Information Aggregation with Costly Information Acquisition*

Spyros Galanis† Sergei Mikhalishchev‡

This draft: January 16, 2024

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, separability is not robust to small changes in the traders' private information. To remedy this problem, we enhance the model by allowing traders to acquire signals with cost κ, in every period. We show that 'κ separable securities' aggregate information and, as the cost decreases, nearly 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. Finally, even with myopic traders, cheaper information may accelerate or decelerate information aggregation for all but Arrow-Debreu securities.

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

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

*We would like to thank Ehud Lehrer and participants at the CEPR Workshop in Turin, Durham York Workshop in Economic Theory, Durham Economic Theory Conference, EWMES2023 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.

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, showing that information gets aggregated in *all Nash equilibria*, as long as the traded securities are *separable* and trading takes place for infinitely many periods. However, separability is not robust to small changes in the composition of the market or to the traders' information structure. 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. 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]

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]. We first characterize the class of securities which are always separable, irrespective of who trades and what is their information structure. We show that this class is very small and uninformative, as it only contains the Arrow-Debreu (A-D) 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.[2] For any other security, there is a market at which there is no information aggregation so that the security's price does not converge to its true value.

Is it the case that the availability of cheap information makes markets aggregate information more efficiently? To study this question, we enhance the model of

[1]For 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].

[2]These 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.

Ostrovsky [2012] by allowing traders to buy a costly signal structure in each period, before trading. We allow for a large class of information cost functions κ, including the Shannon entropy. We show that, as information acquisition costs decrease, almost any security aggregates information, but not all. We define the class of κ separable securities and show that they are necessary and sufficient for information aggregation, in both strategic and non-strategic environments, when the cost of information is κ. Surprisingly, even when the cost of information converges to zero, a small class of securities never becomes κ separable and therefore may not aggregate information.

As costs decrease, a security switches discontinuously from being non-separable to becoming κ separable, hence even a slight reduction can result in a market, which was previously failing, to aggregate information. This is in contrast to a poll, which is the average of the participants' opinions, because its predictive accuracy improves smoothly as costs decrease. Hence, the availability of cheap information leverages the value of the markets, enabling them to aggregate information long before the cost goes to zero.

Finally, we examine whether information acquisition can make markets more efficient by aggregating information faster. We show that this is not always true. Even in non-strategic environments, as long as the separable security is not A-D, information aggregation can happen both faster and slower, depending on the parameters. This implies that, as the information becomes more affordable, markets could become less efficient if the underlying security is not A-D. This could explain why, as financial technology improves and data becomes more abundant, the information content of firm prices may both decline and increase [Farboodi et al., 2022].

Literature.

Our paper contributes to three strands of literature. The first studies the inefficiency of information acquisition and its effect on information aggregation in markets.[3] 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

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

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. Finally, information aggregation can be delayed when information acquisition is cheap, thus introducing another element of inefficiency.

The second strand looks at the information aggregation properties of financial and, in particular, prediction markets.[4] 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 and Sami [2008] and Chen et al. [2010] examine information aggregation by varying the assumptions regarding the traders' information structure. Galanis et al. [2023] study information aggregation with ambiguity-averse traders, whereas Galanis and Kotronis [2021] allow for boundedly rational traders who are unaware of some contingencies.[5] We contribute to this literature by allowing traders to acquire costly signals at every period and we characterize the κ separable securities which aggregate information.

Finally, the paper contributes to the growing literature on the implications of rational inattention, originated by Sims [2003]. We build on the results of [Denti, 2022, Caplin et al., 2019, Matějka and McKay, 2015, Caplin and Dean, 2015] to characterize the traders' optimal behavior in a game with infinitely many periods,

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

[5]Unawareness 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.

where traders have posterior-separable cost functions and can buy signals in every period. We show that, in any Nash equilibrium, almost any security aggregates information for a sufficiently small marginal cost of information. See also Maćkowiak et al. [2023] for a recent survey of the literature on rational inattention.

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]. 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.[6] 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. 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. In Section 2, we provide two examples that illustrate how information acquisition can enable information aggregation but also impact how fast it is achieved. Section 3 describes the model. Finally, Section 4 outlines and discusses the results.

[6]See Ostrovsky [2012] and Galanis et al. [2023] for examples. Automated market makers are widely used in Decentralized Finance, see Schlegel et al. [2022] for an axiomatization of the logarithmic MSR.

2 Two Examples

In this section, we provide two examples which illustrate aspects of our model. The first demonstrates how the ability to acquire costly signals can result in a security aggregating information. The second shows how information acquisition can impact the speed of information aggregation.

Consider two myopic (non-strategic) traders who have a common prior μ on state space $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$. Trader 1's partition is $\{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ and Trader 2's is $\{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}\}$. This means that when ω_1 is the true state, Trader 1 considers ω_1 and ω_2 to be possible, whereas Trader 2 considers ω_1 and ω_4 to be possible. Note that the traders' pooled information reveals the true state. Hence, if the two traders could truthfully reveal their private information, everyone would learn the true state. A security X has intrinsic value $X(\omega) \in \mathbb{R}$ at each state ω. Agents trade X for infinitely many periods. We say that X aggregates information if the equilibrium price converges to the intrinsic value $X(\omega)$, at each ω and for any common prior μ.

Following Ostrovsky [2012], we model trading using the Market Scoring Rule (MSR) (Hanson [2003, 2007]), where a market maker posts an initial price and then the two traders take turns in announcing a price for X. A scoring rule computes a score that increases as the announced price gets closer to the intrinsic value of the security. The period payoff of the MSR is the difference between the expected scores of the current and the previous announcement. Suppose that the announcement of the previous trader (or the market maker) is y_{-1}. Then, Trader i's utility at ω is $s(y, X(\omega)) - s(y_{-1}, X(\omega))$, where $s(y, X(\omega)) = -(y - X(\omega))^2$ is the quadratic scoring rule, for example. His expected score is $E_p[s(y, X) - s(y_{-1}, X)]$, where p is i's belief and y is the announcement.

All proper scoring rules have the property that the period payoff is maximized by announcing the expected value of X, $E_p[X]$, given the trader's belief, p. In the following two examples, we assume that traders are myopic, hence they do not care about future payoffs. This means that a trader with beliefs p will announce $E_p[X]$.

2.1 Information aggregation

The first example shows how allowing traders to acquire additional information can result in the market aggregating information. The security is $X = (0, \frac{10}{7}, 0, 1)$ and the common prior is $\mu = (\frac{1}{4}, \frac{1}{6}, \frac{1}{4}, \frac{1}{3})$. Suppose that the true state is ω_1 and traders cannot acquire costly signals. In period 1, Trader 1 with conditional beliefs $(3/5, 2/5, 0, 0)$ will announce $\frac{4}{7}$. 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 and, furthermore, 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. In such a case, we say that there is no information aggregation at ω_1 because the final announcement $\frac{4}{7}$ is not equal to the intrinsic value of X at ω_1, which is 0. Security X is non-separable because for some prior (such as μ), all traders announce the same expected value of X, at all states in its support, yes there is still uncertainty about its value. Ostrovsky [2012] shows that information aggregation fails for some prior if and only if the security is non-separable.[7]

Suppose now that traders are able to acquire a noisy signal about the value of the security before making their announcement. For instance, Trader 1 can acquire a statistical experiment $\mathcal{R}$ that generates given each state a signal z with probabilities $\mathcal{R}_{\omega_1}(z) = 1$, and $\mathcal{R}_{\omega_2}(z) = \mathcal{R}_{\omega_3}(z) = \mathcal{R}_{\omega_4}(z) = 0.5$.

At ω_1, Trader 1's prior belief is $(\frac{3}{5}, \frac{2}{5}, 0, 0)$, and after receiving signal z, his posterior belief is $(\frac{3}{4}, \frac{1}{4}, 0, 0)$. He then announces the expected value of X, which is $\frac{5}{14}$. Trader 2 considers states ω_1 and ω_4 to be possible. The public announcement of $\frac{5}{14}$ reveals to Trader 2 that the true state is ω_1. The reason is that, irrespective of whether Trader 1 received signal z or not, his posteriors at ω_4 are $(0, 0, \frac{3}{7}, \frac{4}{7})$ and he would have announced $\frac{4}{7}$. 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 aggregation does occur. In summary, the ability to acquire

[7] Information 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 coud update to some posterior μ', at which there is no information aggregation.

extra signals transforms X from non-separable to separable, enabling information aggregation.

We make two observations. First, we have abstracted from the cost of the signal structure. Each trader will acquire the signal structure only if the expected benefit from making a better prediction outweighs the cost. Second, Trader 2 free rides on Trader 1 buying the signal. By moving the price from $\frac{5}{14}$ to 0, he books a profit, without paying the cost of a signal. This example illustrates that the ability to acquire information can turn a non-separable security into a κ separable.

Security X is κ separable if whenever there is agreement about the expected value of X, given a prior that does not put probability 1 to only one value of X, then at least one (myopic) trader finds it profitable to acquire information. It is therefore a generalization of separability, which requires that such a prior does not exist. In Section 4.2, we formally define the notion of κ separability and Theorem 1 shows that this class of securities characterizes information aggregation when the cost is κ. A natural question is whether all securities eventually become κ separable, for sufficiently low cost of information. Surprisingly, Proposition 1 shows that there is a very small class of securities that never become κ separable, even if the cost is negligible but strictly positive. We, therefore, have a discontinuity as costs converge to zero.

2.2 Speed of aggregation

The second example illustrates how acquiring signals can make information aggregation faster or slower. Suppose that the security is $X = (0, 2, 1, 1)$.

We first show that information acquisition may make information aggregation faster. Suppose that the common prior is $\mu = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ and the true state is ω_1. Initially, traders cannot acquire any information. In the first period, Trader 1 announces 1. This reveals no information to Trader 2 because his private information is $\{\omega_1, \omega_4\}$, and the same announcement would be made by Trader 1 in both states. In period 2, Trader 2 announces 0.5. This reveals to Trader 1 that the true state is ω_1, hence, in period 3 he announces 0 and the game ends.

Suppose now that traders are able to acquire the same signal as in the previous example. When the true state is ω_1, Trader 1 acquires signal z, updates his beliefs to $(\frac{2}{3}, \frac{1}{3}, 0, 0)$ and announces $\frac{2}{3}$. The announcement reveals to Trader 2 that the state is ω_1, because at ω_4 Trader 1 would have announced 1. As a result, Trader 2 announces 0 in period 2 and the game ends, hence information aggregation happens faster with information acquisition.

We now show that the converse is also true, so that information acquisition may delay information aggregation. Suppose that the common prior is $\mu = (\frac{3}{16}, \frac{6}{16}, \frac{7}{32}, \frac{7}{32})$. At ω_1, if there is no information acquisition, Trader 1 announces $\frac{4}{3}$. The announcement reveals to Trader 2 that the state is ω_1, hence Trader 2 announces 0 and the game ends in the second round. If there is information acquisition, Trader 1 acquires the signal z and updates his beliefs to $(\frac{1}{2}, \frac{1}{2}, 0, 0)$, and, therefore, he announces 1. This reveals no information to Trader 2, and he announces 0.5. This reveals to Trader 1 that the true state is ω_1, hence in the third round he announces 0 and the game ends.

Interestingly, if the security is the Arrow-Debreu (A-D), then information acquisition has no impact on the speed of information aggregation as shown in Proposition 6 in Section 4.5. To understand why, consider an A-D security that pays 1 at ω and 0 otherwise. At any period, there are two cases. First, the trader knows that either ω is true, or that it is not true. He then announces 1 or 0 and the game ends. Second, he is unsure about whether ω is true or not. Irrespective of whether he buys a signal structure, he will announce some value $0 < v < 1$. Crucially, other traders already know that his partition cell includes ω, hence his announcement does not reveal any public information and the speed of information aggregation is unaffected.

We conclude by observing that our result relies on the assumption that the cost of acquiring a signal that reveals the true state with probability one is infinite. This property is true for many cost functions, such as entropy, and we maintain it throughout the paper. The only way for a trader to learn that a state ω has not occurred is by hearing an announcement that, according to his reasoning about the announcer's private information and equilibrium behavior, would be impossible if ω was true. However, traders may still buy costly signals when their beliefs become more accurate after updating, which enables them to get more of the surplus by making announcements closer to the true value of the security.

3 The Model

3.1 Preliminaries

Uncertainty is described by a finite state space $\Omega = \{\omega_1, ..., \omega_l\}$ and the set of traders is denoted $I = \{1, \ldots, n\}$. 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. That is $\omega \in \Pi_i(\omega) \in \Pi_i$. 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 $\Pi = \{\Pi_1, \ldots \Pi_n\}$ consists of singleton sets so that $\bigcap\limits_{i \in I} \Pi_i(\omega) = \omega$ for all $\omega \in \Omega$, which means that the traders' pooled information always reveals the true state.[8] 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 $\bigcap\limits_{i \in I} \Pi_i(\omega) = \{\omega\}$ for all $\omega \in \Omega$. Traders have a common prior μ_0 over Ω, having full support, and they are risk-neutral.

3.2 Trading environment

Trading is organized as follows. 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 $X : \Omega \to \mathbb{R}$. 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 $t_\infty \equiv \lim_{k\to\infty} t_k = 1$. All predictions are observed by all traders. Each prediction y_k is required to be within the set $[\min\limits_{\omega\in\Omega} X(\omega), \max\limits_{\omega\in\Omega} X(\omega)]$. At some time $t^* > 1$ the true value $x^* = X(\omega^*)$ of the security is revealed.

The traders' payoffs 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,

[8] This 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.

for any probability measure p and any random variable X, the expectation of s is maximized at $y = E_p[X]$. 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.

3.3 Information acquisition

A crucial addition to the model of Ostrovsky [2012] is that we allow traders to acquire information in every period where they make an announcement by buying statistical experiments.

Let T be a Polish space of possible signals with Borel σ-algebra $\mathcal{T}$. Let $\Delta(T)$ be the set of Borel probabilities on T, with generic element ξ. We endow $\Delta(T)$ with the weak* topology: a sequence of probabilities $\{\xi_n\}$ converges to a probability ξ if for every bounded continuous function $f : T \to \mathbb{R}$, we have $\int f d\xi_n \to \int f d\xi$. Following Blackwell [1951], we model the acquisition of information using statistical experiments. A statistical experiment is a function from states into probabilities on signals, $\mathcal{R} : \Omega \to \Delta(T)$.

Following Bloedel and Zhong [2020], we say that an experiment $\mathcal{R}$ is bounded if (i) the conditional signal distributions $\{\mathcal{R}_\omega\}_{\omega \in \Omega}$ are mutually absolutely continuous, and (ii) there exists a constant $B > 0$ such that the Radon-Nikodym derivatives $\frac{d\mathcal{R}_\omega}{d\mathcal{R}_{\omega'}} \in [1/B, B]$ for all $\omega, \omega' \in \Omega$. In other words, a bounded experiment does not definitively rule out any state and, moreover, has uniformly bounded likelihood ratios. Let $\mathcal{E}$ be the set of all experiments and $\mathcal{E}_b \subseteq \mathcal{E}$ be the collection of all bounded experiments. Note that $\mathcal{E}$ contains experiments that are not bounded. We say that $\mathcal{R}' : \Omega \to \Delta(T')$ is a garbling of $\mathcal{R} : \Omega \to \Delta(T)$ given $E \subseteq \Omega$ if there exists $\psi : T \to \Delta(T')$ such that $\mathcal{R}'_{\omega'}(t') = \int_{t \in T} \psi(t'|t)\mathcal{R}_\omega(t)dt$.

Given a prior belief $\mu \in \Delta(\Omega)$, a statistical experiment $\mathcal{R}$ induces via Bayesian updating a probability distribution over posteriors, or random posterior, $B(\mu, \mathcal{R}) = Q \in \Delta(\Delta(\Omega))$. Let $Q(\gamma)$ be the probability of posterior $\gamma \in \Delta(\Omega)$, $\mathcal{Q}(\mu) = \{B(\mu, \mathcal{R}) : \mathcal{R} \in \mathcal{E}\}$ be the set of all random posteriors that can be generated by some experiment $\mathcal{R} \in \mathcal{E}$ and $\mathcal{Q} = \bigcup\limits_{\mu \in \Delta(\Omega)} \mathcal{Q}(\mu)$. Note that $Q \in \mathcal{Q}(\mu)$ if and only if $\int\limits_{\Delta(\Omega)} \gamma Q(d\gamma) = \mu$. We endow $\mathcal{Q}$ with the weak* topology, so that it is a compact and separable topological space. The subsets of $\mathcal{Q}$ are endowed with the appropriate relative topologies.

Given a full support prior μ, experiment $\mathcal{R}$ is bounded if and only if the induced random posterior $Q = B(\mu, \mathcal{R})$ satisfies $Supp(Q) \subseteq \Delta_\epsilon = \{q \in \Delta(\Omega) : q(\omega) > \epsilon \text{ for all } \omega \in \Omega\}$ for some $\epsilon > 0$. Let $\mathcal{D}_\epsilon = \{Q \in \mathcal{Q} : Supp(Q) \subseteq \Delta_\epsilon\}$. Then, $\mathcal{D}_b = \bigcup\limits_{\epsilon > 0} \mathcal{D}_\epsilon$ denotes the set of random posteriors that are induced by some full support prior and bounded experiment. Denote by δ_μ the degenerate random posterior that puts probability 1 on μ.

Definition 1. *A cost on experiments is a map $h : \Delta(\Omega) \times \mathcal{E} \to [0, \infty]$, where $h(\mu, \cdot)$ is Borel measurable for each prior $\mu \in \Delta(\Omega)$ and has the following properties.*

(i) If $B(\mu, \mathcal{R}) = B(\mu, \mathcal{R}')$ and $\mathcal{R}, \mathcal{R}' \in \mathcal{E}_b$, then $h(\mu, \mathcal{R}) = h(\mu, \mathcal{R}')$,

(ii) If $B(\mu, \mathcal{R}) = \delta_\mu$ and $\mathcal{R} \in \mathcal{E}_b$, then $h(\mu, \mathcal{R}) = 0$,

(iii) Let $\{\mu_j, \mathcal{R}_j\}_{j \in \mathbb{N}}$ be a sequence of experiment-prior pairs inducing random posteriors $Q_j = B(\mu_j, \mathcal{R}_j)$. If $Q_j \to Q^*$ and there exists some $\delta > 0$ such that $\{Q_j\}_{j \in \mathbb{N}} \subseteq \mathcal{D}_\delta$, then $h(\mu_j, \mathcal{R}_j) \to h(\mu^*, \mathcal{R}^*)$, where $Q^* = B(\mu^*, \mathcal{R}^*)$,

(iv) If $\mathcal{R}'$ is a garbling of $\mathcal{R}$ then $h(\mu, \mathcal{R}) \geq h(\mu, \mathcal{R}')$.

Point (i) specifies that if two bounded experiments generate the same random posterior given μ, then they have the same cost. Point (ii) says that a bounded experiment that is completely uninformative has a cost of zero. Point (iii) is a continuity condition that is weaker than weak* continuity, in order to allow for important classes of unbounded cost functions.[9] Finally, point (iv) specifies that more informative experiments are more costly.

The cost on experiments h generates a cost structure $\kappa = (K, c)$ on random posteriors, where $c > 0$ is the unit cost of information and $K : \mathcal{F} \to [0, \infty]$ maps elements of $\mathcal{F} = \{(\mu, Q) : \mu \in \Delta(\Omega), Q \in \mathcal{Q}(\mu)\}$, the set of all priors μ and all posterior distributions Q that can be generated by some experiment $\mathcal{R} \in \mathcal{E}$, to the extended real line.

We assume that the cost $cK(\mu, Q)$ of acquiring random posterior Q, given belief μ, is determined by optimally acquiring the relevant experiment $\mathcal{R}$, so that

$$cK(\mu, Q) \equiv \min\{h(\mu, \mathcal{R}) : B(\mu, \mathcal{R}) = Q\}.$$

Note that Point (i) in Definition 1 ensures that all bounded experiments $\mathcal{R}$ which induce Q given μ have the same cost, hence in that case $cK(\mu, Q) = h(\mu, \mathcal{R})$. If only unbounded experiments generate Q given μ, we have $cK(\mu, Q) = \infty$. Moreover, the cost of learning nothing is zero, so that $K(\mu, Q) = 0$ whenever $Supp(Q) = \{\mu\}$. A cost structure $\kappa = (K, c)$ consists of a cost function K and a unit cost of information $c > 0$. Let $\mathcal{K}$ be the collection of cost structures $\kappa = (K, c)$ that are generated from the experiments.

For Propositions 3 and 4, we assume that it is prohibitively costly to acquire an unbounded experiment.

Assumption 1. *If $\mathcal{R} \notin \mathcal{E}_b$, then $h(\mu, \mathcal{R}) = \infty$.*

If all unbounded experiments have infinite cost, then generating a random posterior with support on a posterior δ_ω, which assigns probability 0 to a state $\omega \in Supp(\mu)$, is infinitely costly.

Corollary 1. *Suppose Assumption 1. Given any $\omega \in Supp(\mu)$ and sequences $\{\mathcal{R}_j\}_{j\in\mathbb{N}}$ with $Q_j = B(\mu, \mathcal{R}_j)$ and $\{\gamma_j\}_{j\in\mathbb{N}}$ with $Q_j(\gamma_j) > \epsilon$ for all j and some $\epsilon > 0$, if $\gamma_j \to \delta_\omega$ then $h(\mu, \mathcal{R}_j) \to \infty$.*

[9] See Bloedel and Zhong [2020] for the motivation behind this continuity condition.

This Corollary implies that for any induced sequences $\{Q_j\}_{j\in\mathbb{N}}$ and $\{\gamma_j\}_{j\in\mathbb{N}}$ with $Q_j(\gamma_j) > \epsilon$ for all j and some $\epsilon > 0$, if $\gamma_j \to \delta_\omega$ then $K(\mu, Q_j) \to \infty$.

A widely used class of functions is posterior-separable cost functions (Caplin et al. [2022]).

Definition 2. *A cost of information function K is posterior-separable if, given $\mu \in \Delta(\Omega)$ and any Bayes-consistent posteriors $Q \in \mathcal{Q}(\mu)$,*

$$K(\mu, Q) = \sum_{\gamma \in Supp(Q)} Q(\gamma) W_\mu(\gamma)$$

for some function $W_\mu : \Delta(Supp(\mu)) \to \overline{\mathbb{R}}$ which is strictly convex and continuous in γ, $T_\mu(\gamma) < \infty$ on $int\Delta(Supp(\mu))$ and $W_\mu(\gamma) \geq 0$.

An example of a posterior-separable cost function that satisfies Assumption 1 is the Shannon cost function,

$$K^S(\mu, Q) = \sum_{\gamma \in Supp(Q)} Q(\gamma) \sum_{\omega \in Supp(\gamma)} \gamma(\omega) \ln \gamma(\omega) - \sum_{\omega \in Supp(\mu)} \mu(\omega) \ln \mu(\omega).$$ [10]

We have described two ways of representing the cost of information acquisition. The first is in terms of costly statistical experiments, whereas the second is in terms of costly random posteriors. There is a growing literature that examines the connection between the two approaches. Bloedel and Zhong [2020] and Hébert and Woodford [2021, 2023] show how costly statistical experiments, with a function h that depends both on the agent's prior and the experiment, can generate uniformly posterior-separable cost functions for random posteriors. Denti et al. [2022], however, restrict h to depend only on the experiment and show that uniformly posterior-separable cost functions cannot be generated.

We do not take a stance on which representation is the most suitable. Our results do not require any specific functional forms, such as posterior separability, however,

[10] Another example is expected Tsallis entropy (Tsallis [1988]). Both are examples of weakly uniformly posterior-separable cost functions, where K depends on μ only through the support of μ (Caplin et al. [2022]). Denti [2022] characterizes posterior-separable and uniformly posterior-separable cost functions. Caplin and Martin [2015] and Caplin and Dean [2015] provide necessary and sufficient conditions for a data set to be represented by optimal choice subject to an information cost.

we adopt this framework for most of the paper because it is the most well-known. When we define the game in Section 4.3, we use the standard framework of costly statistical experiments with function h, allowing the more general form that depends on the agent's evolving beliefs.

3.4 Two trading settings

We examine trading in two settings. In the myopic, or non-strategic, setting, each trader does not care about future payoffs when acquiring information and making an announcement. We denote this setting by $\Gamma^M(\Omega, I, \Pi, X, h, y_0, \mu, \underline{y}, \overline{y}, s)$, where I is the set of n players, s is a strictly proper scoring rule, y_0 is the market maker's initial announcement at time t_0, μ is the common prior, h is the cost of statistical experiments which generates a cost of information κ, $[\underline{y}, \overline{y}]$ is the set of possible announcements.

The strategic setting is studied in Section 4.3. Following Dimitrov and Sami [2008], we focus on the discounted MSR, which specifies that the payment at t_k is $\beta^k(s(y_{t_k}, x^*) - s(y_{t_{k-1}}, x^*))$, where $0 < \beta \leq 1$ is the common discount factor. The total payoff of each trader is the sum of all payments for revisions. We denote this setting by $\Gamma^S(\Omega, I, \Pi, X, h, y_0, \mu, \underline{y}, \overline{y}, s, \beta)$.

3.5 The myopic problem

Suppose that at time t it is Trader i's turn to make an announcement. Having observed all previous announcements and using the resulting 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 $\mu_{\Pi_i(\omega)}$.

His myopic problem consists of buying a random posterior $Q \in \mathcal{Q}(\mu_{\Pi_i(\omega)})$ at cost $cK(\mu_{\Pi_i(\omega)}, Q)$, so that when his posterior beliefs are $\gamma \in Supp(Q)$, he optimally announces $E_\gamma[X]$ because the scoring rule is proper. He, therefore, chooses Q that solves the problem

$$\sup_{Q\in\mathcal{Q}(\mu_{\Pi_i(\omega)})}\left(\sum_{\gamma\in Supp(Q)} Q(\gamma)\sum_{\omega'\in\Omega}\gamma(\omega')\Big[s(E_\gamma[X],X(\omega'))-s(z,X(\omega'))\Big]-cK(\mu_{\Pi_i(\omega)},Q)\right), \tag{1}$$

where $c>0$ is the unit cost of information and z is the previous announcement.

In the standard model of Ostrovsky [2012], without information acquisition, the previous announcement z does not influence the myopic best announcement because the scoring rule is proper. The same is true here, where we allow for information acquisition.[11] To see this, note that we can rewrite the expression in the parenthesis as

$$\sum_{\gamma\in Supp(Q)} Q(\gamma)\sum_{\omega'\in\Omega}\gamma(\omega')\Big[s(E_\gamma[X],X(\omega')))\Big]-\sum_{\omega'\in Supp(\mu)}\mu(\omega')s(z,X(\omega'))-cK(\mu_{\Pi_i(\omega)},Q),$$

hence the previous announcement z does not influence the choice of Q or the announcement.

Note that the trader can always receive a payoff of 0 by just repeating the previous announcement z and not acquiring any new information. If his payoff at (1) is smaller or equal to 0 for all $Q\in\mathcal{Q}(\mu_{\Pi_i(\omega)})$, we say that he does not prefer to acquire information at ω. If, given a common prior μ, no Trader i prefers to acquire information at any state $\omega\in Supp(\mu)$, we say that there is no information acquisition for security X.

Definition 3. *There is no information acquisition for security X, given prior μ and previous announcement z, cost structure $\kappa=(K,c)$ and information structure $\Pi\in\mathcal{P}$, if for all states $\omega\in Supp(\mu)$, all traders i, and all $Q\in\mathcal{Q}(\mu_{\Pi_i(\omega)})$,*

$$\sum_{\gamma\in Supp(Q)} Q(\gamma)\sum_{\omega'\in\Omega}\gamma(\omega')\Big[s(E_\gamma[X],X(\omega'))-s(z,X(\omega'))\Big]-cK(\mu_{\Pi_i(\omega)},Q)\leq 0.$$

Otherwise, there is information acquisition.

[11] Interestingly, with ambiguity averse preferences, the myopic best depends on the previous announcement, as shown in Galanis et al. [2023].

3.6 Information aggregation

We say that information is aggregated if the traders' predictions converge to the intrinsic value $X(\omega)$ of security X, for all $\omega \in \Omega$. For every $\omega \in \Omega$, let $y_k(\omega)$ be the announcement of the trader who moves in period t_k. The announcement $y_k(\omega)$ depends on ω because traders have different private information across states. Because $\{y_k\}_{k=1}^{\infty}$ is a sequence of random variables, we need a probabilistic version of convergence.

Definition 4. *Under a profile of strategies in Γ^M or Γ^S, information aggregates if sequence $\{y_k\}_{k=1}^{\infty}$ converges in probability to random variable X.*

3.7 Separability

Ostrovsky [2012] introduced the class of separable securities and showed that they are necessary and sufficient for aggregating information.

Definition 5. *A security X is called non-separable under partition structure Π if there exists probability μ and value $v \in \mathbb{R}$ such that:*

(i) $X(\omega) \neq v$ *for some* $\omega \in Supp(\mu)$,

(ii) $E_\mu[X|\Pi_i(\omega)] = v$ *for all* $i = 1, ..., n$ *and* $\omega \in Supp(\mu)$.

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.

4 Results

In this section, we present our main results. First, we characterize the class of securities that are separable for all information structures and show that it contains only two types. The first is the Arrow-Debreu (A-D) and the second pays a minimum and a maximum value in two states, and a middle value in all other states.[12]

[12] Our motivation for defining the class of κ separable security comes from the fact that very few securities are always separable.

Our second finding is that 'most' securities are κ separable for some κ. An example of such a security pays differently in each state. Because this condition is independent of the information structure, information aggregation is robust to adding or removing traders and to changing their private information.

Third, we show that when the cost of information acquisition is κ, in strategic and non-strategic environments we have that κ separable securities are sufficient for information aggregation. Moreover, if the security is κ non-separable, then there exists an equilibrium where information does not aggregate. This generalizes the result of Ostrovsky [2012] for separable securities.

Fourth, we compare the prediction accuracy of the market with that of a poll, where traders obtain information and announce simultaneously only once. We show that the prediction of the market is strictly better if and only if the security is κ separable.

Finally, we show in non-strategic settings that, for all securities except the A-D, information acquisition can make information aggregation both faster and slower. Combining the first and the last results, we have that the A-D is the only type of security that always aggregates information and for which the speed of aggregation is not impacted by the cost structure in a myopic environment.

4.1 Always Separable Securities

Separable securities can be very useful for a market designer because they always aggregate information. However, separability depends on the information structure Π, which the market maker might not know, so he may not be certain that the equilibrium price is a good predictor of the intrinsic value of the security. In this section, we characterize the securities that are separable for all information structures, so information aggregation is robust to changes in who is trading and what is their private information.

Unfortunately, this class is very small and uninformative. It consists only of the Arrow-Debreu (A-D) and another type of security.[13] An A-D security takes only two values, therefore it is uninformative because when information aggregates, it

[13] The class also includes the uninteresting case of securities that pay the same at all states.

can only predict whether a state has occurred or not. The other type of security takes only three values, hence it is also not very informative. We now characterize the securities that are separable for all information structures in $\mathcal{P}$.

Proposition 1. *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 $a < b < c$ such that $X(\omega_a) = a$ and $X(\omega_c) = c$ for two states ω_a, ω_c, and $X(\omega) = b$ for all $\omega \neq \omega_a, \omega_c$.*

Proof. 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 2 (Ostrovsky [2012])**.** *Security X is separable under partition structure Π if and only if, for every $v \in \mathbb{R}$, there exist functions $\lambda_i : \Pi_i \to \mathbb{R}$ 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 $\omega \neq \omega_a, \omega_c$.[14] Using Proposition 2, we need to show that for every $v \in \mathbb{R}$, there exist functions $\lambda_i : \Pi_i \to \mathbb{R}$ for $i = 1, \dots, n$ such that, for

[14]Our 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.

every state ω with $X(\omega) \neq v$,

$$(X(\omega) - v)\sum_{i \in I} \lambda_i(\Pi_i(\omega)) > 0. \tag{2}$$

If $v \leq a$, then condition (2) 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 (2) 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 (2) is satisfied for ω. Using a symmetric argument, we can show that (2) 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 2, security X is always separable.

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 agent 1 is $\{\{a, d\}, \{b, c\}\}$ for these four states, whereas for any other state, we have $\Pi_1(\omega) = \{\omega\}$. For agent 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*) $X(\omega) \neq v$ for some $\omega \in Supp(p)$,

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

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

$$\begin{cases} ap_1 + d(1-p_1) = bq_1 + c(1-q_1) \\ ap_1 + d(1-p_1) = ap_2 + c(1-p_2) \\ ap_1 + d(1-p_1) = bq_2 + d(1-q_2) \\ bq_1 + c(1-q_1) = ap_2 + c(1-p_2) \\ ap_1 + d(1-p_1) = bq_2 + d(1-q_2) \\ ap_2 + c(1-p_2) = bq_2 + d(1-q_2) \end{cases} \tag{3}$$

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

$$\begin{cases} xp_1 = yp_2 \\ (1-x)q_1 = (1-y)q_2 \\ (1-x)(1-q_1) = y(1-p_2) \\ x(1-p_1) = (1-y)(1-q_2) \end{cases} \tag{4}$$

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 (3 - 4) has the following solution:

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

These posteriors uniquely define the respective prior probabilities. □

Proposition 1 states that the only securities that are always separable are the A-

D and the security that pays the largest payoff in one state of the world, the lowest payoff in another, and in all other states of the world it pays the same payoff which is between the previous two payoffs. This suggests that, for any other security, if a market designer does not know the traders' information structure, then it is not guaranteed that information will aggregate. In the next section, we show how this problem is alleviated if we allow for information acquisition.

4.2 κ Separability

Our result that very few securities are always separable indicates that 'most' markets may not aggregate information. A natural question is whether the possibility of acquiring information fixes this problem. In this section, we define the class of κ separable securities, where $\kappa = (K, c)$ is the cost of information acquisition. In the next section, we show that they are necessary and sufficient for information aggregation.

We say that security X is κ non-separable given an information structure Π and a cost structure κ if there is a prior μ such that no one is acquiring any information, yet 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*

1. *Security X is non-separable at μ,*

2. *There is no information acquisition for X given μ.*

Otherwise, security X is κ separable.

Recall that a non-separable security is non-separable for at least one μ, whereas a separable security is separable for all μ. If security X is κ separable, then for each μ there are two cases. Either X is separable at μ, or X is non-separable at μ but there is information acquisition.

We make the following remarks. First, separability implies κ separability, for all κ. Second, non-separability implies κ non-separability for some κ.

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 $\kappa \in \mathcal{K}$.*

Proof. From Proposition 2, if security is separable then it is separable for any prior μ. Therefore, there does not exist μ such that security X is non-separable at μ. $\square$

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.*

Proof. Let the marginal cost $c \to \infty$, then it follows that information acquisition is guaranteed to be zero. Therefore, security X becomes κ non-separable. $\square$

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 $c' > c$ ($c' < c$).*

It would be natural to assume that for sufficiently low marginal cost c of signals, any security eventually becomes κ separable. But this is not true, as long it is not possible to buy unbounded experiments, which exclude some states with probability 1.

Proposition 3. *Under Assumption 1, there exists a security which is κ non-separable given some Π and for all κ.*

Proof. Consider security $X = (0, 1, 0, 1)$ and two traders with the following information structure. 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 such that no trader acquires any information and X is non-separable at μ. Consider 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 $2m$. 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 $2m$ is very close to 1, say $1 - \epsilon$. By repeating the consensus announcement of $2m = 1 - \epsilon$ and not acquiring any signals, both traders get 0 utility. Suppose that one trader wants to buy a signal structure Q, which, in the case of $X(\omega) = 1$, will move his posterior expected value of X from $1 - \epsilon$ to $1 - \epsilon + \nu$, for some $0 < \nu < \epsilon$. The upper bound of his utility, net of the cost of the signal and when $X(\omega) = 1$ realized, is $s(1 - \epsilon + \nu, 1) - s(1 - \epsilon, 1)$. This is positive because $1 - \epsilon + \nu$ is closer to 1 than $1 - \epsilon$. However, by decreasing ϵ appropriately it can be made as small as needed, and therefore smaller than $cK(\mu, Q)$ for all Q. This is a direct consequence of Assumption 1, which specifies that the cost gets arbitrarily high for posteriors which assign probability that is arbitrarily close to one for some state. For this result, we only need that the cost is non-decreasing. Hence, no trader will buy any signal structure and X is κ non-separable. □

We finally provide two independent and necessary conditions for a security to be κ non-separable for all $\kappa \in \mathcal{K}$.

Proposition 4. *Suppose X is κ non-separable given an information structure Π and for all $\kappa \in \mathcal{K}$. Then, the following two independent conditions are true.*

- *For some $v' \in X(\Omega)$, for all $\epsilon > 0$, there exists $v \in (v' - \epsilon, v' + \epsilon)$ for which there are no functions $\lambda_i : \Pi_i \to \mathbb{R}$ such that for all states ω with $X(\omega) \neq v$,*

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

- *There exist two states $\omega_a \neq \omega_b$ such that $X(\omega_a) = X(\omega_b)$.*

Proof. Suppose X is κ non-separable for all κ. Take κ with very small marginal cost c, so that for any prior μ for which there is no information acquisition, at each state $\omega \in Supp(\mu)$, and for each $i \in I$, $\mu_{\Pi_i(\omega)}$ assigns probability which is arbitrarily close to 1 to some state $\omega' \in \Pi_i(\omega)$. This means that $E_{\mu_{\Pi_i(\omega)}}[X] \in (v' - \epsilon, v' + \epsilon)$, for small enough $\epsilon \in \mathbb{R}$ and $X(\omega') = v'$.

Because X is κ non-separable, there exists a prior μ with no information acquisition and value $v \in (v' - \epsilon, v' + \epsilon)$ for which X is non-separable. This implies that, for any $\lambda_i : \Pi_i \to \mathbb{R}$,

$$E_\mu \left[(X(\omega) - v) \sum_{i \in I} \lambda_i(\Pi_i(\omega)) \right] =$$

$$\sum_{i \in I} E_\mu \left[(X(\omega) - v) \lambda_i(\Pi_i(\omega)) \right] =$$

$$\sum_{i \in I} \sum_{\omega \in Supp(\mu)} \mu(\Pi(\omega)) \lambda_i(\Pi_i(\omega)) E_{\mu_{\Pi_i(\omega)}} \left[(X(\omega) - v) \right] = 0,$$

where the last equality derives from the definition of non-separability. But then we cannot have $(X(\omega) - v) \sum_{i \in I} \lambda_i(\Pi_i(\omega)) > 0$ for all ω, hence there are no such λ_i functions.

For the second condition, we prove the contrapositive. Suppose that there do not exist two states ω_a and ω_b such that $X(\omega_a) = X(\omega_b) = v$. We will show that X is κ separable for some κ.

For any K and for sufficiently low marginal cost c, Assumption 1 implies that the only priors for which there is no information acquisition are the ones where the resulting posterior, for each $i \in I$ and partition element $\Pi_i(\omega)$, assigns probability close to 1 to some state $\omega' \in \Pi_i(\omega)$. Because X assigns different values to different states, if at least two traders' posteriors assign probability close to 1 to different states, their announcements will differ and therefore X is κ separable for that prior. If this is true for all priors for which there is no information acquisition, we are done.

Suppose that for some prior the traders assign the posterior probability close to 1 to the same state, at each state ω. Because of our assumption in Section 3.1 that $\bigcap_{i \in I} \Pi_i(\omega) = \{\omega\}$, if all agents' posteriors agree on the state on which they assign probability close to 1, this has to be ω. We can exclude the case that $\Pi_i(\omega) = \{\omega\}$ for all i and ω, because in that case, the security is trivially separable. We therefore have at least one trader i with $\omega' \in \Pi_i(\omega)$. The same reasoning establishes that given ω', everyone assigns probability close to 1 to ω'. But this is impossible for i, hence we have a contradiction.

To show that the conditions are independent, consider first the A-D security $(1,0,0)$, which satisfies the second condition. Let the partition of agent 1 be $\{\{\omega_1,\omega_2\},\{\omega_3\}\}$ and for agent 2 be $\{\{\omega_1,\omega_3\},\{\omega_2\}\}$. To see that it fails the first condition, set $v'=0$ and $\epsilon=0.5$. If $v\in[0,0.5)$, set $\lambda_1(\{\omega_1,\omega_2\})=\lambda_2(\{\omega_1,\omega_3\})=1$ and $\lambda_1(\{\omega_3\})=\lambda_2(\{\omega_2\})=-2$. If $v\in[-0.5,0)$, set $\lambda_i(\Pi_i(\omega))>0$ for all $i\in I$ and $\omega\in\Omega$. In both cases, (5) is satisfied for all states. A similar set of λ_i functions apply when $v'=1$, hence the first condition fails.

Consider now the security $(-2,-1,0,1,2)$, which violates the second condition. We will show that it satisfies the first. Suppose there are three agents. The partition of agent 1 is $\{\{\omega_2,\omega_4\},\{\omega_1,\omega_3,\omega_5\}\}$, for agent 2 it is $\{\{\omega_1,\omega_5\},\{\omega_2,\omega_3,\omega_4\}\}$ and for agent 3 it is $\{\{\omega_1,\omega_4\},\{\omega_2,\omega_3,\omega_5\}\}$. Let $v'=0$ and suppose that $v\in(0,\epsilon)$, for any $\epsilon>0$. We then have that $X(\omega)-v<0$ for $\omega_1,\omega_2,\omega_3$, and $X(\omega)-v>0$ for ω_4,ω_5.

Suppose there exist functions such that $\lambda_1(\{\omega_2,\omega_4\})=a$, $\lambda_1(\{\omega_1,\omega_3,\omega_5\})=b$, $\lambda_2(\{\omega_2,\omega_3,\omega_4\})=c$, $\lambda_2(\{\omega_1,\omega_5\})=d$, $\lambda_3(\{\omega_2,\omega_3,\omega_5\})=e$ and $\lambda_3(\{\omega_1,\omega_4\})=f$. For (5) to be satisfied for all satisfied, we have the following inequalities.

$$\begin{cases} b+d+f<0 \\ a+c+e<0 \\ b+c+e+<0 \\ a+c+f>0 \\ b+d+e>0 \end{cases} \tag{6}$$

If we add the first with the fifth, we have $e>f$. If we add the second with the fourth, we have $f>e$. This is a contradiction and such equations do not exist, which means that the first condition is satisfied.

□

A Corollary of this Proposition is that if security X provides different payments across different states, then, given any information structure Π, it is κ separable for some κ. In other words, if information is sufficiently cheap to acquire, there are easily describable securities where information aggregation is robust to changes to who participates in the market and what is their private information.

Corollary 2. *If $X(\omega) \neq X(\omega')$ for all $\omega, \omega' \in \Omega$, then, given any Π, X is κ separable for some κ.*

4.3 Strategic traders

Consider a game $\Gamma^S(\Omega, I, \Pi, X, h, y_0, \mu, \underline{y}, \overline{y}, s, \beta)$, where I is the set of n players, s is a strictly proper scoring rule, y_0 is the market maker's initial announcement at time t_0, μ is the common prior, h is the cost of statistical experiments, $[\underline{y}, \overline{y}]$ is the set of possible announcements, and β is the common discount rate. Let $(y_1, \ldots, y_k) \in [\underline{y}, \overline{y}]^k$ be a history of announcements and $(\tau_1, \ldots, y_k) \in T^k$ be a history of signals up to time t_k. Denote by T_i^k the collection of all histories about i's signals up to time t_k.

At time t_k, Player i with belief μ can choose to acquire information in the form of a statistical experiment $\mathcal{R} \in \mathcal{E}$, with cost $h(\mu, R)$. Each experiment $\mathcal{R}$ and belief μ uniquely induce a Bayesian plausible belief Q over posteriors, which costs $cK(\mu, Q) = h(\mu, R)$. We can therefore analyse the game in terms of an information cost structure $\kappa = (K, c)$, which is generated by h.

After receiving a signal τ of experiment $\mathcal{R}$ and updating to a posterior γ in the support of Q, player i makes an announcement y_k. Let $[\underline{y}, \overline{y}]^T$ be the set of all functions from signals to announcements in $[\underline{y}, \overline{y}]$. His mixed strategy at time t_k is a measurable function

$$\sigma_{i,k} : \Pi_i \times T_i^{k-1} \times [\underline{y}, \overline{y}]^{k-1} \times [0, 1] \longrightarrow \mathcal{E} \times [\underline{y}, \overline{y}]^T.$$

It specifies a statistical experiment and an announcement for each signal that is drawn from that experiment, given the element of his partition, the history of his past signals and everyone's announcements up to time t_k, 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 and the draws are independent of each other and of the true state ω.

The full state is $\phi = (\omega, \iota_1, \iota_2, \ldots)$ and describes the initial uncertainty and the randomizations of the players. Let $X(\phi) = X(\omega)$ be the true value of X at $\phi = (\omega, \iota_1, \iota_2, \ldots)$ and let $\Phi = \Omega \times [0,1]^{\mathbb{N}}$ 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 $\sigma = (\sigma_1, \ldots, \sigma_n)$ be a profile of strategies. Given a strategy σ and state ϕ, let $y_{i+nk}(\sigma, \phi)$ be the announcement of Trader i, $\mu_{i+nk}(\sigma, \phi)$ his belief, and $cK(\mu_{i+nk}(\sigma, \phi), Q_{i+nk}(\sigma, \phi))$ his cost of information acquisition in period t_{i+nk} which is generated by the experiment he has chosen and his current beliefs.

Definition 7. *A strategy profile σ is a Nash equilibrium if for every player i and every alternative strategy $\sigma' = (\sigma_{-i}, \sigma'_i)$, we have*

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

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

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

We now show that κ separable securities characterize information aggregation when the cost structure is $\kappa = (K, c)$.[15]

Theorem 1. *Fix information structure Π and cost of experiments h, which generates cost structure κ. Then:*

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

Proof. 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.

[15] Note that Assumption 1 is not needed for this result.

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 Q_ω be the optimal experiment that Trader i acquires at $\Pi_i(\omega)$, given his beliefs $r_{\Pi_i(\omega)}$. Trader i's instant opportunity is

$$\sum_{\omega\in\Omega} r(\omega)\left(\sum_{\gamma\in Supp(Q_\omega)} Q_\omega(\gamma)\sum_{\omega'\in\Omega}\gamma(\omega')\Big[s(E_\gamma[X],X(\omega'))-s(z,X(\omega'))\Big]-cK(r_{\Pi_i(\omega)},Q_\omega)\right).$$

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 1. *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 $\mathcal{R}\in\Delta$. 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 $\omega\in Supp(r)$ and $i\in I$. Because X is κ separable, there is information acquisition. This implies that for some $\omega\in Supp(r)$, 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 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 $E_r[X|\Pi_i(\omega)]=v$.[16] By setting $\chi_3=\min\{\chi_1,\chi_2\}$, we have that

[16]This is true because a proper scoring rule is 'order-sensitive' so that the further away the

at $\Pi_i(\omega)$, Trader i receives a strictly positive payoff of at least χ_3, for all previous announcements z. For any $\omega' \in Supp(r) \setminus \Pi_i(\omega)$, Trader i can repeat v and receive 0. Hence, the instant opportunity of i given r, which is the ex-ante expectation over all $\omega \in Supp(r)$, is at least $\chi = r(\Pi_i(\omega))\chi_3 > 0$ for all previous announcements.

□

Steps 2-4 are identical to the proof of Ostrovsky [2012]. The reason is that we only need to use the uniform lower bound that we have established in Step 1. Note that, once we fix an equilibrium strategy, the only difference from Ostrovsky [2012] is that players receive (up to) an extra signal in each period where they make an announcement.

The proof of the second part of the Theorem is similar to that of Ostrovsky [2012]. 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. In the corresponding game Γ^S where the initial announcement is v, no trader will find it profitable to acquire any information and their best response is to repeat v, in every period and after every history, hence it is a Bayesian Nash equilibrium.

□

4.4 The value of the market

If the cost of information drops significantly, do we even need markets to aggregate information through prices? Each agent could buy the necessary signals and then trade. In this section, we argue that this intuition is not correct, because markets become even more important in an environment with information acquisition. The reason is that they are able to aggregate information fully before it becomes economically viable for each trader to acquire the required information on their own.

To make this point, we compare the prediction accuracy of the market with that of a poll, where traders simultaneously make only one announcement and we then

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]).

take the average.[17] 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.[18]

We first show that κ separability is equivalent to requiring that the market is more accurate than the poll, for all non-degenerate priors given X.

Definition 8. *For each $\omega \in \Omega$ and common prior μ, the prediction of the poll is the average of the expected myopic predictions, where traders optimally obtain some random posterior Q and then announce simultaneously in the first period of the corresponding game Γ^M:*

$$p^p(\omega, \mu) = \frac{\sum_{k=1}^{n} y_n}{n}.$$

The prediction $p^m(\omega, \mu)$ of the market is the last price of Γ^S.

We define the accuracy of the market given ω and μ as $A^m(\omega, \mu) = 1 - |p^m(\omega, \mu) - X(\omega)|$, and similarly for the poll, $A^p(\omega, \mu) = 1 - |p^p(\omega, \mu) - X(\omega)|$. The highest possible accuracy is 1, when $p^m(\omega, \mu) = X(\omega)$. The expected accuracy of the market given μ is defined as $A^m(\mu) = E_\mu A^m(\omega, \mu)$, whereas for the poll it is $A^p(\mu) = E_\mu A^p(\omega, \mu)$.

Recall that a non-degenerate prior μ given X does not assign probability 1 to a unique value of security X. Then, if the security is κ separable for some κ, then the information would be fully aggregated for some positive marginal cost of information, which is not true for polls. We, therefore, have the following remark.

[17] Note that there are many ways of improving the accuracy of a poll by aggregating announcements differently [Baron et al., 2014].

[18] In our framework with myopic traders, 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.

Remark 4. *Under Assumption 1, fix an information structure* $\Pi \in \mathcal{P}$ *and a cost of information acquisition* $\kappa = (K, c)$*. If the security is* κ *separable, then for any non-degenerate* μ *given* X *and for all* $0 < c' \leq c$*, information gets fully aggregated by a market with cost* $\kappa' = (K, c')$*, so that* $A^m(\omega, \mu) = 1$ *for all* $\omega \in Supp(\mu)$*, but it is never fully aggregated by the poll, so that* $A^p(\omega, \mu) < 1$*.*

Proof. The first part follows directly from Theorem 1. Assumption 1 implies that individual traders would never acquire full information in the first round. Because μ is non-degenerate given X, we have that $p^p \neq X(\omega)$. □

We now show that κ separability is equivalent to the market being strictly more accurate than the poll, for all non-degenerate priors given X.

Proposition 5. *Suppose Assumption 1 and cost of information* κ*. Security* X *is* κ *separable given* $\Pi \in \mathcal{P}$ *if and only if* $A^m(\mu) > A^p(\mu)$ *for all non-degenerate priors* μ *given* X*.*

Proof. If X is κ separable, then from Theorem 1 we have that information aggregates at all states and therefore $A^m(\mu) = 1$. Assumption 1 implies that no trader would acquire full information, hence $p^p(\omega, \mu) \neq X(\omega)$ and $A^p(\mu) < 1$.

If 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. This means that everyone agrees on the announcement at all states 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 and $A^m(\mu) = A^p(\mu)$. □

Define the value of the market (with security X) to be $V(X) = \min_{\mu \in \Delta_0} A^m(\mu) - A^p(\mu)$, the minimum improvement in accuracy given all non-degenerate priors given X. We then have that X is κ separable if and only if $V(X) > 0$.

In the following example, we show how the market and the poll compare in prediction accuracy if we fix a prior μ for which the security is non-separable and we vary the marginal cost of information c. Suppose that the security is $X = (0, 1, 2, 3)$ and the common prior is $\mu = (\frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8})$ on state space $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$. 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 because every agent for any state realization announces $v = \frac{3}{2}$, yet there is uncertainty about the value of the security. Traders can acquire information, where K is the Shannon cost of information. If traders acquire information they either agree on the value of the security being $v \neq X(\omega)$ or information gets aggregated.

Figure 4.4 shows the expected accuracy of markets and polls for a given prior. When the marginal cost of information is high (c>4.5), then the market and the poll accuracy is the same and is the lowest. Then, as information becomes cheaper, Trader 2 would acquire information if the state is either 1 or 4. Therefore, his announcement would differ across partitions and, hence, reveal his private information. Then, Trader 1 would know the state and announce $v = X(\omega)$. Thus, a small change in the marginal cost of information allows the market to fully aggregate information. In contrast, the poll's accuracy improves gradually as information gets cheaper. This means that, as the cost of information decreases, the value of the market first skyrockets compared to polls, and then the difference in accuracy gradually disappears.

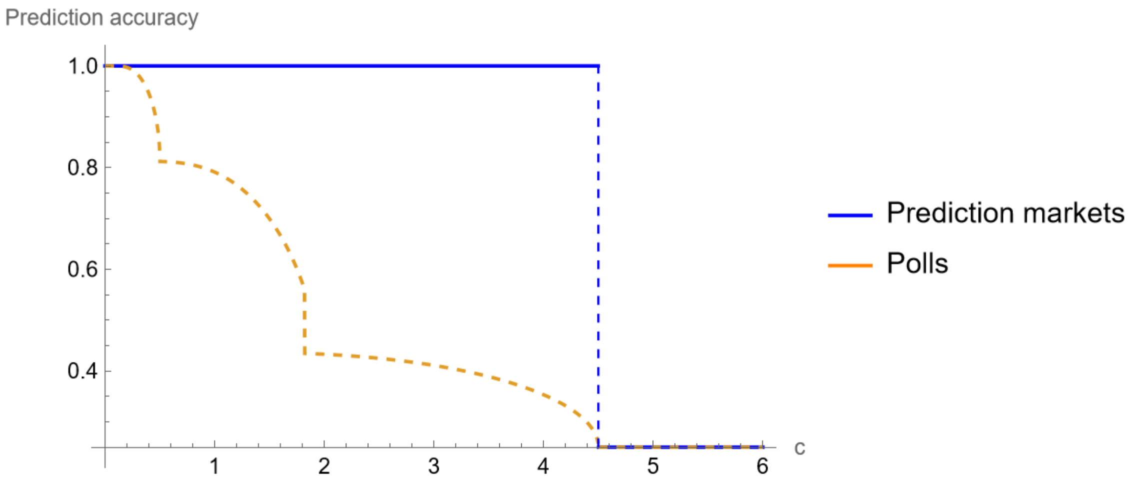

Figure 1: Prediction accuracy ($A(\mu) = E_\mu A(\omega, \mu)$) for markets and polls.

We conclude by observing that markets can incentivize traders to acquire and utilize information more efficiently when the cost is close to the threshold of 4.5, because a small information acquisition could enable the aggregation of all available information, whereas in polls it can only marginally improve individual predictions.

4.5 The speed of information aggregation

As we showed in the example in Section 2.2, information acquisition can make the process of information aggregation both slower and faster, depending on the parameters. In this section, we show that this is true for all securities except for the Arrow-Debreu (A-D) security. In particular, with A-D securities the information aggregation can never be faster (or slower) as compared to the no information case. If the process cannot get faster or slower, the competition of traders to acquire more information in order to get more of the surplus produces neither negative nor positive externalities in society.

Let $t_I^*(\omega)$ be the period where information aggregation is achieved, given a separable security and state ω, in an environment where traders can get information before they make an announcement. Similarly, let $t_{NI}^*(\omega)$ be the period where information aggregation is achieved, in an environment where there is no information acquisition.

Proposition 6. *Suppose non-constant security X is not A-D. Then, there is information structure $\Pi \in \mathcal{P}$ for which X is separable, such that for some information cost $\kappa = (K, c)$ and $\omega \in \Omega$,*

(i) $t_{NI}^*(\omega) < t_I^*(\omega)$ *for some common prior p,*

(ii) $t_{NI}^*(\omega) > t_I^*(\omega)$ *for some common prior p'.*

Proof. Consider non-constant security X which is not A-D. Then, X maps to at least three values, $a < b < c$, and we denote the respective states as a, b, and c.

Let the partition of agent 1 be $\Pi_1 = \{\{a, c\}, \{b\}\}$ for these three states, whereas for any other state, we have $\Pi_1(\omega) = \{\omega\}$. For agent 2 it is $\Pi_2 = \{\{a, b\}, \{c\}\}$ and for any other state we have $\Pi_2(\omega) = \{\omega\}$. Hence, the information structure is in $\mathcal{P}$. For agent 1, let p_1 be the probability of state a conditional on $\{a, c\}$. For agent 2, let p_2 be the probability of state a conditional on $\{a, c\}$.

We first show that X is separable with respect to Π, so that t_I^*, t_{NI}^* are well-defined for all priors. If Ω has more than three states, then for all $\omega' \in \Omega \setminus \{a, b, c\}$ and $i \in I$ we have $\Pi_i(\omega') = \{\omega'\}$ and $E_p[X|\Pi_i(\omega')] = X(\omega')$. Therefore, separability only depends on the three states, a, b, and c. From Proposition 1, the security X is separable with respect to the restriction of Π on $\{a, b, c\}$. But then, X is separable with respect to Π.

Second, we show that (i) is true at state a, given Π and some K. Because $X(a) < X(b) < X(c)$, there exists unique p^* with full support on the first three states and Bayesian updates p_1^* such that $E_{p^*}[X|\Pi_1(\omega)] = v = b$ for all $\omega \in Supp(p^*)$. Since K is strictly convex and continuous, all posteriors in the interior of the support of p are feasible. Then, we can find c and $p_1 \neq p_1^*$ such that p_1^* is the optimal posterior (Lemma 2 in Caplin et al. [2022]).[19] This implies that the first agent optimally acquires information given c and updates his beliefs from p_1 to p_1^* with non-zero probability. Therefore, the second agent would not get certain information about the first's agent partition, and the game proceeds to the next period, so $t_I^*(a) > 2$. However, without information acquisition, p_1 is such that $E_p[X|\Pi_i(\omega)] \neq v$ for all $i = 1, 2$ and $\omega \in Supp(p)$. Therefore, without information acquisition, after the first agent announces his prediction, the second agent will know the state and $t_{NI}^*(a) \leq 2$.

Finally, we show that (ii) is true at state a given the same Π and some K. Following the same argument as before, because $X(a) < X(b) < X(c)$, there exists p with full support on the first three states and resulting unique p_1 such that $E_p[X|\Pi_1(\omega)] = v = b$ for all $\omega \in Supp(p)$. Therefore, after the first agent announces his prediction, the second agent does not know the state and $t_{NI}^*(a) > 2$. With information acquisition, we can find c such that the first agent acquires information, updates his beliefs to $p_1^* \neq p_1$, and then announces $E_{p^*}^*[X|\Pi_i(\omega)] \neq v$ for all $i = 1, 2$ and $\omega \in Supp(p)$ with non zero probability. To see that, consider infinitely small c and note that the state is a. Then, the agent announces $X(a) + \epsilon < X(b)$, where $\epsilon > 0$. Therefore, the second agent will know the state and $t_I^*(a) = 2$. □

Finally, we show that the speed of aggregation is unchanged if the security is A-D.

Proposition 7. *If X is A-D, then $t_{NI}^*(\omega) = t_I^*(\omega)$ for any state $\omega \in \Omega$, cost κ,*

[19] To see that $p_1 \neq p_1^*$, recall that p^* is unique, hence, there does not exist any $p_1 \neq p_1^*$ which would delay information acquisition. Then, the only possibility for p^* and p to coincide is if there is no information acquisition. But we can always find such c that the agent always acquires information for some p and therefore we can always find $p_1 \neq p_1^*$.

information structure $\Pi \in \mathcal{P}$*, and common prior* p*.*

Proof. Without loss of generality we assume that $X(\omega^*) = 1$ and $X(\omega) = 0$ for all other ω. In every period t, there are two cases. First, the agent announces 0 and the process ends. In that case, the agent does not acquire any information, because he knows that the price is 0. Second, the agent considers ω^* with $X(\omega^*) = 1$ to be possible, acquires information, and forms a posterior. As the posterior cannot reveal any state with certainty, his private information $\Pi_i(\omega) \cap E$, where E is the public event revealed by previous announcements, stays the same, but his posterior might change. Hence, no information is revealed to other agents, because he does not make a different announcement based on which partition cell he is in. Because no information is revealed to other agents conditional on the process continuing, the common knowledge event that is created by each announcement is the same with and without information acquisition, hence, the process ends in the same number of periods. Note, however, that announcements may differ across the two environments. $\square$

The property that speed is unaffected is related to the fact that the A-D security is not very informative. An agent will buy a posterior only in the partition cell where 1 is possible, as in all others he knows that the value is 0. This action does not have the positive externality of revealing to other agents what is his partition cell and therefore it does not provide any public information about what the true state is. Note, however, that some positive externality persists. The other agents can solve the announcer's problem and therefore know the optimal signal structure he has purchased. By hearing the announcement, they also update their own posteriors and they can benefit as long as it is their turn to announce and there is still some surplus to be obtained. For example, if the initial trader buys a signal structure and moves the price from 0.5 to 0.99 (when the correct price is 1), then all other agents can only benefit by moving it from 0.99 to 1, hence the remaining surplus is very small. Finally, even though the number of periods for full aggregation remains the same, the price will get faster to the true value, hence the market will still benefit by attaching faster a higher probability to the true value of the security.

References

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.

David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102. University of California Press, 1951.

Alexander W Bloedel and Weijie Zhong. The cost of optimally-acquired information. *Unpublished Manuscript*, 2020.

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.

Andrew Caplin and Mark Dean. Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203, 2015.

Andrew Caplin and Daniel Martin. A testable theory of imperfect perception. *The Economic Journal*, 125(582):184–202, 2015.

Andrew Caplin, Mark Dean, and John Leahy. Rational inattention, optimal consideration sets, and stochastic choice. *The Review of Economic Studies*, 86(3): 1061–1094, 2019.

Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130 (6):1676–1715, 2022.

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.

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.

Tommaso Denti. Posterior separable cost of information. *American Economic Review*, 112(10):3215–59, 2022.

Tommaso Denti, Massimo Marinacci, and Aldo Rustichini. Experimental cost of information. *American Economic Review*, 112(9):3106–3123, 2022.

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.

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, forthcoming*, 2023.

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.

Benjamin Hébert and Michael Woodford. Neighborhood-based information costs. *American Economic Review*, 111(10):3225–3255, 2021.

Benjamin Hébert and Michael Woodford. Rational inattention when decisions take time. *Journal of Economic Theory*, 208:105612, 2023.

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.

Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6): 1315–35, 1985.

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.

Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273, 2023.

Filip Matějka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1): 272–298, 2015.

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*, page rdv017, 2015.

Michael Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–647, 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 information acquisition and aggregation through prices. 2022.

Matthias Pelster and Joel Val. Can chatgpt assist in picking stocks? *Finance Research Letters*, page 104786, 2023.

Jan Christoph Schlegel, Mateusz Kwaśnicki, and Akaki Mamageishvili. Axioms for constant function market makers. *Available at SSRN*, 2022.

Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, April 2003.

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.

Constantino Tsallis. Possible generalization of boltzmann-gibbs statistics. *Journal of statistical physics*, 52:479–487, 1988.

Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–26, 2004.