence of games with finitely many periods and show that the sequence of Perfect Bayesian equilibria in the finite games converges to a Perfect Bayesian equilibrium in the infinite game. Finally, we show that the announcements converge to a limit which is as close as possible to the true value of X , given that only finitely many announcements are available.

A.1 Games with Finitely Many Actions

Let $\mathcal{Y} \subseteq [\underline{y}, \bar{y}]$ be a set of finitely many announcements, evenly distributed on $[\underline{y}, \bar{y}]$, and $\mathcal{E}$ a set of finitely many signals. Each signal $\mathcal{R} \in \mathcal{E}$ has finitely many realizations $T_{\mathcal{R}}$. We denote a game with $k < \infty$ periods and finitely many actions as Γ^k . We denote a game with infinitely many periods and finitely many actions as Γ^∞ . Recall that a game with infinitely many actions and periods is denoted by Γ . A game with finitely many actions (Γ^k or Γ^∞) is a tuple $(\Omega, \Pi, X, \mu, \mathcal{E}, c, K, y_0, \mathcal{Y}, s, \beta)$ where Ω is the state space, Π is the collection of partitions, X is the traded security, μ is the common prior, $\mathcal{E}$ is the finite set of signals, c is the marginal cost and K is the cost function of signals, y_0 is the market maker's initial announcement at time t_0 , $\mathcal{Y}$ is the finite set of possible announcements, s is a strictly proper scoring rule, and β is the common discount rate.

Recall that other traders are not informed about which signal is chosen by Trader i , or of the specific signal realization. Therefore, an information set $\mathcal{I}_k$ for the Trader who announces at t_k only distinguishes between decision nodes that specify a different history of announcements, or a different history of his own signal realizations. Let $\mathcal{J}$ denote the collection of all information sets.

An assessment $\mathcal{A} = \{\sigma, \mathcal{B}\}$ is a strategy profile $\sigma \in \Sigma$ and a system of beliefs $\mathcal{B} = \{\mathcal{B}(\mathcal{I})\}_{\mathcal{I} \in \mathcal{J}}$, where $\mathcal{J}$ is the collection of all information sets $\mathcal{I}$ and $\mathcal{B}(\mathcal{I}) \in$

$\Delta(\Phi)$ is the belief at $\mathcal{I}$. Assessment $\mathcal{A} = \{\sigma, \mathcal{B}\}$ is consistent if the system of beliefs is generated by Bayesian updating, whenever possible, given the common prior μ and the strategy σ .

Definition 8. *The continuation payoff of player $a_k = i$ at time t_k and state ϕ , given strategy profile σ , information set $\mathcal{I}_k$, and system of beliefs $\mathcal{B}$, is*

$$V_i(\mathcal{I}_k, \phi, \sigma, \mathcal{B}) = E_{\mathcal{B}(\mathcal{I}_k)} \left[\sum_{m=0}^{\infty} \beta^{nm} \left(s(y_{k+nm}(\sigma, \phi|\mathcal{I}_k), X(\phi)) - s(y_{k+nm-1}(\sigma, \phi|\mathcal{I}_k), X(\phi)) \right) - cK(\mathcal{R}_{k+nm}(\sigma, \phi|\mathcal{I}_k)) \right].$$

We fix the state ϕ for the remainder of the analysis and omit it to simplify the notation.

Definition 9. *Consistent assessment $(\sigma^*, \mathcal{B})$ is a Perfect Bayesian equilibrium if there is no information set $\mathcal{I}_k$, player $a_k = i$ and alternative strategy $\sigma = (\sigma_i, \sigma_{-i}^*)$, such that*

$$V_i(\mathcal{I}_k, \sigma, \mathcal{B}) > V_i(\mathcal{I}_k, \sigma^*, \mathcal{B}).$$

The existence of a Perfect Bayesian equilibrium in every game Γ^k with finitely many actions and periods follows from standard results, as it is weaker than a sequential equilibrium which always exists (Kreps and Wilson [1982]).

A.2 Truncated Games with Finitely Many Periods

For each game Γ^∞ with infinitely many periods and finitely many actions, we will generate a sequence of truncated games with finitely many periods. Following Fudenberg and Levine [1983, 1986], we assume, without loss of generality, that there is a “do nothing” or null action at each time t , denoted 0. The 0 action means that the Trader repeats, with probability 1, the previous announcement, and buys no signals. Note that the null action guarantees that the agent’s payoff at that period is 0, irrespective of the previous announcement. With each game Γ^∞ , we associate a collection of truncated games. A truncated game $\Gamma(m)$ effectively ends in time t_m , because for all $t_k > t_m$, the only available action for any Trader is the null action of repeating the previous announcement and their beliefs no longer update. Hence, the announcement in t_m is repeated by everyone. Note that a game Γ^m ends at t_m , whereas game $\Gamma(m)$ has infinitely many periods but traders can only choose the null action after t_m .

More formally, let $\Sigma(m)$ be the strategy space of truncated game $\Gamma(m)$. Each strategy profile $\sigma \in \Sigma(m)$ is of the form $(\sigma_1, \sigma_2, \dots, \sigma_m, 0, 0, \dots)$, where σ_k maps

information sets at time t_k to the set of possible (finite) signals and announcements $\mathcal{Y}$, for the Trader who makes an announcement at t_k . Note that $\Sigma(1) \subseteq \Sigma(2) \subseteq \dots \subseteq \Sigma(\infty)$, where $\Sigma(\infty) = \Sigma$ is the strategy space of the infinite horizon game. Moreover, each $\Sigma(m)$ is compact.

Let $\mathcal{A}(m)$ be the collection of assessments $\mathcal{A} = (h, \mathcal{B})$ such that $h \in \Sigma(m)$ and beliefs do not update after m , so that if $k > m$ and information set $\mathcal{I}_{k-1}$ immediately precedes $\mathcal{I}_k$, then $\mathcal{B}(\mathcal{I}_k) = \mathcal{B}(\mathcal{I}_{k-1})$. Let $\mathcal{B}(m)$ be the collection of all systems of beliefs $\mathcal{B}$ of game $\Gamma(m)$ and note that $\mathcal{B}(m)$ is compact.

An assessment $\mathcal{A} = (h, \mathcal{B})$ is an equilibrium in $\Gamma(m)$ only if $\mathcal{A} = (h, \mathcal{B}) \in \mathcal{A}(m)$. Note that $\mathcal{A}(1) \subseteq \mathcal{A}(2) \subseteq \dots \subseteq \mathcal{A}(\infty)$, where $\mathcal{A}(\infty)$ is the collection of all assessments in the infinite horizon game Γ^∞ . Intuitively, after t_m no Trader changes her action, hence there is no information revelation and beliefs do not update. Set $\mathcal{A}(\infty)$ is compact because it is a closed subset of $\times_{m=1}^\infty (\Sigma(m) \times \mathcal{B}(m))$, which is the product of compact sets in the product topology, and therefore also compact.

A.3 Continuity

In order to approximate an equilibrium in the infinite game Γ^∞ we need to define the distance between assessments. Recall that Φ is the space of all uncertainty and let $\delta(p, q) = \sup_{E \in \Phi} |p(E) - q(E)|$ be the distance between two beliefs $p, q \in \Delta(\Phi)$.

We define the distance between two systems of beliefs $\mathcal{B}, \mathcal{B}'$ as

$$d(\mathcal{B}, \mathcal{B}') = \sup_{\mathcal{I}_k \in \mathcal{I}} \left\{ \frac{1}{k} \delta(\mathcal{B}(\mathcal{I}_k), \mathcal{B}'(\mathcal{I}_k)) \right\}.$$

This metric is motivated in [Fudenberg and Levine \[1983\]](#). It specifies that two systems of beliefs are close to each other if they only differ in the distant future. We can similarly measure the distance between two strategies σ, σ' with $d(\sigma, \sigma') = \sup_{\mathcal{I}_k \in \mathcal{I}} \left\{ \frac{1}{k} \delta(\sigma(\mathcal{I}_k), \sigma'(\mathcal{I}_k)) \right\}$, where $\sigma(\mathcal{I}_k) \in \Delta\left(\bigcup_{\mathcal{R} \in \mathcal{E}} \mathcal{Y}^{T_{\mathcal{R}}}\right)$ is the strategy of the Trader who announces at t_k and information set $\mathcal{I}_k$. Finally, we can also measure the distance between truncated systems of beliefs (and strategies) that occur after an information set. Given an information set $\mathcal{I}_k$, let $\mathcal{I}_k(\mathcal{B})$ be the restriction of $\mathcal{B}$ to all information sets that succeed $\mathcal{I}_k$, including $\mathcal{I}_k$. The same notation $\mathcal{I}_k(\sigma)$ applies to a strategy σ . Then, the distances $d(\mathcal{I}_k(\mathcal{B}), \mathcal{I}_k(\mathcal{B}'))$ and $d(\mathcal{I}_k(\sigma), \mathcal{I}_k(\sigma'))$ calculate the distance only with respect to information sets that follow $\mathcal{I}_k$.

We can now define the distance between two assessments $\mathcal{A} = \{\sigma, \mathcal{B}\}, \mathcal{A}' =$

$\{\sigma', \mathcal{B}'\}$:

$$d(\mathcal{A}, \mathcal{A}') \equiv \sup_{\mathcal{I}_k \in \mathcal{I}} \left\{ d(\mathcal{I}_k(\sigma), \mathcal{I}_k(\sigma')), \left\{ \sup_{h_{a_k} \in \Sigma_{a_k}} d(\mathcal{I}_k(h_{a_k}, \sigma_{-a_k}), \mathcal{I}_k(h_{a_k}, \sigma'_{-a_k})) \right\}, d(\mathcal{I}_k(\mathcal{B}), \mathcal{I}_k(\mathcal{B}')) \right\},$$

where a_k is the announcer at time t_k and information set $\mathcal{I}_k$.

This metric is also motivated in [Fudenberg and Levine \[1983\]](#). Two assessments $\mathcal{A}$ and $\mathcal{A}'$ are close to each other if the following three conditions are true at each information set $\mathcal{I}_k$. First, the distributions over actions that are generated by σ and σ' for every subsequent information set are close to each other. Second, the distributions over actions are close to each other even when Trader a_k deviates with h_{a_k} at $\mathcal{I}_k$. Finally, the sequence of sets of beliefs that are generated from $\mathcal{B}, \mathcal{B}'$ given $\mathcal{I}_k$ are also close to each other.

Given an assessment $\mathcal{A} = \{\sigma, \mathcal{B}\}$, let $V_i(\mathcal{I}, \sigma, \mathcal{B})$ be i 's continuation payoff at information set $\mathcal{I}$. We say that the game is uniformly continuous if whenever two assessments are close to each other, the continuation payoffs are also close, for each information set.

Definition 10. *The game Γ^∞ is uniformly continuous if for all information sets $\mathcal{I}$ and all sequences of assessments $\{\mathcal{A}^n\} = \{\sigma^n, \mathcal{B}^n\}$, $\{\mathcal{A}'^n\} = \{\sigma'^n, \mathcal{B}'^n\}$, $\mathcal{A}^n \rightarrow \mathcal{A}'^n$ implies $|V_i(\mathcal{I}, \sigma^n, \mathcal{B}^n) - V_i(\mathcal{I}, \sigma'^n, \mathcal{B}'^n)| \rightarrow 0$ for all $i \in I$.*

[Ostrovsky \[2012\]](#) generates a uniformly continuous game by shortening the time period t_m as $m \rightarrow \infty$. For example, in his version of the [Kyle \[1985\]](#) model, he sets $t_k = 1 - \frac{1}{2^k}$. Although he does not specify how he achieves uniform continuity in the prediction market model, we use something similar. Following [Dimitrov and Sami \[2008\]](#), we assume that the payment at t_m is $\beta^m(s(y_{t_m}, x^*) - s(y_{t_{m-1}}, x^*))$, instead of $s(y_{t_m}, x^*) - s(y_{t_{m-1}}, x^*)$. This is equivalent to shortening the period t_m as $m \rightarrow \infty$, so that the discounting factor needs to decrease accordingly. Although the results in the main text are true with both definitions, the former ensures that the game is uniformly continuous. We refer to this uniformly continuous game with finitely many announcements as Γ^∞ .

Let $\mathcal{B}(\infty)$ be the collection of all systems of beliefs. Let constant w^m be the greatest variation in any agent's payoff and for any system of beliefs, from strategies σ, τ that are identical for all information sets up to time t_{m-1} , written as $\sigma =_{m-1} \tau$.

$$w^m \equiv \sup_{\substack{\mathcal{I} \in \mathcal{I} \\ \mathcal{B} \in \mathcal{B}(\infty)}} \sup_{\substack{i \in I \\ \sigma =_{m-1} \tau}} \left| V_i(\mathcal{I}, \sigma, \mathcal{B}) - V_i(\mathcal{I}, \tau, \mathcal{B}) \right|.$$

Definition 11. *The game Γ^∞ is continuous at infinity if $w^m \rightarrow 0$ as $m \rightarrow \infty$.*

It is straightforward that because Γ^∞ is uniformly continuous, it is also continuous at infinity.

A.4 Existence of Equilibrium in Γ^∞

Following [Fudenberg and Levine \[1983, 1986\]](#), we prove existence in the infinite game Γ^∞ through a series of lemmas. We first define the notion of an ϵ -Perfect Bayesian equilibrium.

Definition 12. *Consistent assessment $(\sigma^*, \mathcal{B})$ is an ϵ -Perfect Bayesian equilibrium if there is no information set $\mathcal{I}_k$, player $a_k = i$ and alternative strategy $\sigma = (\sigma_i, \sigma_{-i}^*)$, such that*

$$V_i(\mathcal{I}_k, \sigma, \mathcal{B}) > V_i(\mathcal{I}_k, \sigma^*, \mathcal{B}) + \epsilon.$$

Lemma 1. *If $(h^*, \mathcal{B}) \in \mathcal{A}(m)$ is an ϵ -Perfect Bayesian equilibrium in $\Gamma(m)$, then $(h^*, \mathcal{B})$ is an $(\epsilon + w^m)$ -Perfect Bayesian equilibrium in Γ^∞ .*

Proof. Suppose $(h^*, \mathcal{B})$ is an ϵ -Perfect Bayesian equilibrium in $\Gamma(m)$ and let $g \in \Sigma(\infty)$. Set $h = (g_1, g_2, \dots, g_m, 0, \dots)$. For any information set $\mathcal{I} \in \mathcal{I}$ where agent i makes an announcement, we have

$$V_i(\mathcal{I}, h_i, h_{-i}^*, \mathcal{B}) - V_i(\mathcal{I}, h^*, \mathcal{B}) \leq \epsilon.$$

Because h and g differ only after t_m , we have

$$V_i(\mathcal{I}, g_i, h_{-i}^*, \mathcal{B}) - V_i(\mathcal{I}, h_i, h_{-i}^*, \mathcal{B}) \leq w^m.$$

Adding the two inequalities we have

$$V_i(\mathcal{I}, g_i, h_{-i}^*, \mathcal{B}) - V_i(\mathcal{I}, h^*, \mathcal{B}) \leq \epsilon + w^m.$$

For any information set $\mathcal{I}_k$ where $t_k \leq t_m$, the consistency on beliefs is satisfied because $(h^*, \mathcal{B})$ is an ϵ -Perfect Bayesian equilibrium in $\Gamma(m)$. For any other information set, the consistency on beliefs is also satisfied because everyone chooses the null action of repeating the previous announcement, hence there is no updating of information or beliefs. Because g was arbitrary, $(h^*, \mathcal{B})$ is an $(\epsilon + w^m)$ -Perfect Bayesian equilibrium in Γ^∞ . □

Lemma 2. Consider a sequence of assessments $\{\mathcal{A}_m\} = \{g_m, \mathcal{B}_m\}$ such that each $\mathcal{A}_m$ is an ϵ -Perfect Bayesian equilibrium in Γ^∞ and $\mathcal{A}_m \rightarrow \mathcal{A} = \{g, \mathcal{B}\}$. Then, $\mathcal{A}$ is an ϵ -Perfect Bayesian equilibrium in Γ^∞ .

Proof. First, note that since $\mathcal{A}_m \rightarrow \mathcal{A} = (g, \mathcal{B})$, the corresponding sets of beliefs $\mathcal{B}_m$ converge to $\mathcal{B}$, hence the consistency condition on beliefs is satisfied. Suppose there is an information set $\mathcal{I}$ and a strategy h_i such that

$$V_i(\mathcal{I}, h_i, g_{-i}, \mathcal{B}) - V_i(\mathcal{I}, g, \mathcal{B}) \geq \epsilon + 3\delta.$$

Because $\mathcal{A}_m \rightarrow \mathcal{A}$ and the game is uniformly continuous, for any δ , there exists large m such that

$$V_i(\mathcal{I}, g_m, \mathcal{B}_m) - V_i(\mathcal{I}, g, \mathcal{B}) < \delta,$$

$$V_i(\mathcal{I}, h_i, g_{-i}, \mathcal{B}) - V_i(\mathcal{I}, h_i, g_{m-i}, \mathcal{B}_m) < \delta.$$

Combining the three inequalities, we have

$$V_i(\mathcal{I}, h_i, g_{-i}, \mathcal{B}) - V_i(\mathcal{I}, g, \mathcal{B}) \geq \epsilon + 3\delta > \epsilon + \delta + V_i(\mathcal{I}, g_m, \mathcal{B}_m) - V_i(\mathcal{I}, g, \mathcal{B}) +$$

$$V_i(\mathcal{I}, h_i, g_{-i}, \mathcal{B}) - V_i(\mathcal{I}, h_i, g_{m-i}, \mathcal{B}_m),$$

which implies

$$V_i(\mathcal{I}, h_i, g_{m-i}, \mathcal{B}_m) - V_i(\mathcal{I}, g_m, \mathcal{B}_m) > \epsilon + \delta.$$

As δ can be taken to be arbitrarily small, we can find big enough m so that

$$V_i(\mathcal{I}, h_i, g_{m-i}, \mathcal{B}_m) - V_i(\mathcal{I}, g_m, \mathcal{B}_m) > \epsilon,$$

which contradicts that $\mathcal{A}_m$ is an ϵ -Perfect Bayesian equilibrium. $\square$

Lemma 3. Suppose that there is a sequence $\{\mathcal{A}_m\} = \{g_m, \mathcal{B}_m\}$ such that $\mathcal{A}_m$ is an ϵ_m -Perfect Bayesian equilibrium in $\Gamma(m)$ and, as $m \rightarrow \infty$, we have $\epsilon_m \rightarrow 0$ and $\mathcal{A}_m \rightarrow \mathcal{A}^* = (g^*, \mathcal{B}^*)$. Then, $\mathcal{A}^*$ is a Perfect Bayesian equilibrium in Γ^∞ .

Proof. From Lemma 1, $\mathcal{A}_m$ is an $(\epsilon_m + w^m)$ -Perfect Bayesian equilibrium in the infinite game Γ^∞ . Because Γ^∞ is uniformly continuous, it is also continuous at infinity and, together with $\epsilon_m \rightarrow 0$, we have $\epsilon_m + w^m \rightarrow 0$. This implies that for each $\delta > 0$ there is M such that $\epsilon_m + w^m < \delta$, whenever $m > M$. From Lemma 2, $\mathcal{A}^*$ is a δ -Perfect Bayesian equilibrium in the infinite game Γ^∞ . Since this is true for every $\delta > 0$, $\mathcal{A}^*$ is a Perfect Bayesian equilibrium in Γ^∞ . $\square$

Proposition 5. *There is a Perfect Bayesian equilibrium in Γ^∞ .*

Proof. Each finite-horizon Γ^m and therefore $\Gamma(m)$ has a Perfect Bayesian equilibrium $\mathcal{A}_m$, for any $m < \infty$. From Lemma 1, $\mathcal{A}_m$ is a w^m -Perfect Bayesian equilibrium in Γ^∞ . Since $\mathcal{A}(\infty)$ is compact, there is a subsequence $\mathcal{A}_k \subseteq \mathcal{A}_m$ with $\mathcal{A}_k \rightarrow \mathcal{A}^*$. From Lemma 3, $\mathcal{A}^*$ is a Perfect Bayesian equilibrium in Γ^∞ . $\square$

A.5 Convergence of Announcements

When $\mathcal{Y}$ is finite, the announcements in equilibrium $\mathcal{A}^*$ do not necessarily converge to the true value $X(\omega)$, for all states, because $X(\omega)$ may not belong to $\mathcal{Y}$. We argue here that, nevertheless, they are as close as possible.²²

Proposition 5 shows that the subsequence of Perfect Bayesian equilibria $\mathcal{A}_k = \{g_k, \mathcal{B}_k\}$, one for each truncated game $\Gamma(k)$, converges to a Perfect Bayesian equilibrium $\mathcal{A}^* = (g^*, \mathcal{B}^*)$ in Γ^∞ . Given that each $\Gamma(k)$ has effectively finitely many periods, the last Trader to announce makes the myopic best announcement, given the available information. Moreover, a proper scoring rule is ‘order-sensitive’ so that the further away the forecast is from the true expected value, the lower is the expectation of the score (see footnote 24). This implies that the myopic strategy in the last period is to announce, with probability one, the element in $\mathcal{Y}$ which is closest to the expected value of X . Given that $\mathcal{A}_k \rightarrow \mathcal{A}^*$, so that both beliefs and strategies converge to $\mathcal{A}^*$, we have that the sequence of announcements in $\mathcal{A}^*$ converge to a limit which is as close as possible to the true expected value of X , given the limit of the beliefs in $\mathcal{B}^*$.²³ Finally, as we increase the number of announcements which are evenly distributed in $\mathcal{Y}$, the distance between the limit of announcements and the true value $X(\omega)$ decreases.

B Proofs

Proof of Theorem 1. We follow the proof of Theorem 1 in Ostrovsky [2012], which proceeds in four steps. In the first step, we show that there is a uniform lower bound on the expected profits that at least one Trader can make by improving the forecast.

²²Note also that, as shown in Example 1, signals with finitely many realizations can generate κ separable securities.

²³The proof of Theorem 1 in Ostrovsky [2012], which we use for the proof of our Theorem 1, shows that in all equilibria, the updating of beliefs is a martingale, hence through the martingale convergence theorem the beliefs of all traders converge to a limit which incorporates all available information, given the equilibrium strategies.

Let r be a distribution over Ω and z be the previous announcement. Following [Ostrovsky \[2012\]](#), we define the instant opportunity of Trader i as the highest expected payoff that he can achieve by changing the forecast from z , if the state is drawn according to r . The difference from [Ostrovsky \[2012\]](#) is that we allow the Trader to acquire information before making an announcement. Let $\mathcal{R}_\omega$ be the optimal signal that Trader i acquires at $\Pi_i(\omega)$, given his beliefs $r_{\Pi_i(\omega)}$, and Q_ω the resulting random posterior. Trader i 's instant opportunity is

$$\sum_{\omega \in \Omega} r(\omega) \left(\sum_{\gamma \in \text{Supp}(Q_\omega)} Q_\omega(\gamma) \sum_{\omega' \in \Omega} \gamma(\omega') \left[s(E_\gamma[X], X(\omega')) - s(z, X(\omega')) \right] - cK(\mathcal{R}_\omega) \right).$$

Let Δ be the set of probability distributions $r \in \Delta(\Omega)$ for which there is uncertainty about security X , so that there are states a, b with $r(a) > 0$, $r(b) > 0$ and $X(a) \neq X(b)$.

Lemma 4. *If security X is κ separable, then for every $r \in \Delta$, there exist $\chi > 0$ and Trader i such that, for every $z \in \mathbb{R}$, the instant opportunity of Trader i given r and z is greater than χ .*

Proof. Fix $r \in \Delta$. There are two cases. First, X is separable with respect to r . Then, the proof of Lemma 1 in [Ostrovsky \[2012\]](#) applies and we have the result. Second, X is non-separable at r and value v , where $v = E_r[X|\Pi_i(\omega)]$ for all $\omega \in \text{Supp}(r)$ and $i \in I$. Because X is κ separable, there is information acquisition. This implies that for some $\omega \in \text{Supp}(r)$, Trader i receives a strictly positive payoff by acquiring information and changing the previous announcement $v = E_r[X|\Pi_i(\omega)]$. Note that the new announcement is not deterministic but depends on the optimal random posterior Q , so he announces $E_\gamma[X]$ with probability $Q(\gamma)$. From continuity, there is a small enough $\epsilon > 0$, so that for all previous announcements $z \in [v - \epsilon, v + \epsilon]$, Trader i receives a strictly positive payoff of at least $\chi_1 > 0$ by acquiring information and changing the announcement. If $z \notin [v - \epsilon, v + \epsilon]$, then Trader i receives a strictly positive payoff of at least $\chi_2 > 0$, by not acquiring any information and announcing the myopically best $E_r[X|\Pi_i(\omega)] = v$.²⁴ By setting $\chi_3 = \min\{\chi_1, \chi_2\}$, we have that at $\Pi_i(\omega)$, Trader i receives a strictly positive payoff of at least χ_3 , for all previous announcements z . For any $\omega' \in \text{Supp}(r) \setminus \Pi_i(\omega)$, Trader i can repeat v and receive 0. Hence,

²⁴ This is true because a proper scoring rule is ‘order-sensitive’ so that the further away the previous announcement is from the true expected value $E_r[X|\Pi_i(\omega)] = v$, the higher the payoff is for Trader i . The lowest payoff from the myopically best announcement is 0 and it is achieved when it is equal to the previous announcement (see p. 2618 in [Ostrovsky \[2012\]](#)).

the instant opportunity of i given r , which is the ex-ante expectation over all $\omega \in \text{Supp}(r)$, is at least $\chi = r(\Pi_i(\omega))\chi_3 > 0$ for all previous announcements. $\square$

Steps 2-3 are identical to the proof of [Ostrovsky \[2012\]](#) and establish that there is Trader i and lower bound $\nu^* > 0$ such that, for infinitely many periods t_{nk+i} , the expected instant opportunity of Trader i is greater than ν^* . The reason they are identical is that, once we fix an equilibrium strategy, the only difference from [Ostrovsky \[2012\]](#) is that traders buy an extra, payoff-irrelevant, signal in each period where they make an announcement.

In Step 4, we show that the presence of a “non-vanishing arbitrage opportunity” is impossible in equilibrium. The proof is very similar to that of [Ostrovsky \[2012\]](#), but we write it here for completeness.

Let $\bar{s}_k$ be the expected score of prediction y_k , where the expectation is over all ϕ and the moves of players according to the mixed equilibrium. The expected payoff to the Trader who moves in period t_k is $\beta^k(\bar{s}_k - \bar{s}_{k-1} - \bar{c}_k)$, where $\bar{c}_k \geq 0$ is the expected cost of information acquisition. It is zero if and only if there is no information acquisition.

Take any period t_k . Let Ψ_k be the sum of all players’ expected continuation payoffs at t_k , divided by β^k :

$$\Psi_k = (\bar{s}_k - \bar{s}_{k-1} - \bar{c}_k) + \beta(\bar{s}_{k+1} - \bar{s}_k - \bar{c}_{k+1}) + \beta^2(\bar{s}_{k+2} - \bar{s}_{k+1} - \bar{c}_{k+2}) + \dots$$

Note that Ψ_k is weakly positive, because each Trader can guarantee a zero payoff by repeating the previous announcement and not buying any costly information. Additionally, it is strictly positive if i ’s expected instant opportunity is strictly positive and it is i ’s turn to make an announcement. That is, with some probability, some history H^{k-1} occurs and i ’s instant opportunity is strictly positive.

Consider now $\lim_{K \rightarrow \infty} \sum_{k=1}^K \Psi_k$. On the one hand, this limit must be infinite, because each Ψ_k is non-negative and an infinite number of them are greater than $\nu^* > 0$,

from Step 3. On the other hand, for any K , we have

$$\begin{aligned}
\sum_{k=1}^K \Psi_k &= (\bar{s}_1 - \bar{s}_0 - \bar{c}_1) + \beta(\bar{s}_2 - \bar{s}_1 - \bar{c}_2) + \beta^2(\bar{s}_3 - \bar{s}_2 - \bar{c}_3) + \dots \\
&\quad + (\bar{s}_2 - \bar{s}_1 - \bar{c}_2) + \beta(\bar{s}_3 - \bar{s}_2 - \bar{c}_3) + \beta^2(\bar{s}_4 - \bar{s}_3 - \bar{c}_4) + \dots \\
&\quad + \quad \vdots \\
&\quad + (\bar{s}_K - \bar{s}_{K-1} - \bar{c}_K) + \beta(\bar{s}_{K+1} - \bar{s}_K - \bar{c}_{K+1}) + \beta^2(\bar{s}_{K+2} - \bar{s}_{K+1} - \bar{c}_{K+2}) + \dots \\
&= (\bar{s}_K - \bar{s}_0 - \sum_{k=1}^K \bar{c}_k) + \beta(\bar{s}_{K+1} - \bar{s}_1 - \sum_{k=2}^{K+1} \bar{c}_k) + \beta^2(\bar{s}_{K+2} - \bar{s}_2 - \sum_{k=3}^{K+2} \bar{c}_k) + \dots \\
&\leq (\bar{s}_K - \bar{s}_0) + \beta(\bar{s}_{K+1} - \bar{s}_1) + \beta^2(\bar{s}_{K+2} - \bar{s}_2) + \dots \\
&\leq 2M/(1 - \beta),
\end{aligned}$$

where costs drop out as they are negative and $M = \max_{y \in [\underline{y}, \bar{y}], \omega \in \Omega} |s(y, X(\omega))|$. But this is impossible, hence y_k must converge in probability to the true value of X .

For the second part, suppose that X is κ non-separable. Then, there exist μ and v at which there is no information acquisition and X is non-separable at μ . Consider the following strategy profile. In every period and after every history, each Trader repeats the previous announcement. The initial announcement by the market maker is v .²⁵

We will show that this strategy profile is a Nash equilibrium. Note that everyone's expected utility is 0 from following the Nash equilibrium strategy. Suppose that Trader i has a profitable deviation, hence his expected utility, the sum from all the periods where he is making an announcement, is strictly positive. From the definition of κ non-separability, the expected utility of Trader i at the first time t_1 where he deviates is weakly negative, even if he chooses the optimal experiment and makes the myopic best announcement. Because the sum over all periods is strictly positive, there is time $t_{k'}$ where the expected utility at $t_{k'}$ from deviating is strictly positive.

We say that Trader i moves the announcement at time t_k if he does not repeat the announcement that Trader $i - 1$ made at t_{k-1} . Suppose that at time $t_{k'}$ he moves the announcement and let t_k be the previous time where he also moved the announcement (hence $t_{k'}$ cannot be the first time he deviates). We will show that i 's total expected utility strictly increases if, instead of moving at t_k and at $t_{k'}$, he combines the two announcements and only moves at t_k . Intuitively, if at t_k he

²⁵The initial announcement by the market maker is irrelevant, as long as Trader 1's first announcement is v and then all other traders repeat the previous announcement.

moves the announcement from v to v' , and at $t_{k'}$ he moves it from v' to v'' , he will be strictly better off if he only moves the announcement from v to v'' at t_k . The reason is that the strictly positive expected utility at $t_{k'}$ is moved earlier at t_k , hence it is discounted with β^{t_k} , which is larger than $\beta^{t_{k'}}$. Recall that all other traders always repeat the previous announcement.

At time t_k , suppose that the previous announcement is $y_{t_{k-1}}$, Trader i has initial belief μ_k and purchases experiment $\mathcal{R}_{\mu_k}$ (for simplicity, we also denote the resulting random posterior as $\mathcal{R}_{\mu_k}$).²⁶ After each signal realization and corresponding posterior $\gamma_k \in \text{Supp}(\mathcal{R}_{\mu_k})$, he makes an announcement y_{γ_k} . His expected utility at t_k , given belief μ_k , is

$$\beta^{t_k} \mu_k \left(\sum_{\gamma_k \in \text{Supp}(\mathcal{R}_{\mu_k})} \mathcal{R}_{\mu_k}(\gamma_k) E_{\gamma_k} [s(y_{\gamma_k}) - s(y_{t_{k-1}})] - cK(\mathcal{R}_{\mu_k}) \right), \quad (1)$$

where we write $\sum_{\omega' \in \Omega} \gamma_k(\omega') [s(y_{\gamma_k}, X(\omega')) - s(y_{t_{k-1}}, X(\omega'))]$ as $E_{\gamma_k} [s(y_{\gamma_k}) - s(y_{t_{k-1}})]$.

At time $t_{k'}$, given each posterior $\gamma_k \in \text{Supp}(\mathcal{R}_{\mu_k})$, he buys another experiment $\mathcal{R}_{\gamma_k}$ and for each posterior $\gamma_{k'} \in \text{Supp}(\mathcal{R}_{\gamma_k})$ he makes an announcement $y_{\gamma_{k'}}$, when the previous announcement is y_{γ_k} (since all other traders keep repeating y_{γ_k} from t_k to $t_{k'}$). His expected utility at $t_{k'}$, given belief μ_k , is

$$\beta^{t_{k'}} \mu_k \left(\sum_{\gamma_k \in \text{Supp}(\mathcal{R}_{\mu_k})} \mathcal{R}_{\mu_k}(\gamma_k) \left(\sum_{\gamma_{k'} \in \text{Supp}(\mathcal{R}_{\gamma_k})} E_{\gamma_{k'}} [s(y_{\gamma_{k'}}) - s(y_{\gamma_k})] - cK(\mathcal{R}_{\gamma_k}) \right) \right). \quad (2)$$

If (2) is strictly positive, Trader i could increase his expected utility by bringing forward these announcements, from $t_{k'}$ to t_k . The sum of the utilities would be

$$\begin{aligned} & \beta^{t_k} \mu_k \left(\sum_{\gamma_k \in \text{Supp}(\mathcal{R}_{\mu_k})} \mathcal{R}_{\mu_k}(\gamma_k) \left[E_{\gamma_k} [s(y_{\gamma_k}) - s(y_{t_{k-1}})] + \right. \right. \\ & \quad \left. \left(\sum_{\gamma_{k'} \in \text{Supp}(\mathcal{R}_{\gamma_k})} E_{\gamma_{k'}} [s(y_{\gamma_{k'}}) - s(y_{\gamma_k})] - cK(\mathcal{R}_{\gamma_k}) \right) \right] - cK(\mathcal{R}_{\mu_k}) \right) = \\ & \beta^{t_k} \mu_k \left(\sum_{\gamma_k \in \text{Supp}(\mathcal{R}_{\mu_k})} \mathcal{R}_{\mu_k}(\gamma_k) \left[\left(\sum_{\gamma_{k'} \in \text{Supp}(\mathcal{R}_{\gamma_k})} E_{\gamma_{k'}} [s(y_{\gamma_{k'}}) - s(y_{t_{k-1}})] - cK(\mathcal{R}_{\gamma_k}) \right) \right] - cK(\mathcal{R}_{\mu_k}) \right). \end{aligned} \quad (3)$$

²⁶Note that μ_k is the belief at t_k after receiving a signal realization at the immediately previous time when he bought an experiment. If he has never bought any experiment before, it is just μ .

If (2) is strictly positive, then (3) is strictly greater than (1) + (2). Moreover, Trader i does not have to make two announcements at t_k . Instead of moving the announcement from $y_{t_{k-1}}$ to y_{γ_k} and then to $y_{\gamma_{k'}}$, he can just move the announcement from $y_{t_{k-1}}$ to $y_{\gamma_{k'}}$. From Assumption 1, he is able to do this with lower cost.

It may be that (2), Trader i 's expected utility at $t_{k'}$ given belief μ_k , is not positive. However, we have assumed that the ex ante utility at $t_{k'}$ (the expectation over all μ_k that arise given all signal realizations before t_k) is strictly positive. This means that bringing forward all announcements from $t_{k'}$ to t_k will strictly increase i 's ex ante expected utility.

Because the sum of expected utilities over all periods is strictly positive, there is finite time t_{k_T} such that the sum of expected utilities over all periods up to t_{k_T} is also strictly positive. Without loss of generality, the expected utility at t_{k_T} is strictly positive, otherwise we can consider an earlier t_{k_T} and the total sum will still be strictly positive. Consider the following process which starts at t_{k_T} and ends at t_1 , the time where Trader i deviates initially. If the expected utility at t_k , $1 < k \leq T$, is negative, remove it from the sum. Then, the remaining sum increases and therefore remains strictly positive. If the expected utility at t_k is positive, bring forward the announcements to t_{k-1} . As described above, the sum of expected utilities increases and therefore it remains strictly positive. As there are finitely many steps, we will reach the first period t_1 , where Trader i starts his deviation. From this process, the expected utility at t_1 is strictly positive. However, this contradicts the fact that at the first time where Trader i deviates, his expected utility is weakly negative, even if he always makes the myopically best announcement, because the security is κ non-separable. Hence, there is no profitable deviation.

□

Proof of Proposition 1. Given Π and K , we will show that if there is no information acquisition at some μ and for sufficiently low c , then traders cannot agree on the announcement, hence the security is κ separable at μ . When the marginal cost of information acquisition c is sufficiently low, and given that X pays differently across all states, the fact that K is sufficiently smooth around zero information (Condition (ii) in Definition 1) implies that there is no information acquisition only if all posterior beliefs assign a probability close to 1 to some state $\omega' \in \Pi_i(\omega)$, for all $i \in I$ and $\omega \in \text{Supp}(\mu)$. Suppose that there exists μ with $E_\mu[X|\Pi_i(\omega)] = v$ for all $i = 1, \dots, n$ and $\omega \in \text{Supp}(\mu)$, where v is arbitrarily close to some m with $X(\omega) = m$.

For the security to be κ non-separable, there should also be uncertainty about its value. This means that for some $\omega' \in \text{Supp}(\mu)$, $X(\omega') \neq v$. From the assumption in Section 2.1 that $\bigcap_{i \in I} \Pi_i(\omega) = \{\omega\}$, there exists Trader i such that $m \notin X(\Pi_i(\omega'))$, so that he considers m impossible. Because he assigns a probability close to 1 to some state, from continuity it must be $E_\mu[X|\Pi_i(\omega')] \neq v$, a contradiction. As this is true for all μ that assign probability close to 1 to some state $\omega' \in \Pi_i(\omega)$, we have that the security X is κ separable for some c . $\square$

Proof of Proposition 2. Suppose Ω has at least four states and let $E = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, where X pays (a, d, b, d) , with either $a, b < d$ or $a, b > d$. What X pays outside E is irrelevant, because we will only consider priors that have full support on E , in order to show that X is κ non-separable given some Π and for all κ . Consider two traders with the following information structure on E . Trader 1's partition is $\{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ and Trader 2's is $\{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}\}$. We need to show that for any $\kappa = (K, c)$, X is κ non-separable. It is enough to show that, as the marginal cost c converges to 0, there is always a common prior μ on E such that no Trader acquires any information and X is non-separable at μ .

Consider prior $\mu = (p, m, 1 - p - 2m, m)$, where $0 < m < 0.5$ and $0 < p < 1 - 2m$. We specify p such that, given any m , each Trader's expected value of X is the same at all states. In the following equation, the left hand-side computes Trader 1's expected value of X at $\{\omega_1, \omega_2\}$ and Trader 2's at $\{\omega_1, \omega_4\}$, whereas the right hand-side computes Trader 1's expected value at $\{\omega_3, \omega_4\}$ and Trader 2's at $\{\omega_2, \omega_3\}$:

$$a \frac{p}{p+m} + d \frac{m}{p+m} = b \frac{1-p-2m}{1-p-m} + d \frac{m}{1-p-m}. \quad (4)$$

There is always a solution for some $p \in (0, 1 - 2m)$ to this equation. To see this, group the four terms of (4) on one side. If p is very close to zero, we have

$$\begin{aligned} d - b \frac{1-2m}{1-m} - d \frac{m}{1-m} &= \frac{d - 2dm - b(1-2m)}{1-m} = \\ \frac{d(1-2m) - b(1-2m)}{1-m} &= \frac{(d-b)(1-2m)}{1-m} \end{aligned}$$

If p is very close to $1 - 2m$ we have

$$a \frac{1-2m}{1-m} + d \frac{m}{1-m} - d = \frac{a(1-2m) - d(1-2m)}{1-m} = \frac{(a-d)(1-2m)}{1-m}$$

In both cases, $a, b < d$ or $a, b > d$, one expression is strictly positive and the other

is strictly negative, because $1 - m$ and $1 - 2m$ are always positive. By continuity, there exists p at which the expression is zero, hence there is a solution.

Let u be the expected value of X at all states and for all traders. Fix any marginal cost c , which may be very close to 0. We can then choose m very close to 0.5 (and corresponding μ) such that u is very close to d , say $d - \epsilon$, assuming, without loss of generality, that $a, b < d$, otherwise we have $d + \epsilon$ and the rest of the proof proceeds accordingly. By repeating the consensus announcement of $u = d - \epsilon$ and not acquiring any signals, both traders get 0 utility. Suppose that one Trader wants to buy a signal $\mathcal{R}$ with random posterior Q , which, in the case of $X(\omega) = d$, will move his posterior expected value of X from $d - \epsilon$ to $d - \epsilon + \nu$, for some $0 < \nu < \epsilon$. The upper bound of his utility, net of the cost of the signal, is realised at $X(\omega) = d$ and it is $s(d - \epsilon + \nu, d) - s(d - \epsilon, d)$. This is positive because $d - \epsilon + \nu$ is closer to d than the previous announcement, $d - \epsilon$, and the proper scoring rule is order-sensitive (see footnote 24). However, by decreasing ϵ appropriately it can be made as small as needed, and therefore smaller than $cK(\mathcal{R})$ for all $\mathcal{R}$. As the benefit becomes smaller and ϵ decreases, a signal must be more precise and ν will decrease so that $d - \epsilon + \nu < d$, hence the cost of the signal will increase weakly, from property (iii) of Definition 1. Because the benefit converges to zero, at some threshold the cost will be higher than the benefit and no signal $\mathcal{R}$ will be acquired. Hence, no Trader will buy any signal and X is κ non-separable.

For the converse, suppose X is κ non-separable given some Π and for all κ . The contrapositive of Proposition 1 implies that there exist two states $\omega_a \neq \omega_b$ such that $X(\omega_a) = X(\omega_b) = m$. Let M be the set of states ω with $X(\omega) = m$. Given that Ω has at least four states, we have the following cases.

Case 0. $M = \Omega$. Security X is constant and therefore (trivially) always separable, for all information structures Π . This is a contradiction because we have assumed X is κ non-separable given some Π and for all κ .

Case 1. There is either a unique $\omega \notin M$, or there are exactly two states $\omega, \omega' \notin M$, such that $X(\omega) > m > X(\omega')$. From Proposition 3, X is always separable, for all information structures Π . As with Case 0, this is a contradiction. The only other remaining case is that there exist $\omega, \omega' \notin M$ with either $X(\omega), X(\omega') < m$, or $X(\omega), X(\omega') > m$, which concludes the proof. $\square$

Proof of Proposition 3. We start by restating a useful characterization of separable securities by Ostrovsky [2012]. It specifies that X is separable if and only if, for any possible announcement v , we can find numbers $\lambda_i(\Pi_i(\omega))$, for each i and ω ,

such that the sum over all traders has the same sign as the difference of $X(\omega) - v$. Intuitively, for any v and at each ω , all traders “vote” and the sign of the sum of the votes has to agree with the sign of the difference between the value of the security and v .

Proposition 6 (Ostrovsky [2012]). *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 $X(\omega) \neq v$,*

$$(X(\omega) - v) \sum_{i \in I} \lambda_i(\Pi_i(\omega)) > 0.$$

If Ω has up to three states, then all securities are of the two types that we have described or the uninteresting case of a constant security. Hence, without loss of generality, we fix a state space Ω with at least four states and a security X . If X is constant, it is trivially separable. Ostrovsky [2012] shows that an A-D security is always separable.

We now show that X is always separable if it is of the following form. Suppose there are $a < b < c$ such that $X(\omega_a) = a$ and $X(\omega_c) = c$ for two states ω_a, ω_c , whereas $X(\omega) = b$ for all $\omega \neq \omega_a, \omega_c$.²⁷ Using Proposition 6, we need to show that 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 $X(\omega) \neq v$,

$$(X(\omega) - v) \sum_{i \in I} \lambda_i(\Pi_i(\omega)) > 0. \quad (5)$$

If $v \leq a$, then condition (5) is satisfied by setting $\lambda_i(\Pi_i(\omega)) = 1$ for all $i \in I$ and $\omega \in \Omega$. Similarly, if $v \geq c$, we set $\lambda_i(\Pi_i(\omega)) = -1$ for all $i \in I$ and $\omega \in \Omega$. Suppose that $a < v \leq b < c$. For all $i \in I$, set $\lambda_i(\Pi_i(\omega_a)) = -1$ and (5) is satisfied for ω_a . For all ω with $\Pi_i(\omega) \neq \Pi_i(\omega_a)$, set $\lambda_i(\Pi_i(\omega)) = k$, where $k = |I|$ is the number of agents. Because of our assumption that the join of all partitions consists of singleton sets, we have that for each $\omega \neq \omega_a$, there exists i such that $\Pi_i(\omega) \neq \Pi_i(\omega_a)$. This implies that if $X(\omega) - v > 0$, we also have $\sum_{i \in I} \lambda_i(\Pi_i(\omega)) \geq k - (k - 1) > 0$ and (5) is satisfied for ω . Using a symmetric argument, we can show that (5) is satisfied for $a < b \leq v < c$, by setting $\lambda_i(\Pi_i(\omega_c)) = 1$ and $\lambda_i(\Pi_i(\omega)) = -k$ for ω with $\Pi_i(\omega) \neq \Pi_i(\omega_c)$, for all $i \in I$. By applying Proposition 6, security X is always separable.

²⁷Our proof for this type of security also applies to an A-D security. Ostrovsky [2012] used Corollary 1 to show that an A-D security is always separable, however, we cannot use it for this type of security.

Suppose that X is not of the three aforementioned types. Then, we can find four distinct states where X assigns values $a \leq b < c \leq d$. For simplicity, we refer to the state with value a as state a and similarly for b, c , and d .

We will show that X is non-separable for an information structure in $\mathcal{P}$ with two agents. The partition of Trader 1 is $\{\{a, d\}, \{b, c\}\}$ for these four states, whereas for any other state, we have $\Pi_1(\omega) = \{\omega\}$. For Trader 2 it is $\{\{a, c\}, \{b, d\}\}$ and for any other state we have $\Pi_2(\omega) = \{\omega\}$. Hence, the information structure is in $\mathcal{P}$.

To show that X is non-separable, it is enough to find a prior p with support on $\{a, b, c, d\}$ such that, for some v ,

$$(i) \quad X(\omega) \neq v \text{ for some } \omega \in \text{Supp}(p),$$

$$(ii) \quad E_p[X | \Pi_i(\omega)] = v \text{ for } i = 1, 2 \text{ and } \omega \in \text{Supp}(p).$$

Let p_1 be 1's probability of state a conditional on $\{a, d\}$, whereas q_1 is 1's probability of state b conditional on $\{b, c\}$. Let p_2 be 2's probability of state a conditional on $\{a, c\}$, whereas q_2 is 2's probability of state b conditional on $\{b, d\}$. Condition (ii) then translates to the following equations

$$\begin{cases} ap_1 + d(1 - p_1) = bq_1 + c(1 - q_1) \\ ap_1 + d(1 - p_1) = ap_2 + c(1 - p_2) \\ ap_1 + d(1 - p_1) = bq_2 + d(1 - q_2) \\ bq_1 + c(1 - q_1) = ap_2 + c(1 - p_2) \\ ap_1 + d(1 - p_1) = bq_2 + d(1 - q_2) \\ ap_2 + c(1 - p_2) = bq_2 + d(1 - q_2) \end{cases} \quad (6)$$

The posteriors of the two agents can be derived by a common prior p if the following conditions hold:

$$\begin{cases} xp_1 = yp_2 \\ (1 - x)q_1 = (1 - y)q_2 \\ (1 - x)(1 - q_1) = y(1 - p_2) \\ x(1 - p_1) = (1 - y)(1 - q_2) \end{cases} \quad (7)$$

where x is the prior probability of (a, d) and y is the prior probability of (a, c) .

When $a \leq b < c \leq d$, the system (6 - 7) has the following solution:

$$\begin{aligned} q_1 &= \frac{a - c}{a + b - c - d}, \\ p_1 &= \frac{b - d}{a + b - c - d}, \\ p_2 &= \frac{b - c}{a + b - c - d}, \\ q_2 &= \frac{a - d}{a + b - c - d}, \\ x &= p_2, \\ y &= p_1. \end{aligned}$$

These posteriors uniquely define the respective prior probabilities. $\square$

Proof of Proposition 4. If X is κ separable, then from Theorem 1 we have that information aggregates at all states and all Nash equilibria, therefore $A^m(\mu) = 1$. Assumption 2 implies that no Trader would acquire full information in one round, hence $p^p(\phi, \mu) \neq X(\phi)$ for all $\phi \in \Phi$ and $A^p(\mu) < 1$.

For the converse, suppose that X is κ non-separable. Then, we can find non-degenerate μ given X for which there is no information acquisition and X is non-separable at μ . Using the proof of the second part of Theorem 1, which requires Assumption 1, there is a Nash equilibrium where everyone agrees on the announcement at all states in the support of μ and the game ends in the first round. The announcement is the same for everyone, so the poll gives the same prediction as the market, contradicting that $A^m(\mu) > A^p(\mu)$. $\square$

References

- Beth Allen. Generic existence of completely revealing equilibria for economies with uncertainty when prices convey information. *Econometrica*, pages 1173–1199, 1981.
- Pavel Atanasov, Phillip Rescober, Eric Stone, Samuel A. Swift, Emile Servan-Schreiber, Philip Tetlock, Lyle Ungar, and Barbara Mellers. Distilling the wisdom of crowds: Prediction markets vs. prediction polls. *Management Science*, 63(3):691–706, 2017.
- Jonathan Baron, Barbara A Mellers, Philip E Tetlock, Eric Stone, and Lyle H Ungar. Two reasons to make aggregated probability forecasts more extreme. *Decision Analysis*, 11(2):133–145, 2014.
- Joyce E. Berg, Forrest D. Nelson, and Thomas A. Rietz. Prediction market accuracy in the long run. *International Journal of Forecasting*, 24(2):285–300, 2008.
- Alexander W Bloedel and Weijie Zhong. The cost of optimally-acquired information. *Unpublished Manuscript*, 2024.
- Colin F Camerer, Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Johan Almenberg, Adam Altmejd, Taizan Chan, et al. Evaluating replicability of laboratory experiments in economics. *Science*, 351(6280):1433–1436, 2016.
- Yiling Chen, Stanko Dimitrov, Rahul Sami, Daniel M Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, 58(4):930–69, 2010.
- Yiling Chen, Mike Ruberry, and Jennifer Wortman Vaughan. Designing informative securities. Mimeo, 2012.
- Brice Corgnet, Cary Deck, Mark DeSantis, Kyle Hampton, and Erik O Kimbrough. When do security markets aggregate dispersed information? *Management Science*, 2022.
- Bo Cowgill and Eric Zitzewitz. Corporate prediction markets: Evidence from Google, Ford, and Firm X. *The Review of Economic Studies*, 82(4):1309–41, 2015.
- Jason Dana, Pavel Atanasov, Philip Tetlock, and Barbara Mellers. Are markets more accurate than polls? the surprising informational value of “just asking”. *Judgment and Decision Making*, 14(2):135–147, 2019.
- Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

- Eddie Dekel, Drew Fudenberg, and Stephen Morris. Topologies on types. *Theoretical Economics*, 1(3):275–309, 2006.
- Peter DeMarzo and Costis Skiadas. Aggregation, determinacy, and informational efficiency for a class of economies with asymmetric information. *Journal of Economic Theory*, 80(1):123–52, 1998.
- Peter DeMarzo and Costis Skiadas. On the uniqueness of fully informative rational expectations equilibria. *Economic Theory*, 13(1):1–24, 1999.
- Stanko Dimitrov and Rahul Sami. Non-myopic strategies in prediction markets. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 200–9, 2008.
- Maryam Farboodi, Adrien Matray, Laura Veldkamp, and Venky Venkateswaran. Where has all the data gone? *The Review of Financial Studies*, 35(7):3101–3138, 2022.
- Rafael Frongillo, Maneesha Papireddygar, and Bo Waggoner. An axiomatic characterization of CFMMs and equivalence to prediction markets. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, pages 51:1–51:21, 2024.
- Drew Fudenberg and David Levine. Subgame-perfect equilibria of finite-and infinite-horizon games. *Journal of Economic Theory*, 31(2):251–68, 1983.
- Drew Fudenberg and David K. Levine. Limit games and limit equilibria. *Journal of Economic Theory*, 38(2):261–79, 1986.
- Spyros Galanis. Syntactic foundations for unawareness of theorems. *Theory and Decision*, 71(4):593–614, 2011.
- Spyros Galanis. Unawareness of theorems. *Economic Theory*, 52(1):41–73, 2013.
- Spyros Galanis. Dynamic consistency, valuable information and subjective beliefs. *Economic Theory*, 71:1467–1791, 2021.
- Spyros Galanis and Stelios Kotronis. Updating awareness and information aggregation. *B.E. Journal of Theoretical Economics*, 2021.
- Spyros Galanis, Christos A Ioannou, and Stelios Kotronis. Information aggregation under ambiguity: theory and experimental evidence. *Review of Economic Studies*, 91(6):3423–3467, 2024.
- John Geanakoplos and Heraklis Polemarchakis. We can’t disagree forever. *Journal of Economic Theory*, 28:192–200, 1982.
- Benjamin Golub and Matthew O Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.
- Sanford Grossman. On the efficiency of competitive stock markets where trades

- have diverse information. *The Journal of Finance*, 31(2):573–85, 1976.
- Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–19, 2003.
- Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.
- Friedrich August Hayek. The use of knowledge in society. *The American Economic Review*, pages 519–30, 1945.
- Alex Kim, Maximilian Muhn, and Valeri Nikolaev. Bloated disclosures: Can chatgpt help investors process financial information? *arXiv preprint arXiv:2306.10224*, 2023.
- Carlo Kraemer, Markus Nöth, and Martin Weber. Information aggregation with costly information and random ordering: Experimental evidence. *Journal of Economic Behavior & Organization*, 59(3):423–432, 2006.
- David M Kreps and Robert Wilson. Sequential equilibria. *Econometrica: Journal of the Econometric Society*, pages 863–894, 1982.
- Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–35, 1985.
- Jean-Jacques Laffont and Eric S Maskin. The efficient market hypothesis and insider trading on the stock market. *Journal of Political Economy*, 98(1):70–93, 1990.
- Kian-Ping Lim and Robert Brooks. The evolution of stock market efficiency over time: A survey of the empirical literature. *Journal of Economic Surveys*, 25(1):69–108, 2011.
- Richard D. McKelvey and Talbot Page. Public and private information: An experimental study of information pooling. *Econometrica*, 58(6):1321–39, 1990.
- Antonio Mele and Francesco Sangiorgi. Uncertainty, information acquisition and price swings in asset markets. *The Review of Economic Studies*, pages 1533–1567, 2015.
- Michael Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647, 2012.
- Lionel Page and Christoph Siemroth. An experimental analysis of information acquisition in prediction markets. *Games and Economic Behavior*, 101:354–378, 2017.
- Lionel Page and Christoph Siemroth. How much information is incorporated into financial asset prices? experimental evidence. *The Review of Financial Studies*, 34(9):4412–4449, 2021.
- Alessandro Pavan, Savitar Sundaresan, and Xavier Vives. (In) efficiency in infor-

- mation acquisition and aggregation through prices. *Working paper*, 2022.
- Matthias Pelster and Joel Val. Can chatgpt assist in picking stocks? *Finance Research Letters*, 2023.
- Roy Radner. Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica*, 47(3):655–78, 1979.
- Jan Christoph Schlegel, Mateusz Kwaśnicki, and Akaki Mamageishvili. Axioms for constant function market makers. *Available at SSRN*, 2022.
- Claude E Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- Erik Snowberg, Justin Wolfers, and Eric Zitzewitz. Prediction markets for economic forecasting. In *Handbook of economic forecasting*, volume 2, pages 657–687. Elsevier, 2013.
- The Economist. How spooks are turning to superforecasting in the cosmic bazaar. <https://www.economist.com/science-and-technology/2021/04/15/how-spooks-are-turning-to-superforecasting-in-the-cosmic-bazaar>, 2021. Accessed: 04-17-2021.
- Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–26, 2004.