

Information Aggregation Under Ambiguity: Theory and Experimental Evidence

Spyros Galanis

Department of Economics, Durham University, UK

Christos A. Ioannou

Centre d'Économie de la Sorbonne, Université Paris 1 Panthéon - Sorbonne, France

and

Stelios Kotronis

Department of Economics, Durham University, UK

First version received June 2019; Editorial decision August 2023; Accepted January 2024 (Eds.)

We study information aggregation in a dynamic trading model. We show theoretically that separable securities, introduced by Ostrovsky in the context of Expected Utility, no longer aggregate information if some traders have imprecise beliefs and are ambiguity averse. Moreover, these securities are prone to manipulation as the degree of information aggregation can be influenced by the initial price set by the uninformed market maker. These observations are also confirmed in our laboratory experiment using prediction markets. We define a new class of strongly separable securities, which are robust to the above considerations and show that they characterize information aggregation in both strategic and non-strategic environments. We derive several testable predictions, which we are able to confirm in the laboratory. Finally, we show theoretically that strongly separable securities are both sufficient and necessary for information aggregation but, strikingly, there does not exist a security that is strongly separable for all information structures.

Key words: Information aggregation, Ambiguity aversion, Financial markets, Prediction markets, Experiments

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

1. INTRODUCTION

Making predictions about future events is an inescapable part of decision-making. Revenues in the forecasting industry are estimated at around \$300 billion in current dollars (Atanasov *et al.*, 2017), hence even slightly better predictions are economically beneficial for individuals, governments, firms, and organizations. Prediction markets constitute one of the most promising

tools to perform forecasts as they leverage the wisdom of the crowds by aggregating information that is dispersed among individuals.

In fact, in several cases, prediction markets perform significantly better than other conventional forecasting methods, such as polls or expert opinion. [Berg *et al.* \(2008\)](#) compared the predictions in the five presidential elections between 1988 and 2004 of the Iowa Electronic Markets and those of 964 polls. They found that 74% of the time, the prediction market was closer to the truth, whereas for forecasts 100 days before the actual election, the prediction market outperformed the polls at every election. [Cowgill and Zitzewitz \(2015\)](#) examined data from prediction markets ran by Google, Ford, and an anonymous basic materials' conglomerate, and found that the internal prediction markets conducted improved upon the forecasts of experts in all three firms by as much as a 25% reduction in the mean squared errors.

Interestingly, in the case of a “once-in-a-lifetime” event, prediction markets *may* fare significantly worse. For instance, Cultivate Labs designed a prediction market on the outcome of the Brexit referendum. It ran for 10 days prior to the polling day and the closing prediction was a 15% probability of “leave,” suggesting that the most likely outcome would be “remain.”¹ On the contrary, an average of all polls, reported by the Financial Times on the day of the referendum, found 48% in favour of “remain” and 46% in favour of “leave,” suggesting a probability of Brexit closer to 50%.² Given that the actual result was 48.1% in favour of “remain” and 51.9% in favour of Brexit, the almost even split between the two outcomes reported by the Financial Times seems more accurate than the heavy favourite outcome of “remain” suggested by the prediction market.

Clearly, there are limitations to the forecasting ability of prediction markets. In this study, we examine the conditions under which prediction (and, more generally, financial) markets are successful at aggregating information. In particular, can they aggregate information for events that are rare or uncommon and for which beliefs are imprecise? The literature has so far focused exclusively on traders with precise probabilities about events and objective Expected Utility (EU) preferences. Specifically, [Ostrovsky \(2012\)](#) has shown that with unique priors and EU preferences, when payoffs are determined using the Market Scoring Rule (MSR) ([Hanson, 2003, 2007](#)), even if there are a few large and strategic traders, information aggregates for a large class of securities, called *separable*, which includes the Arrow–Debreu securities. More importantly, there are securities that are separable for all information structures; thus, a market designer can be sure that the prediction market will always aggregate information.

These results rely heavily on the assumption that traders share a unique (and common) prior. However, Brexit is a once-in-a-lifetime event for which no historical data exist. How can we be sure that the traders have precise probabilities for such a hard-to-quantify and unfamiliar event?³ If we cannot maintain the hypothesis of a unique prior and EU, it is no longer the case that markets aggregate information even if the traders' multiple priors are common. More importantly, a slight departure from a unique prior could result in the traders agreeing on a security price that is far from its intrinsic value.

1. The market can be found at [Cultivate Forecasts \(2016\)](#).

2. A closer look at the individual polls suggests a similar story. Throughout 2016, neither of the two outcomes was a consistent winner and margins were always small. In the telephone polls, No Brexit was a consistent winner but with a margin that was declining over time. The results can be found at [The Financial Times \(2016a, 2016b\)](#).

3. Additional evidence is provided by [Kilka and Weber \(2001\)](#) who study experimentally the investment decisions of German subjects on stocks of (self-reported) familiar German banks and less familiar Japanese ones. The authors report that German subjects have more ambiguous beliefs about the associated outcomes of the Japanese banks. Earlier, [Heath and Tversky \(1991\)](#) found similar evidence.

TABLE 1
Information Structure

Outcome	Trader 1's signal	Trader 2's signal
Brexit	Referendum Not Cancelled	Either Brexit or Cancelled
No Brexit	Referendum Not Cancelled	No Brexit
Referendum Cancelled	Referendum Cancelled	Either Brexit or Cancelled

Notes: This table depicts the private signals of Trader 1 and Trader 2. The two traders' pooled information always reveals the true state.

To show this, consider the ambiguity aversion model of [Gilboa and Schmeidler \(1989\)](#), where a decision maker acts as if maximizing the minimum expected utility over a set of multiple priors (henceforth, referred to as MEU preferences).⁴ An important insight, which we prove in Lemma 1 and use heavily, is that with multiple priors, the optimal announcement of a myopic trader is still unique and the expectation of the security according to *one* of her beliefs. The choice of the belief, however, depends on the previous announcement, thus introducing path dependence (which is absent if the prior is unique). If the previous announcement happens to be the expectation of the security according to some of *i*'s beliefs, then *i*'s optimal myopic strategy is to repeat it. As we show in the example of Section 2, path dependence implies that the security is susceptible to manipulation, for instance, by the market maker who sets the initial price and can thus influence the degree of information aggregation.

To build some intuition, consider two individuals who trade an Arrow–Debreu security *X* in a dynamic setting.⁵ Suppose there are three possible states: (i) Brexit, (ii) No Brexit, and (iii) Referendum Cancelled. Security *X* pays 1 if Brexit occurs (*i.e.* the intrinsic value in that state is 1) and 0 otherwise. The information structure is such that Trader 1 cannot distinguish between Brexit and No Brexit but knows if the referendum is cancelled. Trader 2 cannot distinguish between Brexit and the referendum being cancelled, but knows if the referendum result is No Brexit.⁶ To simplify the exposition, we assume that the two traders are non-strategic and take turns (*i.e.* alternate) in announcing the price to maximize their period payoff according to the MSR. A scoring rule, like the quadratic, 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. Each announcement reveals some information about the intrinsic value of the security, which may prompt the other trader to revise her announcement. We say that information aggregates if the announcements converge to the intrinsic value of the security.

Suppose that the true state is Brexit and the market maker's initial announcement is 0. In the EU framework, the initial announcement plays no role. When Trader 1 learns that the referendum is not cancelled, so that either Brexit or No Brexit is true, she updates her unique belief and announces the expected value of *X*, which is a number strictly between 0 and 1. This reveals to Trader 2 that the referendum is not cancelled, otherwise Trader 1 would know it and announce a value of 0. Trader 2 already knows that no Brexit is not true, hence announces 1. This informs Trader 1 that Brexit is true, and she also announces 1 leading to information aggregation.

4. [Anantanasuwong et al. \(2019\)](#) use an incentivized survey on a representative sample of investors to study ambiguity attitudes across different assets. They find that around 65% of investors are ambiguity averse. Moreover, ambiguity aversion is highly and positively correlated across these assets.

5. The formal treatment of this example for the case of a quadratic scoring rule is presented in Section 2. However, the arguments work for any proper scoring rule.

6. Table 1 in Section 2 summarizes the information structure. Note that the traders' combined information always reveals the true state.

In the MEU framework, the initial announcement is crucial and may prevent information aggregation due to path dependence. Suppose that at least one prior (but not all) assigns zero probability to Brexit.⁷ When Trader 1 learns that the referendum is not cancelled, she knows that either Brexit or No Brexit is true and updates each of her priors. She then announces the expected value of X according to one of her updated beliefs. Given that the previous announcement was 0 and the expected value of X according to one of her beliefs is 0, she makes the exact same announcement.⁸ If there was no Brexit, her information would be the same and so would make the same announcement. If the referendum was cancelled, Trader 1 would know this (through her private signal) and again announce 0. Given that the same announcement of 0 would be made in all possible states, no public information is revealed from Trader 1's announcement. As a result, Trader 2 does not learn anything from Trader 1's announcement and her announcement is, for similar reasons, 0. In turn, Trader 1 also announces 0.

The market fails to aggregate information because both traders do not want to deviate from an announcement of 0. However, if the initial announcement was different, there would be information aggregation. In particular, in this example, any non-zero initial announcement would prompt Trader 1 to announce something other than 0, which would then reveal to Trader 2 that the referendum is not cancelled, hence Trader 2 would learn that there is Brexit and information would thus aggregate.⁹ In summary, the example shows that information which would be revealed under EU preferences fails to be revealed under MEU preferences.

We make three key contributions in this article. First, to the best of our knowledge, we are the first to analyse dynamic prediction markets with ambiguity aversion. We propose a new class of strongly separable securities and show that in a prediction market which implements the MSR, they are necessary and sufficient for information aggregation. Theorem 1 characterizes information aggregation in terms of strongly separable securities for the case of myopic players. For the case of strategic players, the trading procedure is an infinite-horizon game with incomplete information. Given that traders are ambiguity averse, they might be dynamically inconsistent. This means that Trader i might devise an optimal continuation strategy at time t , which may not be optimal for her at a later time. To tackle with this problem, we generalize the Revision-Proof equilibrium, first studied by [Asheim \(1997\)](#) and [Ales and Sleet \(2014\)](#) in the context of infinite-horizon complete information games with time-inconsistent preferences, to games with incomplete information. Theorem 2 shows that strongly separable securities are both necessary and sufficient for information aggregation for all Revision-Proof equilibria. Although we prove these results for the MEU preferences model of [Gilboa and Schmeidler \(1989\)](#), in the [Supplementary Appendix](#), we show that they are also true for the much larger class of Variational preferences ([Maccheroni et al., 2006a, 2006b](#)), which includes, for example, the Multiplier preferences of [Hansen and Sargent \(2001\)](#), and the class of Smooth Ambiguity preferences ([Klibanoff et al., 2005](#)). Interestingly, the set of strongly separable securities stays the same as we move from MEU to the general class of Uncertainty Averse preferences ([Cerreià-Vioglio et al., 2011](#)).

Our second contribution is an impossibility result. In Proposition 3, we show that no security is strongly separable for *all* information structures and this result extends to all Uncertainty

7. This simplifies the exposition but it is not necessary. In Appendix C, we show how path dependence and no information aggregation can arise when all priors have full support.

8. This is due to Lemma 1 as we explain above. Intuitively, due to her MEU preferences, she wants to minimize the expected difference between her score and the score of the previous announcement. This difference is minimized when the announcements coincide.

9. In Appendix C, we show that one can easily construct examples where information aggregation fails for multiple initial announcements.

Averse preferences.¹⁰ This property is not true for separable securities. For example, Arrow–Debreu securities are always separable. Given that strongly separable securities characterize information aggregation, we show that if we move away from EU and precise beliefs, there is no prediction market that can aggregate information for all possible information structures. Furthermore, if we cannot find a security that can always aggregate information in the special case of prediction markets, we cannot hope to find one in the more general class of financial markets. In other words, imprecise beliefs can severely constrain the ability of markets to generically aggregate information, which goes against the general consensus of the literature that starts with Hayek (1945).¹¹ This is detrimental to both investors and policy makers who can no longer trust that market prices “incorporate all available information.” Mispricing can have distortionary effects on investment, stemming from under- or over-investment. Moreover, path dependence and manipulation can lead to persistent price bubbles.

Our third contribution is to investigate and confirm our testable predictions in an incentivized laboratory experiment that we conducted, where subjects assumed the role of traders in prediction markets forecasting the value of a security in sequential trading. Specifically, we examined the impact on information aggregation of three dimensions: the market type (unique priors and EU preferences versus multiple priors and MEU preferences), the security type (separable versus strongly separable), and the initial price announcement of the market maker. Our first set of results finds that in the case of separable securities, information aggregation is significantly worse in environments with imprecise beliefs and ambiguity-averse individuals compared to that in environments with precise beliefs and EU preferences. This is not the case in the mirrored environments with strongly separable securities; specifically, information aggregation across the two environments is not significantly different. The latter result is in line with our Theorems 1 and 2. Our second set of results, finds that, in the case of separable securities, the initial price announcement of the market maker in an environment with imprecise beliefs and ambiguity-averse individuals can influence subjects’ behaviour and, thereby, the degree of information aggregation. On the contrary, in the case of strongly separable securities, the initial announcement does not influence subjects’ behaviour in the same environment, which is again consistent with our theory. Taken together, these results suggest that strongly separable securities aggregate information and are resilient to manipulation by the market maker in environments with imprecise beliefs and ambiguity aversion.

Our article contributes to two main strands of the literature. The first strand looks at ambiguity and information aggregation (revelation) in various contexts. The underlying themes here revolve around information transmission, interpreting information, and information acquisition. Condie and Ganguli (2011) demonstrate a failure of information transmission with ambiguity averse agents in standard heterogeneous information exchange economies. In the context of common values voting games with ambiguity averse voters, Ellis (2016) finds that there is no equilibrium in which information aggregates. Chen (2022) allows informational ambiguity to occur naturally in a sequential learning problem to find that it can result in an information cascade. Mailath and Samuelson (2020) study agents who have different and incomplete models

10. See Section 1 of the Supplementary Appendix.

11. Grossman (1976) showed that, in equilibrium, prices aggregate information. Radner (1979) introduced the concept of Rational Expectations Equilibrium (REE) and proved that generically prices aggregate information dispersed among traders. Several results regarding the convergence of REE in dynamic settings have been shown by Hellwig (1982), McKelvey and Page (1986), Dubey *et al.* (1987), Wolinsky (1990), and Golosov *et al.* (2014) among others. Siga and Mihm (2021) provide microfoundations for REE using common-value auctions and study when prices aggregate information.

to elucidate the sense in which interpretations can effectively aggregate information and generate approximate consensus. Finally, [Mele and Sangiorgi \(2015\)](#) analyse costly information acquisition in asset markets with ambiguity averse traders to show that when uncertainty is high enough, information acquisition decisions become strategic complements and lead to multiple equilibria.¹²

The second strand looks at the increasingly extensive literature on prediction markets.¹³ The first theme in this strand studies the degree and conditions of information aggregation of prediction markets in various frameworks. [Ostrovsky \(2012\)](#) and [Chen *et al.* \(2012\)](#) show that in a market with dynamically consistent traders, separable securities, introduced by [DeMarzo and Skiadas \(1998, 1999\)](#), are both necessary and sufficient for information aggregation. [Dimitrov and Sami \(2008\)](#) and [Chen *et al.* \(2010\)](#) also look at information aggregation but focus instead on varying the assumptions regarding the traders' information structure. The second theme asks whether prediction markets can be manipulated. In the theoretical literature, [Ottaviani and Sørensen \(2007\)](#) provide the first formal analysis of outcome manipulation in a corporate prediction market setup, where traders are able to influence the outcome. In the empirical literature, most studies find very little evidence of price manipulation, both in the actual markets (see [Camerer, 1998](#); [Wolfers and Leigh, 2002](#); [Rhode and Strumpf, 2004](#)), and in the laboratory ([Hanson *et al.*, 2006](#); [Hanson and Oprea, 2009](#)). However, [Zitzewitz \(2007\)](#) and [Snowberg *et al.* \(2013\)](#) document a case from actual markets where a manipulator was able to influence the price on an Arrow–Debreu contract. Along similar lines, [Veiga and Vorsatz \(2010\)](#) show experimentally that, under some conditions, prices can be manipulated by an uninformed trader, which is also corroborated in [Jian and Sami \(2012\)](#).

Although some aspects and ideas in the aforementioned studies do find common ground in our study, what we propose here is different. First, we depart from (and thus contribute to) the existing literature, by analysing prediction markets with ambiguity averse and dynamically inconsistent traders, not only for the MEU framework but also for the more general ones of Variational and Smooth Ambiguity preferences.¹⁴ We do so because the alternative EU framework with traders that have precise beliefs is unrealistic and highly stylized for events that are rare or unfamiliar. Our analysis culminates in a profound result for asset markets in general: there is no way to build a securitization scheme that will ensure information revelation for all information structures. Second, our approach to investigate price manipulation is different from the existing studies. We thus extend this stream of research by utilizing a new channel, where imprecise beliefs interact with the initial price announcement of the uninformed market maker. Consequently, we are able to vary directly the initial price to determine the effect on the degree of information aggregation of the security. Third, our approach is holistic in the sense that it combines testable predictions with an empirical investigation by means of a controlled laboratory experiment.

The article adheres to the following plan. In Section 2, we provide the formal treatment of our introductory example. Section 3 describes the model. In Section 4, we characterize information aggregation for the case of myopic traders, whereas in Section 5, we examine the case of strategic traders. In Section 6, we describe our experiment and discuss the support for our theory. Finally, in Section 7, we conclude and offer suggestions for future research. All proofs are included

12. Relatedly, [Page and Siemroth \(2017\)](#) study experimentally information acquisition in prediction markets at the individual level to find that traders with larger endowments, existing inconclusive information, lower risk aversion, and less experience in financial markets tend to acquire more information.

13. See [Wolfers and Zitzewitz \(2004\)](#) for an early overview of the literature.

14. [Galanis and Kotronis \(2021\)](#) also study prediction markets with dynamically inconsistent traders. However, the cause is not ambiguity aversion but being boundedly rational and unaware of some signals ([Galanis, 2011, 2013](#)).

in the Appendices. In the [Supplementary Appendix](#), we extend our results to the framework of Uncertainty Averse preferences, show the existence of a Revision-Proof equilibrium, and include the experimental instructions.

2. AN EXAMPLE

In this section, we describe in detail the 2016 Brexit referendum example of Section 1. The dynamic trading mechanism begins with an initial public announcement about the value of the security by the market maker and with nature choosing a state. Then, each trader sequentially announces in public her prediction, which may reveal some of her private information. A score for each prediction, based on a strictly proper scoring rule, is calculated after trading ends and the true state is revealed. For non-strategic traders, who only care about their current payoff, MSR ensures that the optimal strategy is to announce the expected value of the security given their posterior beliefs. The per-period utility of a trader is calculated by subtracting, from the score of her prediction, the score of the prediction made by the previous trader. A potential interpretation could be that each time a trader makes a prediction, she “buys out” the previous one.

The state space has three states, $\Omega = \{\omega_1, \omega_2, \omega_3\}$, which correspond to Brexit, No Brexit, and Referendum Cancelled, respectively. Trader 1's information partition is $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}\}$, whereas Trader 2's is $\Pi_2 = \{\{\omega_1, \omega_3\}, \{\omega_2\}\}$. They trade an Arrow–Debreu security X that pays 1 at ω_1 (*i.e.* if Brexit occurs) and 0 otherwise. The information structure is depicted in Table 1. In particular, Trader 1 is informed whether the referendum is cancelled or not. Trader 2 is informed whether the referendum's result is against Brexit or not. Notice that the two traders' pooled information always reveals the true state.

The two traders are non-strategic, they have MEU preferences and share a common set of priors \mathcal{P} , which is the convex hull of $p^1 = (0, \frac{1}{2}, \frac{1}{2})$ and $p^2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. If trader i 's announcement is y , the intrinsic value of the security is $x^* = X(\omega)$, and the announcement of the previous trader (or the market maker) is z , then i 's utility is $s(y, x^*) - s(z, x^*)$, where $s(y, x^*) = -(y - x^*)^2$ is the quadratic scoring rule (or, more generally, a proper scoring rule).

Trader i announces y that solves her myopic problem

$$\max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}} E_p[s(y, X) - s(z, X)],$$

where \underline{y}, \bar{y} are the minimum and maximum announcements, respectively.

By announcing $y = z$, she can secure a payoff of zero. Because s is a proper scoring rule, and sets $[\underline{y}, \bar{y}]$, \mathcal{P} are convex and compact, the minimax theorem applies and we can consider the dual problem

$$\min_{p \in \mathcal{P}} \max_{y \in [\underline{y}, \bar{y}]} E_p[s(y, X) - s(z, X)].$$

The inner max problem is solved for each unique prior p , hence the unique solution is $y = E_p[X]$, as in [Ostrovsky \(2012\)](#), because s is a proper scoring rule. We can therefore simplify the problem to

$$\min_{p \in \mathcal{P}} E_p[s(E_p[X], X) - s(z, X)].$$

Note that if the prior is unique, as in the EU framework, the optimal announcement is independent of the previous announcement z . With MEU preferences, the announcement is still the expectation of X , but according to some $p \in \mathcal{P}$ that depends on the previous announcement z , thus introducing path dependence. Moreover, if the previous announcement is the expectation of X according to one of her beliefs $p \in \mathcal{P}$, so that $E_p[X] = z$, then she will repeat the

same announcement securing a payoff of zero.¹⁵ The reason is that she wants to minimize the expected difference between the score of her prediction and the score of the previous announcement because she evaluates this difference using the worst possible probability due to her MEU preferences.¹⁶ However, this creates inertia as traders try to announce as close as possible to the previous announcement given the constraint that their announcement must be the expectation of X according to one of their beliefs. If Trader 1's announcement is the expected value of X according to one of Trader 2's beliefs, then she will repeat it, prompting Trader 1 to do the same so that there is no more updating of information and, consequently, no information aggregation.

Suppose that the true state is ω_1 so that the correct price to be inferred is $x^* = X(\omega_1) = 1$. Moreover, suppose that the initial price of the security is $y_0 = 0$ set by the market maker. Trader 1 is informed that $E_1 = \{\omega_1, \omega_2\}$ has occurred and maximizes her utility myopically. If she announces 1, her payoff is the difference between the expected score of 1 and the expected score of 0. With MEU preferences, she considers the worst-case scenario by choosing p that minimizes her expected payoff. This means that she will choose p that maximizes the score of announcing 0 and minimizes the score of announcing 1. The reason is that scoring rules are "order-sensitive," so that the further away the forecast is from the true expected value, according to the chosen p , the lower is the expectation of the score. This means that she will get a negative payoff by announcing 1. If her announcement is closer to 0, the negative expected payoff decreases, irrespective of which p she uses to evaluate it. In fact, the optimal announcement is to repeat 0 because her payoff will then be zero for all p .

Formally, using Lemma 1 and letting p_{E_1} be the conditional of p given E_1 , the solution to her maximin problem is the same as the solution to her minimax problem. Hence, she minimizes over her priors and for each prior she maximizes her expected utility by announcing the expected value of the security $E_{p_{E_1}}[X]$. We therefore have $\min_{p \in \mathcal{P}} E_{p_{E_1}}[s(E_{p_{E_1}}[X], X(\omega)) - s(0, X(\omega))] = \min_{p \in \mathcal{P}} [p_{E_1}(\omega_1)^2(2 - p_{E_1}(\omega_1) - p_{E_1}(\omega_2))] = \min_{p \in \mathcal{P}} p_{E_1}(\omega_1)^2$. We conclude that the solution is p^1 with $p^1(\omega_1) = 0$ and her prediction is $y_1 = 0$. If the true state was ω_3 , she would know that the intrinsic value of X was 0, and she would announce 0.

The above imply that Trader 2 cannot learn anything from Trader 1's announcement, hence can only rely on her private signal $E_2 = \{\omega_1, \omega_3\}$. Maximizing myopically her utility, she solves $\min_{p \in \mathcal{P}} E_{p_{E_2}}[s(E_{p_{E_2}}[X], X(\omega)) - s(0, X(\omega))] = \min_{p \in \mathcal{P}} [p_{E_2}(\omega_1)^2(2 - p_{E_2}(\omega_1) - p_{E_2}(\omega_3))] = \min_{p \in \mathcal{P}} p_{E_2}(\omega_1)^2$. The solution is again p^1 , with $p^1(\omega_1) = 0$, and her prediction is $y_2 = 0$.

Each trader learns nothing from the other's announcement, which is always 0. Hence, both traders agree on repeating a price of 0 for the security. Given that the intrinsic value of the security at ω_1 is 1, there is no information aggregation, even though their pooled information would reveal that the true state is ω_1 and the intrinsic value is 1. However, if the state is either ω_2 or ω_3 , an initial announcement of 0 will lead to information aggregation as the traders will agree on that price.

We make the following observations. First, the same result of no aggregation can be obtained if the common set of priors is the convex hull of $p^1 = (0, \frac{1}{2}, \frac{1}{2})$ and $p^2 = (\epsilon, \frac{1-\epsilon}{2}, \frac{1-\epsilon}{2})$, where $0 < \epsilon \leq \frac{1}{3}$. Hence, even if belief imprecision is vanishingly small, a prediction market may fail to aggregate information. Second, in this example, there is a belief p that assigns probability zero to the true state ω_1 . However, this is not necessary. Example 1 in Appendix C shows that information aggregation can also fail when all priors have full support.

15. We summarize these points in Lemma 1.

16. In Section 3.3, we provide intuition about this property.

Third, the initial announcement by the market maker is crucial. An announcement of 1 when the true state is ω_1 leads to information aggregation. The reason is that Trader 1 would announce 1 at ω_1 or ω_2 and 0 at ω_3 , thus revealing to Trader 2 that the true state is ω_1 .¹⁷ However, it is impossible for an uninformed market maker to know whether 1 or 0 is the “correct” initial announcement. More importantly, information aggregation fails only when the initial announcement is 0. Nevertheless, this is due to the simplicity of the example. In Appendix C, we show how to easily construct examples where information aggregation fails for multiple initial announcements. Finally, the result of no aggregation does not depend on the quadratic scoring rule, but it is true for all proper scoring rules. The third claim of Lemma 1 shows that as long as the market maker’s announcement is 0 and the expectation of X according to one of Trader 1’s beliefs is 0, then Trader 1 will also announce 0.

To accommodate the case of imprecise probabilities, consider security X' that pays 0 if there is no referendum and 1 otherwise. Then, Trader 1 always knows the value of the security. This implies that, irrespective of the initial announcement, Trader 1 will announce 0 if there is no referendum and 1 otherwise. If the true state is Brexit, Trader 1’s announcement reveals to Trader 2 that the referendum took place. Since she already knows that No Brexit is not true, she deduces that the true state is Brexit and repeats the announcement of 1. Irrespective of what the true state is and what the beliefs of the traders are, the two traders always agree eventually on the intrinsic value of the security, hence there is information aggregation. It is straightforward to check that, whatever the public information is generated from previous announcements, one of the two traders knows the intrinsic value of the security. As we show in Corollary 1, this is a sufficient condition for strong separability.

Security X' aggregates information for any initial announcement of the market maker, irrespective of whether market participants have precise probabilities, or they are ambiguity averse and have multiple priors. Hence, it is robust as compared to the separable securities of [Ostrovsky \(2012\)](#). Moreover, security X' is immune to manipulation by the market maker. We call such securities strongly separable and show that they are always separable, but the converse is not true. Theorems 1 and 2 characterize information aggregation in terms of strongly separable securities for the non-strategic and strategic environments, respectively.

We conclude the example by commenting on the generality and applicability of the MSR. A prediction market with a 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. [Ostrovsky \(2012\)](#) establishes such a justification and Example 2 in Appendix C provides the details for the case of ambiguity aversion. 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 U.K. and the Czech Republic ([The Economist, 2021](#)).¹⁸

17. Note that, in this case, Trader 2 would not be able to do prior-by-prior updating on the specific p that assigns zero probability to state ω_1 . As this case is not the main focus of our example, we nevertheless chose to keep it. We could present here Example 1 of Appendix C with full support priors in order to avoid the issue; however, it is a more complicated example without any further insights.

18. Firms and governments use companies such as Cipher and Cultivate Labs to implement MSR-based prediction markets. See [Cultivate Labs \(2021\)](#) for an explanation of how the logarithmic MSR is implemented in practice and [Schlegel et al. \(2022\)](#) for axiomatic foundations.

3. THE MODEL

In this section, we describe the ambiguity averse preferences of the traders and the MSR trading environment, which, in turn, is based on proper scoring rules (*e.g.* [Brier, 1950](#)). We next distinguish between two cases. In the first case, all traders are myopic so that they only care about the current period's payoff. In the second case, all traders act strategically and care about the future.

3.1. Preferences and updating

Consider a finite state space $\Omega = \{\omega_1, \dots, \omega_I\}$ and let the powerset 2^Ω be the σ -algebra over Ω . Traders are ambiguity averse and have MEU preferences ([Gilboa and Schmeidler, 1989](#)). In particular, each trader evaluates act $f : \Omega \rightarrow \mathbb{R}$ as

$$V(f) = \min_{p \in \mathcal{P}} \int u(f(s)) dp(s),$$

where $u : \mathbb{R} \rightarrow \mathbb{R}$ is a utility index, and \mathcal{P} is a convex and closed subset of $\Delta(\Omega)$. We assume that \mathcal{P} is common among all traders and, without loss of generality, $\bigcup_{p \in \mathcal{P}} \text{Supp}(p) = \Omega$ so that each state is considered possible by some $p \in \mathcal{P}$. Traders are risk-neutral so $u(x) = x$.

The set of traders is $I = \{1, \dots, n\}$. Trader i 's initial private information is represented by partition Π_i of Ω . When the true state is $\omega \in \Omega$, Trader i considers the set of states $\Pi_i(\omega) \subseteq \Omega$ to be possible. As in [Ostrovsky \(2012\)](#), we assume that the join (the coarsest common refinement) of partitions $\Pi = \{\Pi_1, \dots, \Pi_n\}$ consists of all states in Ω so that $\bigcap_{i \in I} \Pi_i(\omega) = \{\omega\}$ for all $\omega \in \Omega$.

In other words, the traders' pooled information always reveals the true state. This implies that, for any two states $\omega_1 \neq \omega_2$, there exists Trader i who can distinguish between them so that $\Pi_i(\omega_1) \neq \Pi_i(\omega_2)$.

We argue that this is a reasonable assumption for two reasons. First, if the conjunction of the traders' private information cannot distinguish between two states, we cannot expect that a security which pays differently in these two states can achieve information aggregation as this would imply that the market has more information than all traders combined.¹⁹ We therefore do not consider securities that pay differently within an element of the coarsest common refinement. Second, given that restriction, it is without loss of generality to consider each element of the coarsest common refinement to be a state rather than a set of states.

When a trader learns event E , her beliefs are \mathcal{P}_E , the prior-by-prior updating of \mathcal{P} .²⁰ This rule is well-defined as long as each prior assigns positive probability to E . We say that measures $p_1, p_2 \in \mathcal{P}$ are mutually absolutely continuous with respect to a collection of events \mathcal{E} if, for all $E \in \mathcal{E}$, $p_1(E) = 0$ if and only if $p_2(E) = 0$. Compact and convex set $\mathcal{P} \subseteq \Delta(\Omega)$ is *regular* with respect to \mathcal{E} if all $p_1, p_2 \in \mathcal{P}$ are mutually absolutely continuous with respect to \mathcal{E} . We interpret \mathcal{E} as the collection of all events that can be revealed when traders make announcements. Hence, if all priors assign positive probability, prior-by-prior updating is well-defined. However, note that regularity does not imply that priors have the same support. For example, it is possible that measure p assigns probability zero to state ω , whereas other measures do not. If event $\{\omega\}$ does not belong to \mathcal{E} , then, \mathcal{P} can be regular with respect to \mathcal{E} .

19. As we explain in Section 3.4, information aggregation means that the price of the security converges to its intrinsic value. Thus, an outside observer, without any private information and just by observing the price, would be able to distinguish between the two states, effectively having more information than all traders combined.

20. This rule is axiomatized in [Pires \(2002\)](#).

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 $X : \Omega \rightarrow \mathbb{R}$. At time $t_1 > t_0$, Trader 1 makes a revised prediction y_1 , then at $t_2 > t_1$ Trader 2 makes her prediction, and so on. At time $t_{n+1} > t_n$, Trader 1 makes another prediction y_{n+1} . Let a_k be the trader that makes a prediction at time t_k . All predictions are observed by all traders. Each prediction y_k is required to be within the set $[\underline{y}, \bar{y}]$, where $\underline{y} = \min_{\omega \in \Omega} X(\omega)$ and $\bar{y} = \max_{\omega \in \Omega} X(\omega)$.

The process repeats until time $t_\infty = \lim_{k \rightarrow \infty} t_k$. At time $t^* > t_\infty$, the intrinsic value $x^* = X(\omega^*)$ is revealed. The traders' payoffs are computed using a scoring rule $s(y, x^*)$, where x^* is the intrinsic value of the security and y is a prediction. A scoring rule is *proper* if, 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 basic MSR (McKelvey and Page, 1990; Hanson, 2003, 2007), a trader is paid for each revision she makes. In particular, her payoff from announcing y_{t_k} at t_k is $s(y_{t_k}, x^*) - s(y_{t_{k-1}}, x^*)$, where $y_{t_{k-1}}$ is the previous announcement and x^* is the intrinsic value of the security. We then say that the trader “buys out” the previous trader's prediction.²¹

The assumption that the payoff is the difference between the previous and the current scores is an inconsequential normalization with EU preferences and myopic traders because a myopic trader will always announce the expected value of the security according to her measure. With MEU preferences, behaviour can change drastically because the measure that minimizes the expected score of the announcement may not be the same as the one that minimizes the expected difference of the two scores. Although this presents some limitations, we argue that the MSR is a reasonable assumption for two reasons. First, the MSR is used in the real world, in both public and corporate prediction markets. Second, 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, as we explain in Section 2 and in Example 2 of Appendix C.

We examine trading in two settings. The myopic or non-strategic is analysed in Section 4, where each trader does not care about future payoffs when making an announcement.²² We denote this setting by $\Gamma^M(\Omega, I, \Pi, X, \mathcal{P}, y_0, \underline{y}, \bar{y}, s)$. The strategic setting is studied in Section 5. 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 $0 \leq \beta < 1$. The total payoff of each trader is the sum of all payments for revisions. We denote this setting by $\Gamma^S(\Omega, I, \Pi, X, \mathcal{P}, y_0, \underline{y}, \bar{y}, s, \beta)$.

3.3. Properties of scoring rules

In the EU framework, the optimal (myopic) choice of y_{t_k} that maximizes $E_p[s(y_{t_k}, x^*) - s(y_{t_{k-1}}, x^*)]$ does not depend on the previous announcement $y_{t_{k-1}}$ because p is fixed. This is no

21. A trader can be guaranteed a payoff of zero by repeating the previous announcement or by abstaining from the market. It would be interesting to separate the two by providing an explicit outside option to the traders. However, such direction is outside the scope of this study and is thus deferred for future research.

22. With EU preferences, this non-strategic setting effectively turns into the communication process of Geanakoplos and Polemarchakis (1982), where traders sequentially announce posterior beliefs about an event. Several other papers extend this process to other aggregate statistics, such as Cave (1983), Sebenius and Geanakoplos (1983), Nielsen (1984), Bacharach (1985), and Nielsen et al. (1990).

longer the case with MEU preferences and multiple priors \mathcal{P} , further complicating our analysis. However, the following lemma establishes three properties that we use heavily.²³ First, the optimal (myopic) announcement is still unique for continuous strictly proper scoring rules. Second, the announcement is the expectation of X according to some belief in \mathcal{P} . Third, the announcement coincides with the previous one if the latter is the expectation of X according to some belief in \mathcal{P} .

Lemma 1. *Let s be a continuous strictly proper scoring rule on $[\underline{y}, \bar{y}]$ and let $z \in [\underline{y}, \bar{y}]$ be an announcement. Then,*

- $y^* \equiv \arg \max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}} E_p[s(y, X) - s(z, X)]$ is unique,
- $y^* = E_p[X]$ for some (not necessarily unique) $p \in \arg \min_{p \in \mathcal{P}} \max_{y \in [\underline{y}, \bar{y}]} E_p[s(y, X) - s(z, X)]$,
- if $z = E_p[X]$ for some $p \in \mathcal{P}$, then $y^* = z$.

As s is a proper scoring rule, hence the optimal announcement is $E_p[X]$, when the expected score is evaluated using p , the second property implies that $y^* = E_p[X]$ for some (not necessarily unique) $p \in \arg \min_{p \in \mathcal{P}} E_p[s(E_p[X], X) - s(z, X)]$. In other words, the maxmin operation simplifies to choosing probability p that minimizes her expected score given that she announces $E_p[X]$.

To provide some intuition for the third property, first note that given an announcement y^* , MEU preferences imply that the trader will evaluate her period payoff by minimizing over all available beliefs in \mathcal{P} . Moreover, scoring rules are “order-sensitive” so that the further away the forecast is from the true expected value, according to the chosen p , the lower is the expectation of the score. These two properties imply that the trader will minimize her expected utility by picking the probability that maximizes the expected score of the previous announcement, which is subtracted from her payoff, and minimizes the expected score of her own announcement, which is added. The only way of counteracting this worst-case scenario is by announcing as close as possible to the previous announcement, given the constraint that it must be the expectation of X according to some $p \in \mathcal{P}$. Moreover, if it is possible to repeat the previous announcement, she will do that and get 0, which is the minimum payoff when announcing the myopic best response.

3.4. Information aggregation

We say that information aggregates 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 1. 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 .

Note that our definition of information aggregation does not specify how prices will evolve in the middle of the game for some t . In fact, it is perfectly possible that prices will diverge widely

23. Lemma 1 is related to a result in Chambers (2008). The proofs are closely related too.

before they start converging to the intrinsic value of the security. Moreover, the intrinsic value $X(\omega)$ is defined objectively, for each ω , and it does not depend on the traders' multiple priors.²⁴

3.5. Strong separability

Ostrovsky (2012) introduced the notion of separable securities, which are sufficient for aggregating information in an environment with EU.

Definition 2. A security X is called non-separable under partition structure Π if there exists probability p and value $v \in \mathbb{R}$ such that:

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

Otherwise, it is called separable.

A security X is non-separable if, for some belief p that assigns positive probability to a state where X does not pay v , all traders agree on its conditional expected value to be v , irrespective of which private signal they have received. In such a case, even if all traders truthfully and repeatedly announce v , no new information is revealed. However, their pooled information reveals the state, hence information aggregation fails.²⁵ To avoid this, the security must be separable. The most common example is the Arrow–Debreu security, which pays 1 at some state and 0 otherwise. Unfortunately, separable securities may not aggregate information with ambiguity aversion as shown in Section 2.

In order to maintain information aggregation in an environment with ambiguity aversion, we need to strengthen the notion of separability. Treating security X as given, let

$$d_{\mathcal{P}}(E, v) = \arg \max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}_E} E_p[s(y, X) - s(v, X)]$$

be the (unique from Lemma 1) myopic announcement that maximizes the trader's current period's utility if her beliefs are \mathcal{P}_E and the previous announcement was v . Note that if $\mathcal{P} = \{p\}$ is a singleton so that we are back to the EU case, $d_{\mathcal{P}}(E, v) = E_p[X \mid E]$ for any v and proper scoring rule s . Hence, the myopic announcement $d_{\mathcal{P}}(E, v)$ under ambiguity is a direct generalization of the myopic announcement under EU, $E_p[X \mid E]$. Below, we generalize the notion of separability by substituting $E_p[X \mid E]$ with $d_{\mathcal{P}}(E, v)$. To save on notation and since security X is fixed throughout the article, we omit it.

Definition 3. A security X is called not strongly separable under partition structure Π and proper scoring rule s if there exist a regular $\mathcal{P} \subseteq \Delta(\Omega)$ with respect to each $\Pi_i, i = 1, \dots, n$, and $v \in \mathbb{R}$ such that:

- (i) $X(\omega) \neq v$ for some $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$,

24. As Harrison and Kreps (1978) comment, it may not be possible to define an objective intrinsic value for some t in the middle of the game if traders do not share a common unique prior, or if they have multiple priors. This means that our model only deals with what happens in the long run. We thank an anonymous referee for pointing out this issue.

25. An example of a non-separable security is provided by Ostrovsky (2012). Let $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ and suppose $X(\omega_1) = X(\omega_4) = 1$, $X(\omega_2) = X(\omega_3) = -1$. Partitions are $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ and $\Pi_2 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. For p that assigns 1/4 at each state, both players always have an expectation of 0, although their pooled information always reveals the intrinsic value of X , which is never 0.

- (ii) $d_{\mathcal{P}}(\Pi_i(\omega), v) = v$ for all $i = 1, \dots, n$ and $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$.

Otherwise, it is called strongly separable.

The interpretation of a not strongly separable security is similar to that of a non-separable security. The only difference is that \mathcal{P} is not a singleton and, as a result, the myopic announcement $E_p[X \mid \Pi_i(\omega)] = v$ under EU is replaced by the myopic announcement $d_{\mathcal{P}}(\Pi_i(\omega), v) = v$ under MEU. However, in both definitions, each trader announces v given that the previous announcement was v and irrespective of the private signal that she has received. We also require that \mathcal{P} is regular with respect to each Trader's partition so that prior-by-prior updating is well-defined.

A potential issue about the definition of strong separability is that it depends on the particular scoring rule because $d_{\mathcal{P}}(E, v) = \arg \max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}_E} E_p[s(y, X) - s(v, X)]$. This is not the case for separability, which only depends on the information structure. Proposition 2, discussed later in this section, establishes that strong separability is also independent of the particular continuous strictly proper scoring rule.

In the example of Section 2, the Arrow–Debreu security is not strongly separable given the information structure. To see this, note that condition (ii) in the definition is satisfied for all states with $v = 0$. Since some priors put strictly positive probability to ω_1 and $X(\omega_1) = 1 \neq v$, condition (i) is also satisfied.

Observe that if a security is non-separable (for some prior p), then it is not strongly separable as well (for $\mathcal{P} = \{p\}$). This means that strong separability implies separability. Moreover, the converse is not true as shown in Section 2. As we discuss after Proposition 2, for any information structure, there exists a strongly separable security. For example, consider state space $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and security X with $X(\omega_1) = X(\omega_2) = 1$, $X(\omega_3) = 0$. Under the partition structure $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3\}\}$, $\Pi_2 = \{\{\omega_1, \omega_3\}, \{\omega_2\}\}$ and any continuous proper scoring rule, X is strongly separable.

Ostrovsky (2012) proposes a useful characterization of separable securities. 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 1 (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 $X(\omega) \neq v$,*

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

We provide a similar but stronger condition that characterizes strong separability. It specifies that, given any v and conditional on any event E where X is never equal to v , there is a trader who knows at some state in E that X is either always above or always below v . We can interpret E as the public information that is revealed by hearing the previous announcements, and v as the current price of the security. Hence, the condition requires that, at any period, at least one trader knows whether the intrinsic value of the security is either higher or lower than the current price. Note that this trader may not be the one who makes the announcement in the next period.

Proposition 2. *Security X is strongly separable under partition structure Π if and only if for any $v \in \mathbb{R}$, for any non-empty event $E \subseteq \{\omega \in \Omega : X(\omega) \neq v\}$, there exists Trader i , state $\omega \in E$*

and $\lambda \in \mathbb{R}$ such that for all $\omega' \in \Pi_i(\omega) \cap E$,

$$(X(\omega') - v)\lambda > 0.$$

In the [Supplementary Appendix](#), we show that the same condition characterizes strongly separable securities in the much more general framework of Uncertainty Averse preferences. Hence, the set of strongly separable securities is the same and all the properties discussed next (Section 3.6) apply to Uncertainty Averse preferences as well. This is surprising because, when we move from EU to MEU, the set of separable securities is a strict subset of the set of strongly separable securities. We discuss the intuition behind this result in the [Supplementary Appendix](#).

3.6. Properties of strongly separable securities

In order to better understand strongly separable securities, we establish the following properties. Fix an information structure. First, there always exists a non-constant, strongly separable security. In particular, we can construct one which predicts all possible events (Lemma 2). This means that, just by observing the price of the security and as it converges to the security's intrinsic value due to information aggregation, an outside observer learns whether the event has occurred or not. Second, we find a sufficient condition to easily check whether a security is strongly separable (Corollary 1). It requires that for any event E , there exists at least one trader who would know the security's value if she was informed of that event.²⁶ As with Proposition 2, we can interpret E as the public information that is revealed by the previous announcements. This condition allows us to easily construct strongly separable securities for any information structure. We describe such an algorithm below. The final property is an impossibility result and one of our main contributions (Proposition 3). In sharp contrast to the EU framework with separable securities, in the framework with MEU preferences, there does not exist a security that is strongly separable for all information structures.

We say that an event E is predictable by security X if the values it assigns to states in E are different from the values it assigns to states not in E . Formally, if $\omega \in E$ and $\omega' \in E^c$, then $X(\omega) \neq X(\omega')$. We then say that X is informative for E because if the price of X converges to its intrinsic value at all states, then it will be revealed whether E has occurred. For the EU model, [Chen *et al.* \(2012\)](#) show in Theorem 5 that there always exists a separable security that is informative for all events. We show that the same is true for the MEU model. We construct a sequence of strongly separable securities, X_1, \dots, X_n , where the collection of predictable events by X_k is larger than that by X_{k-1} for $k = 2, \dots, n$. The last security X_n , where n is the number of Traders, assigns a different value to each state, hence all events are predictable.

- (1) Fix the order of Traders $1, 2, \dots, n$. Order Trader 1's partition elements $\{\Pi_1(\omega)\}_{\omega \in \Omega}$ from 1 to k_1 . If $\Pi_1(\omega)$ is the j th partition element, assign value $X_1(\omega') = j$ to all $\omega' \in \Pi_1(\omega)$. As X_1 provides a different payoff to each of Trader 1's partition elements, X_1 can predict all events in Trader 1's partition together with any of their unions.
- (2) Security X_2 assigns a different value to each event in the collection $\{\bigcap_{i=1,2} \Pi_i(\omega)\}_{\omega \in \Omega}$. Order 2's partition elements, $\{\Pi_2(\omega)\}_{\omega \in \Omega}$ from 1 to k_2 . If $X_1(\Pi_1(\omega)) = m$ and $\Pi_2(\omega)$ is

26. In practice, one needs to check only events for which the security specifies different values; as for the rest, the condition is automatically satisfied.

- the j th partition element, assign value $X_2(\omega') = m + \frac{j-1}{k_2}$ to all $\omega' \in \bigcap_{i=1,2} \Pi_i(\omega)$. Security X_2 predicts all events in $\{\bigcap_{i=1,2} \Pi_i(\omega)\}_{\omega \in \Omega}$ together with any of their unions. Hence, it predicts more events than X_1 .
- (3) Inductively, for security X_l we order l 's partition elements $\{\Pi_i(\omega)\}_{\omega \in \Omega}$ from 1 to k_l . If $X_{l-1}(\bigcap_{i=1, \dots, l-1} \Pi_i(\omega)) = m$, the next highest value of X_{l-1} is m' and $\Pi_l(\omega)$ is the j th partition element, then assign value $X_l(\omega') = m + \frac{j-1}{k_l}(m' - m)$ to all $\omega' \in \bigcap_{i=1, \dots, l} \Pi_i(\omega)$. By construction, from security X_l to security X_{l+1} the ordering of states is preserved. That is, if $X_l(\omega) < X_l(\omega')$, then $X_{l+1}(\omega) < X_{l+1}(\omega')$.
- (4) The final security X_n assigns a different value to each state because of our assumption that $\bigcap_{i \in I} \Pi_i(\omega) = \{\omega\}$ for all $\omega \in \Omega$. Hence, it can predict any event.

The following lemma shows that these securities are strongly separable. Hence, we can always construct non-trivial strongly separable securities, some of which can predict all possible events. However, it is important to note that not all securities that assign a different value to each state, and therefore can predict any event, are strongly separable. For a counter example, see Example 1 in Appendix C.

Lemma 2. *Securities X_1, \dots, X_n are strongly separable.*

The following Corollary provides a sufficient condition for strong separability. It requires that conditioning on any event E , there is at least one trader who knows the value of the security.

Corollary 1. *Suppose that for any event E , there exist Trader i and state $\omega \in E$ such that $\Pi_i(\omega) \cap E \subseteq X^{-1}(k)$ for some k . Then, security X is strongly separable under partition structure Π .*

Using this Corollary, we can construct a strongly separable security in the following way given any information structure. First, fix an order of traders $T = 1, 2, \dots$. A specific trader may appear more than once in T and its cardinality is weakly less than the cardinality of Ω . Pick Trader 1 and a state $\omega_1 \in E_1 \equiv \Omega$, assigning value $X(\omega') = k_1$ for all $\omega' \in \Pi_1(\omega_1)$. Then, pick Trader 2 and a state $\omega_2 \in E_2 = E_1 \setminus \Pi_1(\omega_1)$, assigning value $X(\omega') = k_2$ for all $\omega' \in \Pi_2(\omega_2) \cap E_2$. This process continues using $\omega_{i+1} \in E_{i+1} = E_i \setminus \Pi_i(\omega_i)$, assigning value $X(\omega') = k_{i+1}$ for all $\omega' \in \Pi_{i+1}(\omega_{i+1}) \cap E_{i+1}$ for $i \geq 2$ until E_{i+1} becomes empty. To see how we can apply Corollary 1, take any event E . If there exists $\omega \in E \cap \Pi_1(\omega_1) \neq \emptyset$, then $\Pi_1(\omega) \cap E \subseteq X^{-1}(k_1)$. If $E \cap \Pi_1(\omega_1) = \emptyset$ but there is some state $\omega \in \Pi_2(\omega_2) \cap E$, then $E \subseteq E_2$ and $\Pi_2(\omega) \cap E \subseteq X^{-1}(k_2)$. Continuing inductively, we can find some Trader i and state $\omega \in E$ such that $\Pi_i(\omega) \cap E \subseteq X^{-1}(k_i)$.

The last question is whether there exists a security that is strongly separable for all information structures. Recall that there are several securities that are separable for all information structures, such as the Arrow–Debreu. However, the following proposition shows that there is no security which is strongly separable for all information structures.

Proposition 3. *If state space Ω has at least three states, there is no (non-constant) security X which is strongly separable under all partition structures $\Pi = \{\Pi_1, \dots, \Pi_n\}$, where the join of Π consists of singleton sets.*

As we show in subsequent subsections (Theorems 1 and 2), strong separability is not only sufficient but also necessary for information aggregation under ambiguity. This suggests a negative result as there is no security that aggregates information for all information structures in contrast to the EU case. In other words, if an outside observer does not know the traders' information structure, there is no way of being sure that a particular security is strongly separable and therefore will aggregate information.

More interestingly, a security which has been successful at aggregating information (because of the particular information structure), may subsequently fail to do so, once the composition of the traders and their information changes. Although this negative result is shown for the specific case of prediction markets, it is also a negative result for financial markets in general. This means that markets may fail to predict events and that prices do not incorporate all available information. Moreover, as we show in [Proposition 5 in the Supplementary Appendix](#), the set of strongly separable securities does not change in the much more general framework of Uncertainty Averse preferences. Hence, this negative result is robust.

4. MYOPIC TRADERS

Let $\Gamma^M(\Omega, I, \Pi, X, \mathcal{P}, y_0, \underline{y}, \bar{y}, s)$ be an environment with myopic traders who only care about their period t payoff when making an announcement at t . Suppose ω^* is the true state and y_0 is the market maker's initial announcement. At time t_1 , Trader 1 announces her prediction $y_1 = d_{\mathcal{P}}(\Pi_1(\omega^*), y_0) = \arg \max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}_{\Pi_1(\omega^*)}} E_p[s(y, X) - s(y_0, X)]$. As mentioned above, y_1 depends on the market maker's announcement y_0 , which is not the case with EU.

The prediction of any trader is public, therefore, the new information revealed refines the information partitions of all other traders. In particular, the initial public information at t_0 is $\mathcal{F}^0(\omega^*) = \Omega$. At t_1 , Trader 1 announces $y_1 = d_{\mathcal{P}}(\mathcal{F}^0(\omega^*) \cap \Pi_1(\omega^*), y_0)$. The updated public information is $\mathcal{F}^1(\omega^*) = \{\omega' \in \mathcal{F}^0(\omega^*) : d_{\mathcal{P}}(\mathcal{F}^0(\omega^*) \cap \Pi_1(\omega'), y_0) = y_1\}$. Note that from Lemma 1, the announcement is unique, hence $\mathcal{F}^1(\omega^*)$ is well-defined. Trader i 's new private information is $\mathcal{F}^1(\omega^*) \cap \Pi_i(\omega^*)$.

Trader 2 is next to make a public announcement and her private information is $\mathcal{F}^1(\omega^*) \cap \Pi_2(\omega^*)$. At t_2 , she announces $y_2 = d_{\mathcal{P}}(\mathcal{F}^1(\omega^*) \cap \Pi_2(\omega^*), y_1)$ and the updated public information is $\mathcal{F}^2(\omega^*) = \{\omega' \in \mathcal{F}^1(\omega^*) : d_{\mathcal{P}}(\mathcal{F}^1(\omega^*) \cap \Pi_2(\omega'), y_1) = y_2\}$. Trader 3 updates her private information to $\mathcal{F}^2(\omega^*) \cap \Pi_3(\omega^*)$, makes an announcement, and the process goes on. More generally, player $a_k = i$ at time t_k has private information $F = \mathcal{F}^{k-1}(\omega^*) \cap \Pi_i(\omega^*)$ and announces $y_k = d_{\mathcal{P}}(F, y_{k-1})$.

Let $\mathcal{E} = \{\mathcal{F}^k(\omega) \cap \Pi_{a_k}(\omega)\}_{k \geq 0, \omega \in \Omega}$ be the collection of all events on which the traders update their beliefs given that it is their turn to make an announcement. We say that Γ^M is *regular* if \mathcal{P} is regular with respect to \mathcal{E} .

4.1. Information aggregation

Our first main result is to fully characterize information aggregation in an environment with myopic and ambiguity averse traders.

Theorem 1. *Fix security X , information structure Π and continuous strictly proper scoring rule s . Information aggregates for any regular $\Gamma^M(\Omega, I, \Pi, X, \mathcal{P}, y_0, \underline{y}, \bar{y}, s)$ if and only if X is strongly separable.*

To provide some intuition, we describe briefly the steps of the proof. We first show that the public (and therefore private) information is no longer updated after some time t . This is a direct

consequence of the finiteness of the state space so that all possible states are within a common knowledge event \mathcal{F} .

Second, traders agree on the announcement and they stop updating it. If Dynamic Consistency was satisfied, as it is the case with EU and Bayesian updating, this step would be straightforward. Since Trader i optimally announces y_i in each of her partition cells, irrespective of the previous announcement, Dynamic Consistency, and the law of iterated expectations imply that it is optimal to announce y_i if her information was just \mathcal{F} . Given that this is true for all traders, common priors imply that their announcements must coincide.

However, with MEU preferences and prior-by-prior updating, Dynamic Consistency is violated and there is no longer separability across states as a different belief might be picked at each partition cell. Hence, we cannot apply the law of iterated expectations.²⁷ Moreover, the myopic prediction depends not only on the private information, as in the EU case, but also on the previous trader's prediction. Since there are many possible myopic predictions, it could be the case that traders engage in a never-ending cycle of revised predictions, even though their private information does not change. We show that this does not occur because a monotonicity property of the scoring rule and ambiguity aversion imply that Trader i will want to announce as close as possible to the previous announcement in order to minimize the worst-case scenario. Moreover, the set of all myopic announcements is fixed given \mathcal{F} and do not depend on the previous announcement. We therefore have only two cases. First, there is an announcement that is common to all traders. Once a trader makes this announcement, everyone else will repeat it so there will be agreement. Second, two traders disagree so much that i 's maximum myopic prediction, according to her posterior beliefs, is lower than j 's minimum myopic prediction, according to her posterior beliefs. But if this is true for all partition cells, it will also be true when conditioning on \mathcal{F} , which is impossible because there is at least one common prior. Effectively, this is a generalization of the result of Aumann (1976) for MEU preferences; that is, "we cannot agree to disagree too much."²⁸

Finally, suppose that all traders agree on the prediction, which is the expected value of the security for some posterior for all states in \mathcal{F} . Then, the definition of strong separability implies that this can only happen if there is no uncertainty about the value of the security. That is, all states in \mathcal{F} prescribe the same value for the security, which is then equal to the common prediction and there is information aggregation.

5. STRATEGIC TRADERS

Consider a game $\Gamma^S(\Omega, I, \Pi, X, \mathcal{P}, y_0, \underline{y}, \bar{y}, s, \beta)$, 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 , $[\underline{y}, \bar{y}]$ is the set of possible announcements that a player can make, and β is the common discount rate.

Let $H^k = (y_1, \dots, y_k)$ be a history of announcements up to time t_k , and H^0 be the empty history. Given any two histories $H^k = (y_1, \dots, y_k)$ and $H^l = (y'_1, \dots, y'_l)$, let (H^k, H^l) be their concatenation. Although traders have multiple priors over Ω , a mixed strategy consists of randomizing using a unique probability distribution. Player i trades at periods t_{i+nk} , $k \in \mathbb{N}$, hence $a_{i+nk} = i$. Her mixed strategy at time t_k is a measurable function $\sigma_{i,k} : \Pi_i \times [\underline{y}, \bar{y}]^{k-1} \times$

27. See Galanis (2021) for a discussion of Dynamic Consistency in a general framework with multiple beliefs and convex preferences.

28. In an environment with ambiguity aversion, several papers extend the no trade theorems of Aumann (1976), such as Dominiak and Lefort (2013, 2015), Carvajal and Correia-da Silva (2010), and Kajii and Ui (2005, 2009), whereas Condie and Ganguli (2011) show the existence and robustness of partially revealing REE.

$[0, 1] \rightarrow [y, \bar{y}]$. It specifies an announcement y_k given the element of her partition, the history of announcements (y_1, \dots, y_{k-1}) up to time t_k , and the realization of random variable $\iota_k \in [0, 1]$, which is drawn from the uniform distribution. These draws are independent of each other and of the true state ω . The full state is $\phi = (\omega, \iota_1, \iota_2, \dots)$ and describes the initial uncertainty and the randomizations of the players. Let $\Phi = \Omega \times [0, 1]^{\mathbb{N}}$ be the full state space. Player i 's strategy, denoted σ_i , is a set of strategies at all times where it is her turn to make an announcement. Let $\sigma = (\sigma_1, \dots, \sigma_n)$ be a profile of strategies.

A profile of strategies σ and a full state ϕ determine a sequence of predictions on-path, which we denote $y_1(\sigma, \phi), y_2(\sigma, \phi), \dots$. Let $H^k(\sigma, \phi) = (y_1(\sigma, \phi), \dots, y_k(\sigma, \phi))$ be the history at t_k generated by σ and ϕ on-path. Given a history H^{k-1} , which may not be on-path, let $y_{k-1+m}(\sigma, \phi | H^{k-1})$ be the announcement at time t_{k-1+m} if traders play according to strategy profile σ and full state ϕ , from t_k onwards, where $m \geq 0$. We denote by $H^{k-1+m}(\sigma, \phi | H^{k-1}) = (H^{k-1}, y_k(\sigma, \phi | H^{k-1}), \dots, y_{k-1+m}(\sigma, \phi | H^{k-1}))$ the history that is generated by these announcements.

Let $\omega(\phi)$ and $\iota_k(\phi)$ be the first and $(k+1)$ th components of full state $\phi = (\omega, \iota_1, \dots)$, respectively. At time t_k , Trader i knows component $\iota_l(\phi)$, which denotes the realization of the random variable at t_l , if $a_l = i$ and $l \leq k$. Her private information at time t_k and state ϕ is $\Pi_i^k(\phi) = \Pi_i(\omega(\phi)) \times [0, 1]^k \cap [\phi' : \iota_l(\phi') = \iota_l(\phi) \text{ for all } l \leq k \text{ with } a_l = i]$. Trader i 's information set at decision node (H^{k-1}, ϕ) is denoted $\mathcal{I}(H^{k-1}, \phi) = \Pi_i^k(\phi)$. Let \mathcal{I}_i^k be the collection of all information sets for i at time t_k , and \mathcal{I} be the collection of all information sets.

The public information revealed at time t_{k+m} , $m \geq 0$, after history H^k and given that traders play from t_{k+1} according to σ at full state ϕ is

$$\mathcal{F}^{k+m}(\sigma, \phi | H^k) = \{\phi' \in \Phi : H^{k+m}(\sigma, \phi | H^k) = (y_{k+1}(\sigma, \phi' | H^k), \dots, y_{k+m}(\sigma, \phi' | H^k))\}.$$

If $k = 0$, then we denote by $\mathcal{F}^m(\sigma, \phi | H^0) = \mathcal{F}^{k+m}(\sigma, \phi)$ the public information at t_m that is revealed when everyone plays on-path.

Player $a_{k+m} = i$, who makes an announcement at t_{k+m} , can combine the public information $\mathcal{F}^{k+m}(\sigma, \phi | H^k)$ with her private information $\Pi_i^{k+m}(\phi) \subseteq \Phi$ in order to form her updated private information. We denote the player's updated private information given strategy σ , state ϕ and history H^k , by

$$\mathcal{F}_i^{k+m}(\sigma, \phi | H^k) = \Pi_i^{k+m}(\phi) \cap \mathcal{F}^{k+m}(\sigma, \phi | H^k).$$

A system of beliefs is a collection of compact and convex sets of beliefs, one for each information set.

Definition 4. A system of beliefs is a tuple $\mathcal{P} = \{\mathcal{P}(\mathcal{I})\}_{\mathcal{I} \in \mathcal{I}}$ such that each $\mathcal{P}(\mathcal{I})$ is compact and convex.

To save on notation, we denote the beliefs $\mathcal{P}(\mathcal{I}(H^{k-1}, \phi))$ of agent i who announces at t_k and information set $\mathcal{I}(H^{k-1}, \phi)$ as $\mathcal{P}(H^{k-1}, \phi)$.²⁹

We now define the continuation payoff of player a_k at decision node (H^{k-1}, ϕ) . Note that we define this payoff also in nodes that are not reached given strategy profile σ .

29. If $k = 0$, then we are at the initial time t_0 so that a_0 denotes each $i \in I$ and $\mathcal{P}(H^0, \phi) = \mathcal{P}$.

Definition 5. The continuation payoff of player $a_k = i$ at time t_k and state ϕ , given strategy profile σ , history H^{k-1} and system of beliefs \mathcal{P} is

$$V_i(H^{k-1}, \phi, \sigma, \mathcal{P}) = \min_{p \in \mathcal{P}(H^{k-1}, \phi)} E_p \left[\sum_{m=0}^{\infty} \beta^{nm} (s(y_{k+nm}(\sigma, \phi | H^{k-1}), X(\phi)) - s(y_{k+nm-1}(\sigma, \phi | H^{k-1}), X(\phi))) \right].$$

The expectation is taken over Φ and we set $X(\phi) = X(\omega(\phi))$, where $\omega(\phi) \in \Omega$ is the first component of ϕ . To save on notation, we sometimes denote V_i with V as it is clear in each time t_k who is making the announcement. The only exception is at time t_0 , where only the market maker has made an announcement and all traders have received their private information. In that case, we denote i 's ex ante payoff as $V_i(H^0, \phi, \sigma, \mathcal{P})$.

5.1. Revision-Proof equilibrium

In this section, we define the notion of a Revision-Proof equilibrium and use it to show that strongly separable securities characterize information aggregation. An issue that arises in incomplete information games with ambiguity averse players is that their preferences may not be dynamically consistent. This means that an ex ante optimal plan may be considered suboptimal by the same player at a subsequent period, therefore, choosing not to follow it.³⁰

One way of solving the issue of dynamic inconsistency is by imposing a solution concept similar to the consistent planning of [Strotz \(1955\)](#), which is a refinement of backward induction.³¹ Effectively, the decision maker takes into account the constraint that her future selves might have different preferences and may not follow through a plan that is optimal now. Since in our environment there are infinitely many periods, we cannot impose backward induction so the generalization would be to check for one-shot deviations.

Before formalizing the notion of consistent planning, we define consistency, which imposes prior-by-prior updating at all decision nodes whenever possible.³²

Definition 6. Pair (σ, \mathcal{P}) is consistent if, for any full state $\phi \in \Phi$, history H^k , $k \geq 0$ and player $a_k = i$,

- (i) $\mathcal{P}(H^k, \phi)$ is regular with respect to $F = \mathcal{F}_i^{k+n}(\sigma, \phi | H^k)$,

30. [Pahlke \(2022\)](#) studies games with incomplete information and MEU preferences, finitely many actions and periods. She shows the existence of a Sequential equilibrium with rectangular priors ([Epstein and Schneider, 2003](#)), thus ensuring Dynamic Consistency. In the Smooth Ambiguity model, [Hanany et al. \(2020\)](#) show the existence of a Sequential equilibrium in a setting with finite actions and periods, using the smooth rule ([Hanany and Klibanoff, 2007, 2009](#)). Few other papers study equilibrium notions in general dynamic games under ambiguity, such as [Eichberger et al. \(2019\)](#) and [Battigalli et al. \(2019\)](#). [Ellis \(2018\)](#) argues that in games with incomplete information and MEU preferences that satisfy Dynamic Consistency, consequentialism and a common set of priors \mathcal{P} , players act as if they have EU preferences. [Pahlke \(2022\)](#) avoids such a criticism by allowing for different priors.

31. Consistent planning was further developed by [Peleg and Yaari \(1973\)](#) and [Goldman \(1980\)](#). [Siniscalchi \(2011\)](#) provides behavioural foundations in a single-agent setting. Specific applications with MEU preferences, prior-by-prior updating and some form of consistent planning are provided, among others, by [Bose and Daripa \(2009\)](#), [Bose and Renou \(2014\)](#), [Kellner and Le Quement \(2017, 2018\)](#), and [Beauchêne et al. \(2019\)](#).

32. Consistency adapts the standard definition of consistency in a Perfect Bayesian Equilibrium ([Fudenberg and Tirole, 1991](#)). [Bonanno \(2013, 2016\)](#) examines the relationship between Perfect Bayesian Equilibrium and Sequential Equilibrium, by providing a qualitative notion of AGM-consistency, which is based on the theory of belief revision introduced by [Alchourrón et al. \(1985\)](#).

- (ii) If $\bigcup_{p \in \mathcal{P}(H^k, \phi)} \text{Supp}(p) \cap F \neq \emptyset$, then $\mathcal{P}(H^{k+n}(\sigma, \phi \mid H^k), \phi)$ is the prior-by-prior updating of $\mathcal{P}(H^k, \phi)$ given F .³³

At decision node (H^k, ϕ) , the beliefs of player $a_k = i$ are $\mathcal{P}(H^k, \phi)$. Given that everyone plays according to σ and ϕ for one round of n announcements, i 's private information is updated using new information $F = \mathcal{F}_i^{k+n}(\sigma, \phi \mid H^k)$. Consistency requires that beliefs $\mathcal{P}(H^k, \phi)$ are regular with respect to F and that there is prior-by-prior updating whenever possible.

Definition 7. Consistent pair (σ^*, \mathcal{P}) is a Consistent-Planning equilibrium if there is no decision node (H^{k-1}, ϕ) , player $a_k = i$ and alternative strategy $\sigma = (\sigma_i, \sigma_{-i}^*)$ with $\sigma_{i,k'} = \sigma_{i,k'}^*$ for all $k' \neq k$, such that

$$V(H^{k-1}, \phi, \sigma, \mathcal{P}) > V(H^{k-1}, \phi, \sigma^*, \mathcal{P}).$$

This solution concept (for infinitely many periods) has not yet been studied in games with incomplete information and ambiguity averse preferences. However, in complete information games with time-inconsistent preferences, [Asheim \(1997\)](#) and [Ales and Sleet \(2014\)](#) argue against such a solution concept and provide a refinement, Revision-Proofness, which we adapt in our setting.

A consistent pair (σ^*, \mathcal{P}) is a Revision-Proof equilibrium if it is immune to any “collective” deviations by a trader and her future selves, where every future self evaluates the deviation given her updated beliefs and preferences. This latter condition is crucial because of dynamic inconsistency. Even if Trader i considers a deviation profitable at time t_k , it does not mean that her future self, after r rounds, will also find it profitable at t_{k+nr} .

Definition 8. Consistent pair (σ^*, \mathcal{P}) is a Revision-Proof equilibrium if there is no decision node $(H^{k-1}(\phi, \sigma^*), \phi)$, player $a_k = i$ and alternative strategy $\sigma = (\sigma_i, \sigma_{-i}^*)$ such that for all $r \geq 0$ and H^{nr} ,

$$V((H^{k-1}(\phi, \sigma^*), H^{nr}), \phi, \sigma, \mathcal{P}) \geq V((H^{k-1}(\phi, \sigma^*), H^{nr}), \phi, \sigma^*, \mathcal{P})$$

with the inequality strict for at least one H^{nr} .

Our concept has three differences from that of [Asheim \(1997\)](#) and [Ales and Sleet \(2014\)](#). First, they only consider complete information games, hence they do not specify how beliefs are updated. Second, they consider deviations from any set of subsequent players, whereas we only check deviations of a single player and her future selves. Third, they check deviations from any history, not just the one that is followed on-path.

Note that, as is the case with complete information games, Revision-Proof equilibria may not always exist. In the [Supplementary Appendix](#), we show that Revision-Proof equilibria exist when the game is continuous at infinity. A game is continuous at infinity if strategies that only differ in the distant future have negligible impact on the utility of any player. As with [Ostrovsky \(2012\)](#), this is achieved by shortening the time period t_k as $k \rightarrow \infty$, so that the discount factor decreases.

Our main result in the strategic environment is that strongly separable securities aggregate information in all Revision-Proof equilibria.

Theorem 2. Fix information structure Π and bounds $[\underline{y}, \bar{y}]$.

33. If $k = 0$, then we are at the initial time t_0 so that a_0 denotes each $i \in I$ and $\mathcal{P}(H^0, \phi) = \mathcal{P}$.

- (i) If security X is strongly separable under Π , then for any Γ^S and any Revision-Proof equilibrium, information aggregates.
- (ii) If security X is not strongly separable under Π , then there exist game Γ^S and a Revision-Proof equilibrium such that information does not aggregate.

6. EXPERIMENT

Our experimental design focused on three dimensions. The first dimension was whether beliefs about events were precise to reflect EU preferences or imprecise to reflect MEU preferences. The second dimension related to the type of security that was traded: separable securities, such as Arrow–Debreu securities,³⁴ or strongly separable securities. The third dimension related to the initial price of the security set by the uninformed market maker: we allowed for two initial prices. In summary, we applied a 2X2X2 experimental design to examine the impact on information aggregation of the market type, security type, and initial price.

6.1. Experimental design

Initially, subjects received 6,000 Experimental Currency Units (ECUs) as a show-up fee. The conversion was 2,000 ECUs for €1. There were three parts in the experimental instructions. In the first part, we measured subjects' risk attitudes. Specifically, we used a variant of the Eckel–Grossman test (Eckel and Grossman, 2002, 2008), where subjects were presented with five gambles of varying riskiness and were required to select the one they prefer. In the second part, the game play took place. The instructions here accommodated the underlying assumptions about the nature of beliefs, type of security, and initial price. The second part was the only part that differed across the treatments conducted. In the third part, subjects were asked to complete a questionnaire about their demographic characteristics. With the conclusion of the experimental session, subjects were paid in cash by the experimenter. The experimental instructions are included in the [Supplementary Appendix](#).

In the game-play stage, subjects were recruited to play the role of traders in prediction markets forecasting the value of a stock.³⁵ The traders' forecast could take any integer value from 0 to 100 inclusive. The stock value was a binary outcome taking either the value of high (*i.e.* 100) or low (*i.e.* 0). To determine the stock value and, thereby, the payoffs of the traders in the sequential trading (see below for more details), a random draw took place in the beginning of the round. Specifically, a coloured ball was drawn from a fictional urn containing 90 coloured balls. The colours of the balls in the urn $\{red, green, blue\}$ represented states that mapped onto a high- or low-stock value. Prior to the start of trading, subjects were provided with information on the colour composition of the urn, the mapping of colours to a high- or low-stock value, a private signal about the colour of the drawn ball, and the initial price of the stock. The information on the colour composition of the urn reflected the market type (unique priors and EU preferences or multiple priors and MEU preferences), whereas the mapping of colours to a high- or low-stock value reflected the available type of traded securities (separable or strongly separable). The information structure, presented to subjects in a tabular form as shown in Table 2, was *fixed* and *common knowledge* in all treatments. This information was explicitly discussed in the

34. See also Healy *et al.* (2010), Choo *et al.* (2019, 2022), and Page and Siemroth (2021) for laboratory experiments that study information aggregation with Arrow–Debreu securities.

35. In lieu of the word “security,” in the experimental instructions, we used the word “stock,” which is contextually more familiar to most individuals.

TABLE 2
Information structure

Ball drawn	Private information	
	Trader 1	Trader 2
Red	Not blue	Not green
Green	Not blue	Green
Blue	Blue	Not green

Notes: The table displays the state-conditional signals that were provided to the two traders. Even though the structure was common knowledge, the trader's signal in each round was private.

instructions.³⁶ Furthermore, before trading, subjects were also informed of the initial stock price (0 or 50).

In the treatments with unique priors, subjects were given the exact composition of the urn. Specifically, they were told that there are 90 balls in the urn, where 30 of those are red, 30 are green, and 30 are blue. This information allowed subjects to formulate precise beliefs about events and have EU preferences.³⁷ Henceforth, this market is referred to as *EU*. In the treatments with multiple priors, subjects were not given the exact composition of the urn. In the treatment with multiple priors and separable securities, subjects were informed that the urn contains 90 balls, where between 0 and 30 are red balls, between 20 and 70 are green balls, and between 20 and 70 are blue balls. This setting mimics the example in Section 1 (and Section 2), where one belief puts probability 0 on the first state, which we call “red” in the experiment. In the treatment with multiple priors and strongly separable securities, subjects were informed that the urn contains 90 balls, where between 1 and 30 are red balls, between 20 and 69 are green balls, and between 20 and 69 are blue balls. Notice that, here, we change the composition so that no belief puts zero probability on the red state. The reason is that since our theory predicts that there will be information aggregation on the red state, we need to apply prior-by-prior updating when a red ball is drawn and, therefore, all beliefs must assign strictly positive probability on red. Providing partial information about the composition of the urn enables ambiguity averse subjects to formulate multiple priors that give rise to the MEU preferences. Henceforth, this market is referred to as *Amb*.

The second treated variable was the type of the security. In the case of separable securities, we informed subjects that if the red ball was drawn, then, the stock value would be high (*i.e.* 100), otherwise the stock value would be low (*i.e.* 0). Hence, this is a standard Arrow–Debreu security. In the case of strongly separable securities, we informed subjects that if the red or green ball was drawn, then, the stock value would be high (*i.e.* 100), otherwise the stock value would be low (*i.e.* 0). Note that the security is constant on the partition cells of Trader 1. From Proposition 2, this is a sufficient condition for the security to be strongly separable.

The initial price of the security was another treated variable. In the myopic setting, theoretically, the two security types exhibit the same information aggregation, in *every* single state, for all initial prices with the exception of 0; at the 0 initial price, the information aggregation

36. For instance, subjects were told that “if the drawn ball is red, Trader 1 will be informed that the drawn ball is not blue, whereas Trader 2 will be informed that the drawn ball is not green.” Analogous descriptions were provided for the other colours.

37. It should be noted that simply providing explicit information about the composition of the urn does not guarantee that subjects have EU preferences. Rather, the assumption that individuals have EU preferences, in this setting, is a joint hypothesis of the alternative tests that are being run. We thank an anonymous referee for pointing out the necessity of this clarification.

should still be the same across the two security types in the green and blue states, but worse in the red state for the separable security with ambiguity.^{38,39} We thus chose to investigate experimentally information aggregation at the initial price of 0 as well as at an initial price where the two security types perform the same. We chose 50 as the midpoint between 0 and 100.

Subjects were asked to take part in 12 rounds of prediction markets. In each round, traders made sequential predictions about the stock value. Specifically, Trader 1 would make a prediction in the first trading period, then Trader 2 would provide her prediction in the second trading period, then Trader 1, and so on and so forth. Although the number of rounds was common knowledge, the number of trading periods *within* each round was unknown. However, subjects were informed that there was a 95% chance of having an extra trading period within a given round.⁴⁰

The draws for the number of trading periods within each round were done *ex ante* to ensure that all treatments would have the same number of trading periods. The states were also drawn *ex ante* and hard coded. We did so to enable a consistent comparison across treatments without invoking variability in learning effects. The actual numbers of trading periods in each round were $\{(4,16,17,12,9,15,12,8,17,16,21,5)\}$. Thus, the round with the highest number of trading periods was round #11 with 21 trading periods, and the round with the lowest number was round #1 with 4 trading periods. The realized states were $\{(\text{Red}, \text{Blue}, \text{Blue}, \text{Blue}, \text{Red}, \text{Blue}, \text{Red}, \text{Green}, \text{Red}, \text{Green}, \text{Blue}, \text{Blue})\}$.⁴¹ The realized colour of the ball was revealed to the subjects at the end of the respective round. Recall that depending on the type of security, the green colour, for instance, could reflect a low-stock value (in the case of separable securities) or a high-stock value (in the case of strongly separable securities). Furthermore, the trading pairs were fixed for the duration of the round, but new pairs were formed in every new round. This information was common knowledge.

At the beginning of each round, traders were given an endowment of 1,500 ECUs. Payoffs were calculated based on the MSR *at the end of each trading period*. Thus, the trader's payoff was a function of (1) the stock value (high or low), (2) the trader's own prediction, and (3) the previous trader's reported prediction.

- When the value of the stock was high (*i.e.* 100), the trader's payoff was given by the formula:

$$0.01[(100 - \text{previous trader's reported prediction})^2 - (100 - \text{trader's prediction})^2].$$

- When the value of the stock was low (*i.e.* 0), the trader's payoff was calculated by the formula:

$$0.01[(\text{previous trader's reported prediction})^2 - (\text{trader's prediction})^2].$$

The round payoff was then the summation of all the payoffs of the trading periods in the round. Crucially, the round payoff was determined at the end of the round when the stock value was

38. In general, a security that does not aggregate information in the myopic case would be difficult to aggregate information in the strategic case.

39. Note that the failure of information aggregation in the *Amb* setting with a separable security at the 0 initial price is special to the particular example we use. In general, information aggregation can fail at multiple initial prices.

40. This assumption is similar to that made in Fréchette and Yuksel (2017), Ioannou and Romero (2014), Vespa (2019), and Ioannou *et al.* (2023). It is necessary in order to simulate the infinitely many-periods assumption of the theoretical setting and to avoid having subjects implement backward induction reasoning.

41. The respective signals (i, j), where i is the signal of Trader 1 and j is the signal of Trader 2, were $\{(\text{Not Blue}, \text{Not Green}), (\text{Blue}, \text{Not Green}), (\text{Blue}, \text{Not Green}), (\text{Blue}, \text{Not Green}), (\text{Not Blue}, \text{Not Green}), (\text{Blue}, \text{Not Green}), (\text{Not Blue}, \text{Not Green}), (\text{Not Blue}, \text{Green}), (\text{Not Blue}, \text{Not Green}), (\text{Not Blue}, \text{Green}), (\text{Blue}, \text{Not Green}), (\text{Blue}, \text{Not Green})\}$.

TABLE 3
Characteristics of the experimental sessions

No. of Subj.	No. of Ses.	Market type	Security type	Acronym
Initial price is 0				
36	2	EU	Separable	EUS0
36	2	Amb	Separable	AmbS0
36	2	EU	Str. Separable	EUS0S
36	2	Amb	Str. Separable	AmbS0S
Initial price is 50				
36	2	EU	Separable	EUS50
36	2	Amb	Separable	AmbS50
36	2	EU	Str. Separable	EUS50S
36	2	Amb	Str. Separable	AmbS50S

Notes: In the first column, we provide the total number of participants in each treatment. In the second column, we provide the number of sessions per treatment. In every session, we had 18 participants. Treatments differed in the market type, the type of securities traded, and the initial price. The acronyms in the last column consist of the *market type* (*EU* for the market with EU preferences or *Amb* for the market with MEU preferences), the *security type* (*S* for separable securities or *StS* for strongly separable securities), and the *initial price* (0 or 50).

revealed to the traders. It was possible that based on the payoffs of a subject's predictions in the round that her funds would go down to zero or even negative.⁴² In that case, we would zero their round payoff. Specifically, subjects were told that "if your round payoff is a negative number, then we will zero your round payoff for that round. In the new round, you will be given once again your starting 1,500 ECUs." The final payoff of a trader was the summation of all the round payoffs of the trader in the 12 rounds played. To ensure that subjects understood the environment, before the actual game play, they had to complete a quiz with 15 questions.

The experimental sessions took place in February of 2019. For each treatment, we conducted two sessions. The 288 subjects were recruited from the database of the Université Paris 1 Panthéon–Sorbonne. We sent emails publicizing the experiment; interested individuals replied by email. We had participants from a variety of majors, such as business, computer science, economics, history, political science, engineering, biology, finance, art, physics, and mathematics. Participants were allowed to participate in *only* one session. The sessions lasted around an hour and a half. Average earnings per participant were €12.90. The experimental codes were programmed using the experimental software z-Tree (Fischbacher, 2007). Some general characteristics of the sessions are shown in Table 3. Note that each treatment is denoted by an acronym. In particular, the acronym (*market type*, *security type*, *initial price*) consists of the *market type* (*EU* for the market with EU preferences or *Amb* for the market with MEU preferences), the *security type* (*S* for separable securities or *StS* for strongly separable securities), and the *initial price* (0 or 50).

6.2. General hypotheses

Recall that we aim to investigate the impact on information aggregation of three dimensions. The first is the market type (unique priors and EU preferences or multiple priors and MEU preferences). The second relates to the type of security that is traded (separable or strongly separable). The third relates to the initial price announcement of the uninformed market maker (0 or 50).

42. In the actual experiments, no subject lost the entire endowment of the round.

To measure the degree of information aggregation in a market, we use the intrinsic value of the security as a benchmark. This is the most natural candidate to serve as a benchmark for two main reasons. First, by construction, the intrinsic value of the security is always revealed if the private information of the two traders is aggregated. Second, [Ostrovsky \(2012\)](#) showed that in any environment with EU preferences, the predictions of Bayesian Traders always converge to the intrinsic value for separable securities. The same holds true in environments with MEU preferences and strongly separable securities (Theorems 1 and 2). We therefore use the intrinsic value of the security as our baseline and measure its absolute difference (AD) in distance from the final prediction. We say that, given a state (*i.e.* the colour of the drawn ball in the experiment), information aggregation in market B is at least as good as that in market A, if the AD in market A is greater or equal to that in market B.

We now formulate our hypotheses. Hypothesis 1 extends and interprets the main result of [Ostrovsky \(2012\)](#) to an environment with ambiguity aversion assuming an initial price of 0.

Hypothesis 1. *Assuming an initial price of 0 and separable securities, information aggregation in the Amb market is at least as good as that in the EU market regardless of the colour of the drawn ball.*

Hypothesis 2 is a direct implication of Theorems 1 and 2, which show that strongly separable securities always aggregate information, in both EU and Amb markets with myopic or strategic traders. Here, it is formulated assuming an initial price of 0.

Hypothesis 2. *Assuming an initial price of 0 and strongly separable securities, information aggregation in the Amb market is at least as good as that in the EU market regardless of the colour of the drawn ball.*

We now test the degree of information aggregation of the two security types when the initial price is 50.

Hypothesis 3. *Assuming an initial price of 50 and separable securities, information aggregation in the Amb market is at least as good as that in the EU market regardless of the colour of the drawn ball.*

Hypothesis 4. *Assuming an initial price of 50 and strongly separable securities, information aggregation in the Amb market is at least as good as that in the EU market regardless of the colour of the drawn ball.*

The next pair of hypotheses investigates whether separable and strongly separable securities, respectively, are prone to manipulation by the uninformed market maker. Note that we again interpret and extend the main result of [Ostrovsky \(2012\)](#) to an environment with ambiguity aversion. Specifically, we test whether, holding the Amb market fixed, changing the initial price from 0 to 50 has any impact on the degree of information aggregation of each security type.

Hypothesis 5. *In the Amb market with separable securities and for any colour of the drawn ball, the information aggregation under an initial price of 50 is at least as good as that under an initial price of 0.*

Hypothesis 6. *In the Amb market with strongly separable securities and for any colour of the drawn ball, the information aggregation under an initial price of 50 is at least as good as that under an initial price of 0.*

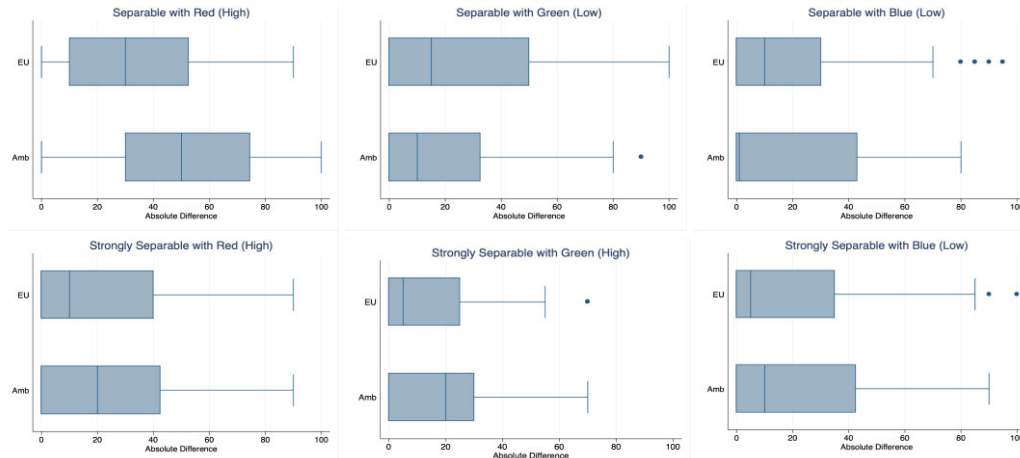


FIGURE 1

Box plots for initial price of 0

Notes: We display the box plots of the ADs across the market and security types conditional on the realized state (red, green, blue) when the initial price is 0.

6.3. Results

6.3.1. Descriptive statistics. We report next some descriptive statistics about the absolute difference (AD) in distance of the final prediction from the intrinsic value of the security. On one hand, when the stock value is low (*i.e.* in the green and blue states of the separable securities, and in the blue state of the strongly separable securities), the median AD also indicates the median last reported prediction. On the other hand, when the stock value is high (*i.e.* in the red state of the separable securities, and in the red or green states of the strongly separable securities), one needs to subtract the median AD from 100 to get the median last reported prediction.

In Figure 1, we display the box plots of the ADs across the market types when the initial price is 0, and in Figure 2, we display the box plots when the initial price is 50. It is evident from the box plots that there was a lot of variability in the reports of the subjects. This could be attributed to the nature of the game which allows for strategic behaviour, and thus results in noisier predictions.

Looking at the median ADs, typically the red state had the largest value, then the green state and, finally, the blue state. For instance, in the treatment EUS0, the median AD for the red state was 30 (*i.e.* the median last reported prediction was 70), the median AD for the green state was 15, and for the blue state it was 10. The last two values were also the median last reported predictions. There was also one treatment where the median AD of the red state was equal to that of the green state; specifically, in the treatment AmbStS0, the red and green states had a median AD of 20. In another treatment, EUS0, the green state and the blue state both had a median AD of 5. The highest median AD was 50 in treatments AmbS0 and AmbS50 for the red states. The fact that subjects consistently had trouble aggregating information with the red state should not be surprising given that it was the only state that did not explicitly reveal the colour of the drawn ball to any trader, in contrast to the green and blue states.

6.3.2. Information aggregation. To investigate the impact on information aggregation of the treated variables, we perform the following statistical analysis. In particular, we use the Mann–Whitney test, where the H_0 states that the AD in the EU market is greater or equal to

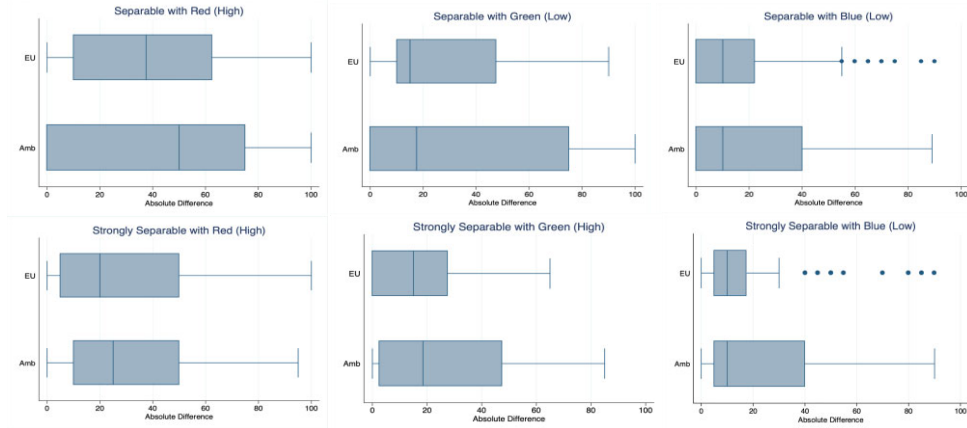


FIGURE 2
Box plots for initial price of 50

Notes: We display the box plots of the ADs across the market and security types conditional on the realized state (red, green, blue) when the initial price is 50.

the AD in the Amb market when fixing the realized state. Thus, rejecting the H_0 signifies that information aggregation is significantly worse in the Amb market relative to that in the EU market. The p -values are displayed in Table 4.

The first hypothesis dealt with the case of separable securities and initial price of 0. The first result is formalized next.

Result 1. *For an initial price of 0 and separable securities, information aggregation in the Amb market is at least as good as that in the EU market when the drawn balls are green or blue. When the drawn ball is red, information aggregation in the Amb market is significantly worse.*

Support. Contrary to our hypothesis, we find that in the red state, information aggregation in the Amb market is significantly worse (p -value is 0.001). Therefore, the H_0 can be rejected at the conventional 5% level of statistical significance in the red state.

Next, we investigate the hypothesis of strongly separable securities when the initial price is 0. Our second result sheds light to the strength of the strong separability condition.

Result 2. *For an initial price of 0 and strongly separable securities, information aggregation in the Amb market is at least as good as that in the EU market regardless of the colour of the drawn ball.*

Support. The p -values in the red, green and blue states are 0.107, 0.133 and 0.195, respectively. We thus fail to reject the H_0 .

Hypotheses 3 and 4 investigate the effect on information aggregation of separable and strongly separable securities, respectively, but this time for an initial price of 50.

Result 3. *For an initial price of 50 and separable securities, information aggregation in the Amb market is at least as good as that in the EU market regardless of the colour of the drawn ball.*

Support. We fail to reject the H_0 as the p -values in the red, green, and blue states are 0.394, 0.342, and 0.265, respectively.

TABLE 4
Mann–Whitney tests on information aggregation

Alternative hypothesis:	Initial price: 0	
	Separable	Strongly separable
	$AD_i < AD_j$	
	<i>p</i> -values	
Panel A		
<i>Red state</i>		
EU versus Amb	0.001	0.107
<i>Green state</i>		
EU versus Amb	0.479	0.133
<i>Blue state</i>		
EU versus Amb	0.447	0.195
Panel B		
	Initial price: 50	
<i>Red state</i>		
EU versus Amb	0.394	0.316
<i>Green state</i>		
EU versus Amb	0.342	0.167
<i>Blue state</i>		
EU versus Amb	0.265	0.262

Notes: We utilize the Mann–Whitney tests to determine whether the AD of the security in the EU market is greater or equal to its AD in the Amb market when fixing the realized state. In Panel A, we report the *p*-values of the comparisons in the ADs when the initial price is 0. In Panel B, we report the *p*-values of the comparisons in the ADs when the initial price is 50.

Result 4. *For an initial price of 50 and strongly separable securities, information aggregation in the Amb market is at least as good as that in the EU market regardless of the colour of the drawn ball.*

Support. The *p*-values in the red, green and blue states are 0.316, 0.167 and 0.262, respectively. Hence, we fail to reject the H_0 .

Hypotheses 5 and 6 test the degree of information aggregation in an environment with ambiguity for separable and strongly separable securities, respectively, when the initial price changes. For the analysis, we again use the Mann–Whitney test, where the H_0 states that, in the Amb market when fixing the realized state, the AD of the security when the initial price is 0 is greater or equal to the AD of the security when the initial price is 50. Therefore, rejecting in this context the H_0 means that information aggregation is significantly worse under an initial price of 50.

Result 5. *In the Amb market with separable securities, information aggregation under an initial price of 50 is at least as good as that under an initial price of 0 in the red and green states, but, in the blue state, information aggregation under an initial price of 50 is significantly worse.*

Support. None of the *p*-values is statistically significant in the red and green states (*p*-values are 0.144 and 0.194, respectively). However, in the blue state, the *p*-value is 0.068; thus, we reject the H_0 at the 10% level of statistical significance in this state.

Result 6. *In the Amb market with strongly separable securities, information aggregation under an initial price of 50 is at least as good as that under an initial price of 0 regardless of the colour of the drawn ball.*

Support. The *p*-values in the red, green and blue states are 0.111, 0.186 and 0.231, respectively. We thus fail to reject the H_0 .

7. CONCLUDING REMARKS

In 1969, Clive W. J. Granger and John M. Bates established in their seminal study that combining different forecasts was more accurate than trying to find the best one (see [Bates and Granger, 1969](#)). Those discoveries, combined with the earlier work of Friedrich Hayek, laid the foundations of prediction markets. Our primary purpose in this study has been to investigate the information aggregation properties of prediction markets with ambiguity-averse traders that have imprecise beliefs. The inability of prediction markets (financial markets in general) to aggregate information when ambiguity is present is a negative result.

We find theoretically that separable securities, which aggregate information in environments with precise beliefs and EU preferences are no longer sufficient when beliefs are imprecise. This implies that utilizing prediction markets to get a better prediction for events that are hard to quantify might backfire as traders could converge to the wrong price of the security. We introduce a new class of strongly separable securities, and show that they aggregate information in an environment with ambiguity, irrespectively of whether traders play strategically or not. Similar to [Ostrovsky \(2012\)](#), we study information aggregation only for sufficiently high t without examining what happens to prices in the middle of the game where it is possible, in an equilibrium, to diverge widely from the intrinsic value of the security.

We take our testable predictions to the laboratory where we simulate trading in prediction markets between pairs of subjects. We find that in environments with imprecise beliefs and ambiguity-averse individuals, separable securities *do not* aggregate information and are prone to manipulation by the market maker's initial price announcement. In sharp contrast, in the same environments, strongly separable securities *do* aggregate information and are resilient to such manipulation. The results for strongly separable securities are in line with the theoretical predictions.

Our emphasis and concern for dynamic prediction markets with ambiguity aversion is not only partisan, but also culminates in a profound result for asset markets in general. Proposition 3 states that there is no way to build a securitization scheme that will ensure information revelation for all information structures; that is, given that strongly separable securities are both sufficient and necessary for information aggregation, we show that there exists no security that can deliver information aggregation for all information structures. This is a negative result not only for the ability of prediction markets to aggregate information with ambiguity, but of financial markets in general.

The article leaves several open questions for future research. First, to alleviate the negative result of Proposition 3, a natural next question is whether a subset of strongly separable securities can deliver information aggregation for large classes of information structures that are of interest. Second, given that a fixed security in prediction markets cannot ensure information aggregation for all information structures, is there a different market design that can? A third direction for future research is to examine whether strongly separable securities aggregate information under ambiguity in the widely used model of [Kyle \(1985\)](#), which includes noise traders and competitive market makers. In that model, the question of information aggregation is intertwined with the question of information revelation so that even with one informed trader, it is not straightforward that her information will be eventually revealed.⁴³

43. [Ostrovsky \(2012\)](#) characterizes separable securities in that framework as well, whereas [Lambert et al. \(2018\)](#) extend it to informationally complex environments to show that, under some conditions, prices in large markets aggregate all available information. Information aggregation has also been studied in the context of other settings, such as elections (see [Barelli et al., 2020](#); [Ekmekci and Lauermaann, 2020](#)).

APPENDIX

A. Proofs for the non-strategic environment

In this section, we present the proofs for the characterization of strongly separable securities and the information aggregation in the non-strategic environment.

Proof of Lemma 1. Where convenient, we use the notation $s(y)(\cdot) \equiv s(y, X(\cdot))$. We first show that $\arg \max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}} E_p[s(y) - s(z)]$ does, in fact, exist. This is true because s is a continuous function, therefore, $\min_{p \in \mathcal{P}} E_p[s(y) - s(z)]$ is upper semicontinuous (as infimum of continuous functions) as a function of y . Since $[\underline{y}, \bar{y}]$ is compact, a maximum exists and the set $\arg \max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}} E_p[s(y) - s(z)]$ is not empty.

Next, we define V to be the convex hull of $\{s(y)\}_{y \in [\underline{y}, \bar{y}]}$. The set $\{s(y)\}_{y \in [\underline{y}, \bar{y}]}$ is compact in \mathbb{R}^I because s is continuous in y and $[\underline{y}, \bar{y}]$ is compact, hence V is compact. Consider the function $G : \mathcal{P} \times V \rightarrow \mathbb{R}$ defined by $G(p, v) = E_p[v - s(z)]$. The function is linear in p and affine in v . Moreover, it is continuous both in p and in v .

By Sion's Minimax Theorem (Berge, 1963, p. 210), there exists $p^* \in \mathcal{P}$ and $v^* \in V$ such that for all $(p, v) \in \mathcal{P} \times V$, $E_{p^*}[v - s(z)] \leq E_{p^*}[v^* - s(z)] \leq E_p[v^* - s(z)]$. Then, we get that $\min_{p \in \mathcal{P}} \max_{v \in V} E_p[v - s(z)] = \max_{v \in V} \min_{p \in \mathcal{P}} E_p[v - s(z)]$ and it is achieved at $p = p^*, v = v^*$.

For a fixed p , as s is a strictly proper scoring rule, the unique maximizer of $E_p[v - s(z)]$ over V is $s(E_p[X])$ so that $v^* = s(E_{p^*}[X])$. Hence, we may conclude that $\min_{p \in \mathcal{P}} \max_{y \in [\underline{y}, \bar{y}]} E_p[s(y) - s(z)] = \max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}} E_p[s(y) - s(z)]$ and it is achieved at $p = p^*, y = E_{p^*}[X]$. This proves the second point.

We claim that $y = E_{p^*}[X]$ is a unique element of $\arg \max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}} E_p[s(y, X(\omega)) - s(z, X(\omega))]$. To see that, let $y' \neq E_{p^*}[X]$. Then,

$$\begin{aligned} \min_{p \in \mathcal{P}} E_p[s(y', X(\omega)) - s(z, X(\omega))] &\leq E_{p^*}[s(y', X(\omega)) - s(z, X(\omega))] \\ &< E_{p^*}[s(E_{p^*}[X], X(\omega)) - s(z, X(\omega))] \\ &= \max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}} E_p[s(y, X(\omega)) - s(z, X(\omega))]. \end{aligned}$$

Hence, the maximizer is unique.

For the third claim, note that $E_p[s(z, X) - s(z, X)] = 0$ for all $p \in \mathcal{P}$, hence $\max_{y \in [\underline{y}, \bar{y}]} \min_{p \in \mathcal{P}} E_p[s(y, X) - s(z, X)] \geq 0$. As $z = E_p[X]$ for some $p \in \mathcal{P}$, we have that $p \in \arg \min_{p \in \mathcal{P}} \max_{y \in [\underline{y}, \bar{y}]} E_p[s(y, X) - s(z, X)]$ and $y^* = z$. \square

Proof of Proposition 2. Suppose that X is not strongly separable for \mathcal{P} and v . Then, from Lemma 1, we have that for each $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p) = E$ and for each $i \in I$, $E_p[X(\omega) - v \mid \Pi_i(\omega)] = 0$ for some $p \in \mathcal{P}$ ignoring, without loss of generality, states ω' for which $X(\omega') = v$. Given that $\text{Supp}(p) \subseteq E$, it cannot be that for some Trader i , state $\omega \in E$ and $\lambda \in \mathbb{R}_+$, $(X(\omega') - v)\lambda > 0$ for all $\omega' \in \Pi_i(\omega) \cap E$.

Conversely, suppose that for some $v \in \mathbb{R}$ and $E \subseteq \{\omega \in \Omega : X(\omega) \neq v\}$, for any Trader i and state $\omega \in E$, we have both $(X(\omega') - v) > 0$ and $(X(\omega'') - v) < 0$ for some $\omega', \omega'' \in \Pi_i(\omega) \cap E$. Then, for each i there exists p'' with $\text{Supp}(p'') = E$ such that $E_{p''}[X(\omega) - v \mid \Pi_i(\omega)] = 0$. To see this, let $E_1 = \{\omega' \in \Pi_i(\omega) : X(\omega') > v\}$ with k_1 elements and $E_2 = \{\omega' \in \Pi_i(\omega) : X(\omega') < v\}$ with k_2 elements. Then, $k \sum_{\omega' \in E_1} X(\omega') + (1 - k) \sum_{\omega' \in E_2} X(\omega')$ is strictly above v for big enough $k \in (0, 1)$ and strictly below v for small enough k . From the Intermediate Value Theorem, for some k , we have $E_{p'}[X(\omega) - v] = 0$, where p' assigns $\frac{k}{k_1}$ to each state $\omega' \in E_1$ and $\frac{k}{k_2}$ to each state $\omega' \in E_2$. We can then extend p' to a belief p'' with full support on E such that its conditional given $\Pi_i(\omega)$ is p' .

Collect all these beliefs p'' for each i and $\omega \in E$, letting \mathcal{P} be their convex hull. Note that \mathcal{P} is regular with respect to each Π_i . From the third result of Lemma 1, given that the previous announcement is v , every trader at each state ω will also announce v . Hence, X is not strongly separable for v and \mathcal{P} , a contradiction. \square

Proof of Lemma 2. Given a value $v \in \mathbb{R}$ and a security X , let $\Omega_{X \neq v} = \{\omega \in \Omega : X(\omega) \neq v\}$. We use induction on the number of securities. For $k = 1$, take any v and any event $E \subseteq \Omega_{X \neq v}$. There are two cases. First, there exist ω, ω' such that $\Pi_1(\omega) \cap E$ and $\Pi_1(\omega') \cap E$ are non-empty and disjoint. Second, $E \subseteq \Pi_1(\omega)$ for some ω . By construction, at each state $\omega \in E$, Trader 1 knows the value of security X_1 , hence we have that for all $\omega' \in \Pi_1(\omega) \cap E$, $(X(\omega') - v)\lambda > 0$ for

some $\lambda \in \mathbb{R}$. Using Proposition 2, X_1 is strongly separable. Note that in the first case, if v is strictly between $X_1(\Pi_1(\omega))$ and $X_1(\Pi_1(\omega'))$, there is a Trader i (in particular 1) who knows at state ω that the value of X_1 is strictly above v , and in some other state ω' that the value is strictly below v . We will use this fact as our induction hypothesis.

Suppose that X_k is strongly separable. We show that X_{k+1} is also strongly separable. Take any v and $E \subseteq \Omega_{X_k \neq v}$. There are two cases. First, there exist ω, ω' such that $\bigcap_{i=1, \dots, k} \Pi_i(\omega) \cap E$ and $\bigcap_{i=1, \dots, k} \Pi_i(\omega') \cap E$ are non-empty and disjoint. From the induction hypothesis, if v is strictly between $X_k(\bigcap_{i=1, \dots, k} \Pi_i(\omega))$ and $X_k(\bigcap_{i=1, \dots, k} \Pi_i(\omega'))$, there is a Trader $i = 1, \dots, k$ who knows at state ω that the value of X_k is strictly above v and a possibly different Trader $j = 1, \dots, k$ who knows in some other state ω' that the value is strictly below v . For any value $X_k v' v'' > v \omega' X_k v'' X_k(\bigcap_{i=1, \dots, k} \Pi_i(\omega)) X_k(\bigcap_{i=1, \dots, k} \Pi_i(\omega')) X_k$

By construction, when going from security X_k to security X_{k+1} the ordering of states is preserved. That is, if $X_k(\omega) < X_k(\omega')$, then $X_{k+1}(\omega) < X_{k+1}(\omega')$. The reason is that, for each $\omega \in \Omega$, X_k assigns the same value to all states in $\bigcap_{i=1, \dots, k} \Pi_i(\omega)$, whereas X_{k+1} partitions $\bigcap_{i=1, \dots, k} \Pi_i(\omega)$ using the elements of $k+1$'s partition and assigns different values to each, respecting the order for states ω, ω' such that $X_k(\omega) \neq X_k(\omega')$. This implies that for any v , there is Trader i and state $\omega \in E$ such that i knows that the value of X_{k+1} is above v or a possibly different trader who knows that it is below v .

The second case is that $E \subseteq \bigcap_{i=1, \dots, k} \Pi_i(\omega)$ for some ω . By construction, X_{k+1} partitions $\bigcap_{i=1, \dots, k} \Pi_i(\omega)$ using the elements of $k+1$'s partition and assigns different values to each. Moreover, given E , Trader $k+1$ knows the value of X_{k+1} at every state in E . If $v \in [\min_{\omega \in E} X_{k+1}(\omega), \max_{\omega \in E} X_{k+1}(\omega)]$, then Trader $k+1$ knows whether the value of X_{k+1} is above or below v at some state in E . The same is true if $v \notin [\min_{\omega \in E} X_{k+1}(\omega), \max_{\omega \in E} X_{k+1}(\omega)]$.

We have shown that in all cases, for any v and any $E \subseteq \Omega_{X_k \neq v}$, there is a Trader i who knows at some state $\omega \in E$ whether the value of X_{k+1} is strictly below or strictly above v . Hence, we have that for all $\omega' \in \Pi_i(\omega) \cap E$, $(X(\omega') - v)\lambda > 0$ for some $\lambda \in \mathbb{R}$. Using Proposition 2, X_{k+1} is strongly separable. \square

Proof of Corollary 1. Take any v and non-empty event $E \subseteq \{\omega \in \Omega : X(\omega) \neq v\}$. We then have that there exist Trader i and state $\omega \in E$ such that $\Pi_i(\omega) \cap E \subseteq X^{-1}(k)$ for some k . By construction of E , $k \neq v$. This implies that for some $\lambda \in \mathbb{R}$, $(X(\omega') - v)\lambda > 0$ for all $\omega' \in \Pi_i(\omega) \cap E$. Using Proposition 2, the security is strongly separable. \square

Proof of Proposition 3. Take any (non-constant) security X and consider the partition \mathcal{X} generated by its values: for each $\omega \in \Omega$, $\omega' \in \mathcal{X}(\omega)$ if $X(\omega) = X(\omega')$. The partition \mathcal{X} has at least two partition cells. Let A be the partition cell generated by the lowest value of X , call it v_A , and B the partition cell generated by the highest value of X . Since Ω has at least three states, we assume, without loss of generality, that the complement of A , denoted A^c , has at least two states (if not, then the complement of B must have at least two states and the same argument applies).

Consider an information structure with two traders. Trader 1's partition cell at state $a \in A$ also includes state $b \in A^c$ so that $\Pi_1(a) = \{a, b\}$. For any other state $\omega \neq a, b$, $\Pi_1(\omega) = \{\omega\}$. Trader 2's partition cell at $a \in A$ also contains state $c \in A^c$ so that $\Pi_2(a) = \{a, c\}$ with $b \neq c$. For any other state $\omega \neq a, c$, $\Pi_2(\omega) = \{\omega\}$. Hence, the join of the two traders' partitions consists of singleton sets.

Let v be strictly higher than v_A and strictly lower than all other values of X . If we let event $E = \{a, b, c\} \subseteq \{\omega \in \Omega : X(\omega) \neq v\} = \Omega$, then $\Pi_1(a) \cap E = \{a, b\}$ and $\Pi_2(a) \cap E = \{a, c\}$. For v, E and state $\omega = a$, we have that for $i = 1, 2$ there is no $\lambda \in \mathbb{R}$ such that for all $\omega' \in \Pi_i(\omega) \cap E$, $(X(\omega') - v)\lambda > 0$. The reason is that both traders consider possible a state where X has a value strictly higher than v and a state where X has a value strictly lower than v . Applying Proposition 2, we have that X is not strongly separable. \square

Proof of Theorem 1. (\Leftarrow) Suppose X is strongly separable. By construction, $\mathcal{F}^0(\omega) \supseteq \mathcal{F}^1(\omega) \supseteq \dots \supseteq \mathcal{F}^k(\omega)$. As Ω is finite, there exists t_k such that $\mathcal{F}^{k'}(\omega) = \mathcal{F}^k(\omega)$ for every $t_{k'} \geq t_k$. We denote this set by $\mathcal{F}(\omega) \equiv \mathcal{F}^k(\omega)$.

Define the function $g(E_p[X]) = \min_{q: E_q[X] = E_p[X]} E_q[s(E_q[X], X) - s(z, X)]$. We first show that g is strictly convex in $\{E_p[X] : p \in \Delta(\Omega)\}$. Let $g(k) = E_p[s(E_p[X], X) - s(z, X)]$ and $g(m) = E_q[s(E_q[X], X) - s(z, X)]$. We have

$$\begin{aligned} g(ak + (1-a)m) &= g(aE_p[X] + (1-a)E_q[X]) = g(E_{ap+(1-a)q}[X]) \\ &= \min_{r: E_r[X] = E_{ap+(1-a)q}[X]} E_r[s(E_r[X], X) - s(z, X)] \\ &\leq E_{ap+(1-a)q}[s(E_{ap+(1-a)q}[X], X) - s(z, X)] \end{aligned}$$

$$\begin{aligned}
&= aE_p[s(E_{ap+(1-a)q}[X], X) - s(z, X)] + (1-a)E_q[s(E_{ap+(1-a)q}[X], X) - s(z, X)] \\
&< aE_p[s(E_p[X], X) - s(z, X)] + (1-a)E_q[s(E_q[X], X) - s(z, X)] \\
&= ag(k) + (1-a)g(m).
\end{aligned}$$

Note that $g(E_p[X]) \geq 0$ for all $E_p[X]$ and its unique minimizer is at z , where $g(z) = 0$.⁴⁴ We then have that g is decreasing at $[a, z]$ and increasing at $[z, b]$, where $a = \min\{E_p[X] : p \in \Delta(\Omega)\}$ and $b = \max\{E_p[X] : p \in \Delta(\Omega)\}$. From Lemma 1, the myopic announcement of Trader i with beliefs $\mathcal{P}_{\mathcal{F}(\omega) \cap \Pi_i(\omega')}$ and previous announcement z is given by $d_{\mathcal{P}}(\mathcal{F}(\omega) \cap \Pi_i(\omega'), z) = \arg \min_{x \in \{E_p[X] : p \in \mathcal{P}_{\mathcal{F}(\omega) \cap \Pi_i(\omega')}\}} g(x)$.

Define $A_{\omega'}^i = \{E_p[X] : p \in \mathcal{P}_{\mathcal{F}(\omega) \cap \Pi_i(\omega')}\}$ for every $i = 1, \dots, n$ and $\omega' \in \mathcal{F}(\omega)$. If z (the unique minimizer of $g(E_p[X])$ for all $p \in \Delta(\Omega)$) is less than or equal to the minimum value of $A_{\omega'}^i$, then that minimum value is the minimizing value of $g(x)$ given $\mathcal{P}_{\mathcal{F}(\omega) \cap \Pi_i(\omega')}$ and therefore the myopic announcement. Similarly, if z is greater or equal to the maximum value of $A_{\omega'}^i$, then that maximum value is the minimizing value of $g(x)$ given $\mathcal{P}_{\mathcal{F}(\omega) \cap \Pi_i(\omega')}$ and therefore the myopic announcement. If z is inside $A_{\omega'}^i$, then the myopic announcement is z . This is due to the strict convexity of g and the fact that z is the global minimum.

Note that for all $t \geq t_k$, $A_{\omega'}^i$ is constant for all i and $\omega' \in \mathcal{F}(\omega)$ because information is no longer updated. We now show that traders agree on the myopic announcement. There are three cases.

Case 1: For some i , $A^i = \bigcap_{\omega' \in \mathcal{F}^k(\omega)} A_{\omega'}^i = \emptyset$.

This implies that $A_{\omega'}^i \cap A_{\omega''}^i = \emptyset$ for some states $\omega', \omega'' \in \mathcal{F}(\omega)$. From Lemma 1, Trader i will either make an announcement in $A_{\omega'}^i$ (if she considers ω' to be true) or $A_{\omega''}^i$ (if she considers ω'' to be true). As $A_{\omega'}^i \cap A_{\omega''}^i = \emptyset$, either ω' or ω'' will be revealed not to be true, which means that not all information has been aggregated yet, a contradiction.

Case 2: $A^i \neq \emptyset$ for all $i \in I$ and $\bigcap_{j \in I} A^j \neq \emptyset$.

There are two subcases. First, there is v such that for all $\omega' \in \mathcal{F}(\omega)$, $X(\omega') = v$. From Lemma 1, the myopic announcement for every Trader is v , hence Traders agree on the announcement. Second, there are two states $\omega', \omega'' \in \mathcal{F}(\omega)$ such that $X(\omega') \neq X(\omega'')$. This means that the first part of Definition 3 of a not strongly separable security is satisfied for $\mathcal{P}_{\mathcal{F}(\omega)}$. Let $z \in \bigcap_{j \in I} A^j \neq \emptyset$. From Lemma 1, part (iii), if the previous announcement is z and $z \in A^i$, then Trader

i will also announce z . We then have that $d_{\mathcal{P}_{\mathcal{F}(\omega)}}(\Pi_i(\omega'), z) = z$ for all $i = 1, \dots, n$ and $\omega' \in \bigcup_{p \in \mathcal{P}_{\mathcal{F}(\omega)}} \text{Supp}(p)$. But this implies that X is not strongly separable, a contradiction.

Case 3: $A^i \neq \emptyset$ for all $i \in I$ but $\bigcap_{j \in I} A^j = \emptyset$.

We first make two observations. From the second property of Lemma 1, each Trader j makes an announcement in $A_{\omega'}^j$ for some $\omega' \in \mathcal{F}(\omega)$. As all information has been aggregated after t_k , any such myopic announcement must be in A^j . Second, the third property of Lemma 1 shows that if the previous announcement of Trader $i-1$ is in A^i , then Trader i will repeat the same announcement. Combining these two observations, we have that if the announcement changes, from Trader $i-1$ to Trader i , then it must be that Trader i is announcing either the right-hand side extreme point of A^i (i.e. the maximum) or the left-hand side extreme point of A^i (i.e. the minimum). In that case, if she announces the maximum (minimum) of A^i , then this is equal to the maximum (minimum) of $A_{\omega'}^i$ for all $\omega' \in \mathcal{F}(\omega)$, otherwise there would be further information aggregation.

Define $i_0 = \min\{i : \bigcap_{j \in \{1, \dots, i\}} A^j = \emptyset\}$. Given that $\bigcap_{j \in I} A^j = \emptyset$, i_0 exists. Moreover, A^{i_0} has an empty intersection with $\bigcap_{j \in \{1, \dots, i_0-1\}} A^j$ and, without loss of generality, suppose that A^{i_0} is on the left-hand side of $\bigcap_{j \in \{1, \dots, i_0-1\}} A^j$. Given that $\bigcap_{j \in \{1, \dots, i_0-1\}} A^j$ is an interval, we can conclude that there are A^{i_1} and A^{i_2} such that one of them defines the left-hand side extreme point of the interval and the other one the right-hand side extreme point.

44. We can observe that there exists $p \in \Delta(\Omega)$ such that $E_p[X] = z$. In addition, the set $\{E_p[X] : p \in \mathcal{P}\}$ is an interval as a convex and closed set of the real numbers.

From the second property of Lemma 1, each Trader j makes an announcement in A^j . Hence, for any value y_{k-1} , trader $i_3 = \max\{i_1, i_2\}$ makes a prediction belonging in the set $\bigcap_{j \in \{1, \dots, i_0-1\}} A^j$. For the same reason, any subsequent announcement up to $i_0 - 1$ also belongs to $\bigcap_{j \in \{1, \dots, i_0-1\}} A^j$. From the convexity of g and the fact that A^{i_0} is to the left of that interval, the prediction of i_0 is always the right-hand side extreme point of A^{i_0} , which we denote by v_0 . Moreover, it cannot be that the right-hand extreme point of A^{i_0} is different from the right-hand extreme point of $A_{\omega'}^{i_0}$ for some $\omega' \in \mathcal{F}(\omega)$, otherwise it would be revealed that ω' is not true.

As i_3 makes some announcement $v_{i_3} > v_0$ and i_3 will announce in the next round, it must be that some Trader j after i_0 will change the announcement to some $v' > v_0$. Then, it must be that v' is the left-hand extreme point of A^j . Using the same argument as in the previous paragraph, it cannot be that the left-hand extreme point of A^j is different from the left-hand extreme point of $A_{\omega'}^j$, for some $\omega' \in \mathcal{F}(\omega)$, otherwise it would be revealed that ω' is not true.

We then have that for all $p \in \mathcal{P}_{\mathcal{F}(\omega)}$ for all $\omega' \in \mathcal{F}(\omega)$, $E_p[X | \Pi_{i_0}(\omega')] \leq v_0 < v' \leq E_p[X | \Pi_j(\omega')]$. Integrating over all $\omega' \in \mathcal{F}(\omega)$ and since $\mathcal{F}(\omega)$ is common knowledge at ω , we have that $E_p[X] \leq v_0 < E_p[X]$ for all $p \in \mathcal{P}_{\mathcal{F}(\omega)}$, a contradiction.

(\Rightarrow) Suppose that for any regular Γ^M , information aggregates so that $y_k(\omega) = d\mathcal{P}(\Pi_{a_k}(\omega) \cap \mathcal{F}^{k-1}(\omega), y_{k-1}) \rightarrow X(\omega)$ for every $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$. We show that, for any regular \mathcal{P} and $v \in \mathbb{R}$, if (ii) in Definition 3 is satisfied, then

(i) is violated.

Suppose there exist regular \mathcal{P} and $v \in \mathbb{R}$ such that $d\mathcal{P}(\Pi_i(\omega), v) = v$ for all $i = 1, \dots, n$ and $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$.

Consider regular $\Gamma^M(\Omega, I, \Pi, X, \mathcal{P}, y_0, \bar{y}, s)$ with initial announcement $y_0 = v$. Then, the predictions $y_{t_k}(\omega)$, $k = 0, 1, \dots$, are equal to v for all $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$. If we have $X(\omega) \neq v$ for some $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$, then at ω all traders agree on v , which is the wrong value of the security. This implies that there is no information aggregation, a contradiction. Hence, condition (i) in Definition 3 is violated and X is strongly separable. \square

B. Proofs for the strategic environment

Before proving Theorem 2, we state the following auxiliary result, which shows that a trader's continuation value is always greater than her one-period payoff.

Proposition 4. *In a Revision-Proof equilibrium, the continuation value for Trader i who plays at t_k is at least as much as her utility from the one-period payoff from playing the myopic best response.*

Proof. We construct a deviation strategy that guarantees a continuation value at least as much as that of the one-period payoff from playing the myopic strategy. We will show that for each t_k , the continuation payoff of Trader i who makes the announcement is weakly more than χ_k , her one-period payoff from playing the myopic strategy at t_k .

We define a deviation strategy $\sigma = (\sigma_i, \sigma_{-i}^*)$, where all traders $j \neq i$ follow the equilibrium strategy σ^* and σ_i is identical to σ_i^* up to time t_{k-1} . At t_k , σ_i specifies that Trader i plays the myopic best response. Given that i deviates and all other traders stick to the equilibrium strategy σ^* , let H^1, \dots, H^m be the possible paths of announcements by all other traders $j \neq i$ from t_k to t_{k+n-1} , together with the common history of announcements up to t_{k-1} . They are finitely many because we consider mixing over finite actions. At t_{k+n} , σ_i specifies that:

- (a) If $V(H^m, \phi, \sigma, \mathcal{P}) \geq 0$ by playing what σ_i^* prescribes at H^m , then σ_i coincides with σ^* in every succeeding information set,
- (b) If $V(H^m, \phi, \sigma, \mathcal{P}) < 0$, then σ_i repeats the previous trader's prediction.

If (a) occurs, then σ_i coincides with σ^* in every succeeding information set, so Trader i follows the recommendation of σ_i . If (b) occurs, then Trader i repeats the previous announcement and in every succeeding information set, σ_i is determined using the two cases (a) and (b). For every other information set not specified by the above procedure, σ_i is identical to σ_i^* .

We now show that at (b), Trader i will follow the recommendation to repeat the previous announcement and get a period payoff of zero. This is true if her continuation value, excluding her current period payoff, is weakly positive, as long as all future selves follow σ . We now show that this is true at all $t \geq t_k$. We show this for $t = t_k$, without loss of generality, and note that, from i 's perspective, there are two types of subsequent paths, given that everyone follows σ . The first type is a path that specifies some zero payoffs initially, and at some time $t > t_{k+n}$ the continuation value of i 's future self is weakly positive by playing σ^* onwards. The second type is a path where the future selves just repeat the previous announcement because from σ^* they would get a negative continuation value, hence the payoffs along this

path are zero always. This means that all paths have a weakly positive continuation value at some time $t > t_k$, and the previous payoffs between t_k and t are zero. Hence, it is without loss of generality to assume that the future selves at period t_{k+n} and at each path, compute weakly positive continuation value. However, because of Dynamic Inconsistency the continuation value at some path at t_{k+n} may be evaluated at a different prior than the one that Trader i uses at t_k to evaluate her own continuation value. The collection of all paths generates a partition Π of state space Φ and σ generates a sequence of acts f_m , for each $t > t_k$. We therefore have, for each $E \in \Pi$, and from the perspective of the future selves in time t_{k+n} , that

$$0 \leq \min_{p \in \mathcal{P}} E_{pE} \sum_{m=0}^{\infty} \beta^{nm} u(f_{k+n+nm}) = E_{qE} \sum_{m=0}^{\infty} \beta^{nm} u(f_{k+n+nm}).$$

At every partition cell E , the future self at t_{k+n} chooses a potentially different belief q_E . Let p be the belief that Trader i uses at t_k to compute her continuation value. We then have that

$$0 \leq E_{qE} \sum_{m=0}^{\infty} \beta^{nm} u(f_{k+n+nm}) \leq E_{pE} \sum_{m=0}^{\infty} \beta^{nm} u(f_{k+n+nm}).$$

By multiplying with β and $p(E)$, and adding over all $E \in \Pi$, we have

$$0 \leq \beta \sum_{E \in \Pi} p(E) E_{pE} \sum_{m=0}^{\infty} \beta^{nm} u(f_{k+n+nm}) \leq E_p \sum_{m=0}^{\infty} \beta^{nm+1} u(f_{k+n+nm}).$$

This shows that the continuation value at any $t \geq t_k$ is weakly positive if Trader i repeats the previous announcement at t and gets a period payoff of zero. Therefore, she will always follow the recommendation at (b), if by sticking to σ^* her continuation value is strictly negative.

At t_k , Trader i plays her myopic best response $E_q[X]$ and gets a period payoff of $\chi_k(q)$, which is weakly positive, where q solves $\min_{q \in \mathcal{P}} \chi_k(q)$. Her continuation value is evaluated at some p , and therefore, we have $\chi_k(q) \leq \chi_k(p)$. Because $E_p \sum_{m=0}^{\infty} \beta^{nm+1} u(f_{k+n+nm}) \geq 0$, we have $\chi_k(p) + E_p \sum_{m=0}^{\infty} \beta^{nm} u(f_{k+n+nm}) \geq \chi_k(q)$. Hence, her continuation value is always weakly greater than her period payoff by playing the myopic best response. \square

Proof of Theorem 2. For part (i), the proof closely follows that of [Ostrovsky \(2012\)](#) and proceeds in four steps. The main innovations are in Step 1, where the arguments for establishing the lower bound of the instant opportunity are very different, and in Step 4, where we need to account for the multiplicity of beliefs.

Step 1: We show that if the security is strongly separable and its value is not constant for each state in the support of the set of beliefs, at least one trader can achieve a strictly positive payoff at some state and a weakly positive payoff at all other states, whatever the previous announcement.

Let \mathcal{P}_k be the beliefs over Ω of an outside observer who hears the announcements up to t_{k-1} and updates the initial set of beliefs \mathcal{P} given the equilibrium strategies but has no private information about Ω . Let $\text{Supp}(\mathcal{P}_k)$ be the union of the supports of all $p \in \mathcal{P}_k$. For each $\omega \in \text{Supp}(\mathcal{P}_k)$ and $i \in I$, define $A_{\omega}^{ik} = \{E_p[X \mid \Pi_i(\omega)] : p \in \mathcal{P}_k\}$ to be the set of all myopic best responses of Trader i and let $\min A_{\omega}^{ik}$ ($\max A_{\omega}^{ik}$) be the minimum (maximum) value. We first show that, in any equilibrium, the announcement of Trader i gets arbitrarily close to the announcement of Trader $i-1$ and to A_{ω}^{ik} , for all $\omega \in \text{Supp}(\mathcal{P}_k)$, as $t_k \rightarrow t_{\infty}$. Note that $A^{ik} = \bigcap_{\omega \in \text{Supp}(\mathcal{P}_k)} A_{\omega}^{ik}$ cannot be empty, otherwise the outside

observer would understand that some state ω is not true, because the announcements do not get arbitrarily close to A_{ω}^{ik} as $t_k \rightarrow t_{\infty}$. Hence, the announcements get arbitrarily close to A^{ik} as well.

Lemma 3. For any $\epsilon > 0$ and Trader i , there is period t' such that for all $t_k > t'$ where i makes an announcement, $|y_k - y_{k-1}| < \epsilon$ and $y_k \in [\min A_{\omega}^{ik} - \epsilon, \max A_{\omega}^{ik} + \epsilon]$, for all $\omega \in \text{Supp}(\mathcal{P}_k)$.

Proof. We first show that, for any $\epsilon > 0$, if Trader $i-1$'s announcement z is outside $[\min A_{\omega}^{ik} - \epsilon, \max A_{\omega}^{ik} + \epsilon]$, for some state ω and time t_k , then i 's expected payoff from playing her myopic best response is greater than some $\chi_k > 0$. For all $z < \min A_{\omega}^{ik} - \epsilon$, i 's myopic best response at ω is $\min A_{\omega}^{ik} = E_{p_1}[X]$ for some $p_1 \in \mathcal{P}_{\Pi_i(\omega)}^k$, where $\mathcal{P}_{\Pi_i(\omega)}^k$ are i 's beliefs at time t_k and state ω , given the equilibrium play up to time t_{k-1} . Her period utility is $E_{p_1}(s(\min A_{\omega}^{ik}, X) - s(z, X)) > E_{p_1}(s(\min A_{\omega}^{ik}, X) - s(\min A_{\omega}^{ik} - \epsilon, X)) = l_{1k} > 0$. The inequality follows because $\min A_{\omega}^{ik} - \epsilon$ is closer to $E_{p_1}[X]$ than z so that the score s increases because any scoring rule is order sensitive. Similarly, for all $z > \max A_{\omega}^{ik} + \epsilon$, i 's myopic best response is $\max A_{\omega}^{ik} = E_{p_2}[X]$ for some $p_2 \in \mathcal{P}_{\Pi_i(\omega)}^k$. Her period utility is $E_{p_2}(s(\max A_{\omega}^{ik}, X) - s(z, X)) > E_{p_2}(s(\max A_{\omega}^{ik}, X) - s(\max A_{\omega}^{ik} + \epsilon, X)) = l_{2k} > 0$. Hence,

for all $z \notin [\min A_{\omega}^{ik} - \epsilon, \max A_{\omega}^{ik} + \epsilon]$, i 's period utility at ω from playing the myopic best response is higher than $\chi_k = \min\{l_{1k}, l_{2k}\} > 0$.

We next show that χ_k cannot converge to 0 as $t_k \rightarrow t_{\infty}$. Consider the set of beliefs $\{p_k\}$ for which the myopic best response is calculated for each t_k . Since the set of all beliefs is compact, there is a converging sequence $\{p_k\}$ of beliefs. If $\lim_{p_k \rightarrow p} E_{p_k}(s(E_{p_k}[X] + \epsilon, X) - s(E_{p_k}[X], X)) = 0$, the continuity of the scoring rule implies that $E_p(s(E_p[X] + \epsilon, X)) = E_p(s(E_p[X], X))$ so that both announcements $E_p[X] + \epsilon$ and $E_p[X]$ are optimal given p , contradicting that s is a strictly proper scoring rule.

From Proposition 4, Trader i 's continuation payoff in equilibrium must be weakly higher than her one-period payoff χ_k . This implies that if Trader $i - 1$ makes announcements outside of $[\min A_{\omega}^{ik} - \epsilon, \max A_{\omega}^{ik} + \epsilon]$ for infinitely many t_k , then i 's expected continuation payoff (which is greater than χ_k) does not converge to zero. We now show that this is impossible.

Suppose not. Then, the expected continuation payoff for i is bounded below by a positive number. For all other traders it is weakly positive, again using Proposition 4 and because their one-period payoff is always weakly positive. Given that the continuation payoff is minimized over all beliefs in \mathcal{P}_k , we can pick any $p \in \mathcal{P}_k$ and define Ψ_k to be the sum of all traders' expected continuation payoffs (given that p) at t_k , divided by β^k ,

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

The \bar{s}_k is the expected score of prediction y_k , where the expectation is over all ϕ given some $p \in \mathcal{P}_k$ and the moves of players according to the mixed equilibrium.

For any K , we have

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

where $M = \max_{y \in [y, \bar{y}], \omega \in \Omega} |s(y, X(\omega))|$. But this contradicts the fact that i 's expected continuation payoff is bounded below by a positive number. We then have that, in equilibrium, Trader $i - 1$ makes announcements that are arbitrarily close to A_{ω}^{ik} for each $\omega \in \text{Supp}(\mathcal{P}_k)$, hence arbitrarily close to A^{ik} . Given that $\sum_{k=1}^K \Psi_k$ is bounded above by a positive

number for any K , and each Ψ_k is weakly positive, we have that $\lim_{K \rightarrow \infty} \sum_{k=1}^K \Psi_k = \chi_0$ for some finite χ_0 .

We finally show that, given that $i - 1$ announces arbitrarily close to A^{ik} , the announcement of i gets arbitrarily close to the announcement of $i - 1$ in equilibrium, and therefore the announcements of i get arbitrarily close to A^{ik} . Suppose not, so that $|y_k - y_{k-1}| > \epsilon$ for a fixed ϵ and for infinitely many t_k , where i makes an announcement. Suppose that in every t_k , where i makes an announcement, we evaluate i 's period payoff at t_k using $p_k \in \mathcal{P}_{\Pi_i(\omega)}^k$, such that $E_{p_k}[X] = y_{k-1}$ if $y_{k-1} \in A_{\omega}^{ik}$, $E_{p_k}[X] = \min A_{\omega}^{ik}$ if $y_{k-1} < \min A_{\omega}^{ik}$ (but arbitrarily close to it) or $E_{p_k}[X] = \max A_{\omega}^{ik}$ if $y_{k-1} > \max A_{\omega}^{ik}$ (but arbitrarily close to it). In all cases and since $i - 1$'s announcement is arbitrarily close to A_{ω}^{ik} , we have that i 's period payoff $E_{p_k}(s(y_k, X) - s(y_{k-1}, X))$ is strictly negative. As scoring rules are order sensitive, we have that the period payoff will also be strictly negative if i 's announcement is exactly ϵ away from the announcement of $i - 1$. By collecting these p_k for all such t_k , we have that $E_{p_k}(s(E_{p_k}[X] + \epsilon, X) - s(E_{p_k}[X], X)) < 0$, where $s(y_{k-1}, X)$ is arbitrarily close to $s(E_{p_k}[X], X)$, by continuity.⁴⁵ Since the set of all beliefs is compact, there is a converging sequence $\{p_k\}$ of beliefs. If $\lim_{p_k \rightarrow p} E_{p_k}(s(E_{p_k}[X] + \epsilon, X) - s(E_{p_k}[X], X)) = 0$, the continuity of the scoring rule implies that $E_p(s(E_p[X] + \epsilon, X)) = E_p(s(E_p[X], X))$ so that both announcements $E_p[X] + \epsilon$ and $E_p[X]$ are optimal given p ,

45. We assume, without loss of generality, that i 's announcement is always ϵ -higher than the previous announcement. Similar arguments can be employed if it was always ϵ -lower or it was alternating as we can always get a subsequence of p_k for which i 's announcement is always higher (or lower).

contradicting that s is a strictly proper scoring rule. If $\lim_{p_k \rightarrow p} E_{p_k}(s(E_{p_k}[X] + \epsilon, X) - s(E_{p_k}[X], X)) < 0$, i 's period payoff given some beliefs p_k is bounded above by a strictly negative number. But, this is the first element of some Ψ_k . We have already shown that $\sum_{k=1}^K \Psi_k$ is bounded above by a positive number for each K and each Ψ_k is weakly positive because a trader can always repeat the previous announcements, as shown in Proposition 4. Therefore, we have that $\lim_{k \rightarrow \infty} \Psi_k = 0$, which contradicts that the first term can be bounded above by a negative number. Since this is true for all states in $\text{Supp}(\mathcal{P}_k)$, the above statements are also true for A^{ik} and the result follows. That is, given that $i - 1$ announces arbitrarily close to A^{ik} , the announcement of i gets arbitrarily close to the announcement of $i - 1$ in equilibrium, and therefore the announcements of i get arbitrarily close to A^{ik} . \square

Given an equilibrium, the updating of beliefs \mathcal{P}_t may never stop for sufficiently high t_k , as traders play their mixed strategies and do prior-by-prior updating. Let \mathcal{P} be a set of limit beliefs of this sequence $\{\mathcal{P}_k\}$ with some probability. Let \mathcal{D} be the collection of these sets of limit beliefs that describe some uncertainty about the value of the security. That is, for each $\mathcal{P} \in \mathcal{D}$, there exist $\omega, \omega' \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$ such that $X(\omega) \neq X(\omega')$.

From Lemma 1, we know that given beliefs $\mathcal{P} \in \mathcal{D}$ and at any state $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$, each Trader j can achieve a weakly positive payoff by making the myopic announcement $\min_{p \in \mathcal{P}_{\Pi_j(\omega)}} E_p[s(E_p[X], X) - s(z, X)]$, where z is the previous announcement.

Generalizing the notion of [Ostrovsky \(2012\)](#), we define the instant opportunity of Trader i , given beliefs $\mathcal{P} \in \mathcal{D}$ and previous announcement z , to be

$$\min_{q \in \mathcal{P}} \sum_{\omega \in \Omega} q(\omega) \left[\min_{p \in \mathcal{P}_{\Pi_i(\omega)}} E_p[s(E_p[X], z), X) - s(z, X)] \right].$$

Note that at each partition cell $\Pi_i(\omega)$, Trader i chooses a possibly different $p \in \mathcal{P}_{\Pi_i(\omega)}$ that minimizes her expected utility. The instant opportunity is the ex ante (minimal over \mathcal{P}) expected utility aggregated over all partition cells.

The following lemma shows that if the security X is strongly separable and beliefs $\mathcal{P} \in \mathcal{D}$ describe some uncertainty about X , then the instant opportunity of some Trader i is strictly positive irrespective of what the previous announcement is.

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

Proof. Note that the expression for the instant opportunity inside the brackets,

$$\min_{p \in \mathcal{P}_{\Pi_i(\omega)}} E_p[s(E_p[X], z), X) - s(z, X)], \quad (1)$$

is i 's expected payoff given $\Pi_i(\omega)$ when making the myopic announcement and the previous announcement is z . From Lemma 1, this is weakly positive for all $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$. Moreover, given that \mathcal{P} is regular, each $p \in \mathcal{P}$ assigns positive probability to each $\Pi_i(\omega)$ where $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p) = E$. Therefore, we only need to show that there exists some trader $i \in I$, such that for any z , there is some $\Pi_i(\omega)$ for which the expression in (1) is above a strictly positive lower bound. Note that the lower bound must be the same for all z .

For each ω and $i \in I$, define $A_\omega^i = \{E_p[X | \Pi_i(\omega)] : p \in \mathcal{P}\}$ and let $\min A_\omega^i$ ($\max A_\omega^i$) be the minimum (maximum) value. Let $A^i = \bigcap_{\omega \in \mathcal{F}} A_\omega^i$. There are three cases.

Case 1: For some i , $A^i = \emptyset$.

Let i be such that $A^i = \emptyset$. Given that each A_ω^i is a convex set, there exist states $a, b \in E$ with $A_a^i = [c, d]$, $A_b^i = [c', d']$ such that $c' > d$. Let $k = (c' - d)/2$ and z be the previous announcement. If $z > k$, then $\min_{p \in \mathcal{P}_{\Pi_i(a)}} \max_{y \in [\underline{y}, \bar{y}]} E_p[s(y, X) - s(z, X)] = \min_{p \in \mathcal{P}_{\Pi_i(a)}} E_p[s(E_p[X], X) - s(z, X)] = E_{p^*}[s(E_{p^*}[X], X) -$

$s(z, X) \geq E_{p^*}[s(E_{p^*}[X], X) - s(k, X)] \geq \min_{p \in \mathcal{P}_{\Pi_i(a)}} E_p[s(E_p[X], X) - s(k, X)] \equiv \chi_1 > 0$.⁴⁶ Similarly, if $z \leq k$, then $\min_{p \in \mathcal{P}_{\Pi_i(b)}} \max_{y \in [\underline{y}, \bar{y}]} E_p[s(y, X) - s(z, X)] \geq \min_{p \in \mathcal{P}_{\Pi_i(b)}} \max_{y \in [\underline{y}, \bar{y}]} E_p[s(y, X) - s(k, X)] \equiv \chi_2 > 0$. The lower bound $\chi > 0$ is just the minimum of χ_1 and χ_2 . Moreover, it is independent of the previous announcement z .

Case 2: $A^i \neq \emptyset$ for all $i \in I$ and $\bigcap_{j \in I} A^j \neq \emptyset$.

This is the same as Case 2 in the proof of Theorem 1. There are two subcases. First, in all states that are considered possible, security X pays the same. This is impossible because we have assumed that there is uncertainty about X given \mathcal{P} . Second, there is uncertainty about X . As we showed in Case 2 in the proof of Theorem 1, this implies that X is not strongly separable, a contradiction.

Case 3: $A^i \neq \emptyset$ for all $i \in I$ but $\bigcap_{j \in I} A^j = \emptyset$.

We will show that this case is impossible. Lemma 3 shows that in any set of beliefs \mathcal{P}_t that can arise in equilibrium after a sufficiently large t_k , i 's announcements get arbitrarily close to A^{it} . Moreover, Trader i 's announcements get arbitrarily close to the announcements of $i - 1$, which get arbitrarily close to A^{i-1t} . At the limit set of beliefs \mathcal{P} , we have that $A^{i-1} \cap A^i \neq \emptyset$ for each $i \in I$. Continuing inductively over all traders, we have that $\bigcap_{j \in I} A^j \neq \emptyset$, a contradiction. \square

Step 2: We construct a stochastic process describing how the beliefs of an outside observer about the realized state ϕ are updated and establish its martingale properties. Let \mathcal{P} be the common set of priors given a (possibly mixed) strategy σ . Consider the following stochastic process, which is the same as in step 2 of the proof of Theorem 1 of [Ostrovsky \(2012\)](#) with the only difference that it is applied to each $p \in \mathcal{P}$ instead of the unique p . Nature draws a state $\phi \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$ and each player i observes $\Pi_i(\omega(\phi))$. Based on her private information and her strategy, player 1 announces y_1 . An outside observer, who shares the same set of beliefs \mathcal{P} and knows strategy σ but has no private information about the state ω , updates each $p \in \mathcal{P}$ using Bayes' rule. Note that the regularity of \mathcal{P} implies that all elements of \mathcal{P} are updated. Denote this set as \mathcal{P}_1 .

At time t_k , the outside observer updates these beliefs, denoted \mathcal{P}_k , using the public announcements up to t_k and the equilibrium strategies. Note that from the regularity of \mathcal{P} , each \mathcal{P}_k is compact and convex. As explained in [Ostrovsky \(2012\)](#), the process Q of updating $p \in \mathcal{P}$ at each time t is a martingale due to the law of iterated expectations. Given that it is also bounded (as it is between 0 and 1), the martingale convergence theorem implies that each Q converges to some random variable q_∞ . Since this is true for all $p \in \mathcal{P}$ and all corresponding martingales, we denote the set of the limits of all convergent beliefs by Q_∞ .

Step 3: We show that if the statement of Theorem 2 does not hold for this equilibrium, then we can identify a “non-vanishing arbitrage opportunity:” there is a player i^* and a positive number η^* , such that the continuation value of player i^* exceeds η^* at infinitely many trading times t_k .

Step 3, Case 1: Suppose that for some $\phi \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$, there is positive probability that some random variable $q \in Q_\infty$ assigns positive likelihoods to two states a and b with $X(a) \neq X(b)$, where q_k converges to q . As shown by [Ostrovsky \(2012\)](#), there exists probability distribution r assigning positive probability to both a and b , such that the following is true. For any $\varepsilon > 0$, there exist K and $\zeta > 0$ such that, for any $k > K$, the probability that q_k , which converges to q , is in the ε -neighbourhood of r is greater than ζ . This can be done for every $q \in Q_\infty$ and, in that case, the K can be selected uniformly because it is affected only by the uncertainty due to mixed strategies.⁴⁷

Any compact and convex set of beliefs \mathcal{P} which contains these limit probability distributions describes some uncertainty about X , hence it belongs to \mathcal{D} . Lemma 4 shows that there is player i and $\chi > 0$, such that i 's instant opportunity is greater than χ given \mathcal{P} and any previous announcement z .

As the definition of instant opportunity minimizes over all available beliefs, there is player i and $\chi > 0$ such that i 's instant opportunity (using any combination of q and p in the definition of instant opportunity) is greater than χ for any previous announcement z . By continuity, we can conclude that this is true for any combination of q_t and p_t (of the

46. The inequality is true because 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 p. 2618 in [Ostrovsky \(2012\)](#)).

47. Note that the only change in beliefs after some time t arises because they are weighted by the unique mixed strategy. Therefore, the convergence to q_∞ is uniform.

definition of instant opportunity), hence we get that the instant opportunity at t (for t big enough) is greater than $\chi > 0$ for any previous announcement for some probability $\zeta > 0$.⁴⁸

Concluding, for some i , $\chi > 0$, t_K and $\zeta > 0$, i 's instant opportunity at any time $t_{nk+i} > t_K$ is greater than χ with probability at least ζ , and thus for i , t_K and $\eta = \chi\zeta > 0$, the expected instant opportunity of player i at any time $t_{nk+i} > t_K$ is greater than η .

Step 3, Case 2: Suppose that there is zero probability that some $q \in \mathcal{Q}_\infty$ assigns positive likelihoods to two states a and b with $X(a) \neq X(b)$. That is, at the limit, the outside observer believes with certainty that the value of the security is equal to some x . As shown in Section A.2.4 of [Ostrovsky \(2012\)](#), almost surely (with probability 1), $\bigcup_{p \in \mathcal{Q}_\infty} \text{Supp}(p)$

contains the true state h . Hence, with probability 1, all $q \in \mathcal{Q}_\infty$ assign probability 1 to the value of the security being $X(h) = x$. In other words, the outside observer's belief about the value of the security converges to the intrinsic value.

Suppose that y_k does not converge in probability to the intrinsic value of the security. Then, there exist state h , numbers $\epsilon, \delta > 0$ such that when h is the true state and for any K , there exists $k > K$ such that the probability that $|y_k - X(h)| > \epsilon$ is greater than δ . As all players have more information than the outside observer, their beliefs about the value of the security also converges to the intrinsic value. This implies that for some player i and some $\eta > 0$, for any K , there exists $t_{nk+i} > t_K$ such that her expected instant opportunity is greater than η .

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

Step 4: This step concludes the proof by showing that the presence of a “non-vanishing arbitrage opportunity” is impossible in equilibrium.

Let $\mathcal{P}(H^{k-1})$ be the set of updated beliefs for the outside observer at time t_k , given the mixed equilibrium, the set of prior beliefs \mathcal{P} and history H^{k-1} . Note that with mixed strategies, H^{k-1} occurs with some probability. Moreover, because the equilibrium profile may consist of mixed strategies, $\mathcal{P}(H^k)$ may not be the same as $\mathcal{P}(H^{k-1})$, however for big enough t_k , they will have the same support on the state space Ω as it is finite. Consider such a big enough t_k .

Fix t_k , history H^{k-1} and suppose i makes an announcement. Her continuation payoff given history H^{k-1} and state ϕ is $V(H^{k-1}, \phi) = \min_{p \in \mathcal{P}(H^{k-1})} E_p \sum_{m=0}^{\infty} \beta^{nm} (s_{k+nm}(\phi') - s_{k+nm-1}(\phi'))$, where $s_{k+nm}(\phi')$ is the score at state ϕ' and time t_{k+nm} .

Using Proposition 4, we have that her continuation payoff $V(H^{k-1}, \phi)$ is greater than the one-period payoff from playing the myopic strategy at t_k . Because this is true for all states $\phi \in \bigcup_{p \in \mathcal{P}(H^{k-1})} \text{Supp}(p)$ that the outside observer considers possible at t_k , given history H^{k-1} , we have that $\min_{p \in \mathcal{P}(H^{k-1})} E_p V(H^{k-1}, \phi)$ is greater than i 's instant opportunity given beliefs $\mathcal{P}(H^{k-1})$.

Again using Proposition 4, the continuation payoff at t_k of each Trader $j \neq i$, who announces at t_k , is weakly positive at each state ϕ and history H^{k-1} . Since this is true for all states $\phi \in \bigcup_{p \in \mathcal{P}(H^k)} \text{Supp}(p)$, we have that

$$\min_{p \in \mathcal{P}(H^{k-1})} E_p V(H^{k-1}, \phi) \geq 0.$$

Since $\min_{p \in \mathcal{P}(H^{k-1})} E_p V(H^{k-1}, \phi)$ is weakly positive for each $i \in I$, we have that $\sum_{i \in I} E_p V(H^{k-1}, \phi)$ is weakly positive for any $p \in \mathcal{P}(H^{k-1})$. Moreover, it is strictly positive if i 's instant opportunity is strictly positive given $\mathcal{P}(H^{k-1})$. Since this is true for all $p \in \mathcal{P}(H^{k-1})$ and any previous announcement, by fixing prior $q \in \mathcal{P}$ and considering the (unique) probability over histories H^{k-1} that can arise at t_k , generated by the (possibly) mixed equilibrium, we can let Ψ_k be the sum of all players' expected continuation payoffs at t_k , divided by β^k as

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

The \bar{s}_k is the expected score of prediction y_k , where the expectation is over all ϕ , given the fixed $q \in \mathcal{P}$ and the moves of players according to the mixed equilibrium. We keep $q \in \mathcal{P}$ constant for all t_k . We then have that Ψ_k is weakly positive. 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.

48. A proper scoring rule may not be continuous. However, [Ostrovsky \(2012\)](#) shows, in footnote 19 of p. 2620, that the instant opportunity is continuous, the way he defines it. This implies that our instant opportunity is also continuous.

The last step is identical to that of [Ostrovsky \(2012\)](#) because all Ψ_k are calculated using the same prior $q \in \mathcal{P}$. The proof of Lemma 3 shows that $\lim_{K \rightarrow \infty} \sum_{k=1}^K \Psi_k = \chi_0$ for some finite χ_0 . From Step 3, this limit must be infinite because each Ψ_k is weakly positive and an infinite number of them is greater than η^* . Hence, both cases of Step 3 are impossible and y_k must converge to the intrinsic value of security X .

For Part (ii), suppose X is not strongly separable under Π and s . Then, there exist $\mathcal{P} \subseteq \Delta(\Omega)$, regular with respect to each Π_i and $v \in \mathbb{R}$, such that (a) $X(\omega) \neq v$ for some $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$ and (b) $d_{\mathcal{P}}(\Pi_i(\omega), v) = v$ for all $i = 1, \dots, n$ and $\omega \in \bigcup_{p \in \mathcal{P}} \text{Supp}(p)$.

Consider game $\Gamma^S(\Omega, I, \Pi, X, \mathcal{P}, y_0, y, \bar{y}, s, \beta)$, where the initial announcement of the market maker is $y_0 = v$. We will construct a Revision-Proof equilibrium (σ^*, \mathcal{P}) , where information does not aggregate. Define pair (σ^*, \mathcal{P}) , where σ^* specifies that each Trader i announces v after any history. At each information set \mathcal{I} of Trader i , set $\mathcal{P}(\mathcal{I}) = \mathcal{P}_{\Pi_i(\omega)}$. A player may deviate by not announcing at some period t_k the myopic best response v . All other players will continue announcing v in all subsequent periods and no information is revealed. Hence, she will not gain anything and her best response would be the myopic announcement v .

Since v is announced irrespective of whether a player deviates, there is never any information revealed and (σ^*, \mathcal{P}) is consistent. We now argue that it is not possible to find an alternative strategy that will make i 's future selves weakly better off and at least one strictly better off. If Trader i deviates, everyone else plays v and there is no updating of information so her future selves have the same beliefs as i . Since the myopically optimal is to play v for every future self, then it is not possible for such a deviation to exist.

C. Examples

In this section, we discuss the robustness of our results. We first show that the negative result that separable securities may not aggregate information under ambiguity does not depend on some priors assigning probability zero to the true state as in the example of Section 2. Such a case is illustrated in Example 1, where all priors have full support. We also use Example 1 to show that if a security assigns different values to each state, and therefore can predict all events, it is not necessarily strongly separable.

Example 1. Consider state space $\Omega = \{\omega_1, \dots, \omega_6\}$ and information structure with $\Pi_1 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}, \{\omega_5, \omega_6\}\}$, $\Pi_2 = \{\{\omega_1, \omega_2, \omega_6\}, \{\omega_3, \omega_4, \omega_5\}\}$ and $\Pi_3 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_5\}, \{\omega_4, \omega_6\}\}$. The security is $X(\omega_1) = X(\omega_5) = 0$, $X(\omega_2) = X(\omega_6) = 2$, $X(\omega_3) = 1$, and $X(\omega_4) = -1$.

To show that the security is separable, we show that the condition of Proposition 1 is always satisfied. In particular, for each $v \in \mathbb{R}$, we specify $\lambda_i : \Pi_i \rightarrow \mathbb{R}$ for $i = 1, 2, 3$ such that, for every state ω with $X(\omega) \neq v$,

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

Whenever $\lambda_i(\Pi_i(\omega))$ is not specified, it is implicitly set to 0.

- For $v \geq 2$, set $\lambda_1(\Pi_1(\omega)) < 0$ for all $\omega \in \Omega$,
- For $v \in [1, 2)$, set $\lambda_1(\Pi_1(\omega_1)) = -2$, $\lambda_2(\Pi_2(\omega_1)) = 1$, $\lambda_2(\Pi_2(\omega_3)) = -1$,
- For $v \in [0, 1)$, set $\lambda_1(\Pi_1(\omega_1)) = 1.4$, $\lambda_1(\Pi_1(\omega_2)) = 1.6$, $\lambda_1(\Pi_1(\omega_5)) = 1$, $\lambda_2(\Pi_2(\omega_1)) = -0.5$, $\lambda_2(\Pi_2(\omega_3)) = -4$, $\lambda_3(\Pi_3(\omega_1)) = -1$, $\lambda_3(\Pi_3(\omega_3)) = 2.7$, $\lambda_3(\Pi_3(\omega_4)) = 2$,
- For $v \in [-1, 0)$, set $\lambda_1(\Pi_1(\omega_1)) = 1$, $\lambda_1(\Pi_1(\omega_2)) = 1$, $\lambda_1(\Pi_1(\omega_5)) = 1$, $\lambda_2(\Pi_2(\omega_3)) = -1.5$, $\lambda_3(\Pi_3(\omega_3)) = 1$,
- For $v < -1$, set $\lambda_1(\Pi_1(\omega)) > 0$ for all $\omega \in \Omega$.

However, the security is not strongly separable. To see this, suppose that the market maker's initial announcement is $y_0 = 0.5$ and consider any strictly proper scoring rule. Given y_0 , consider any compact and convex set of priors that includes the priors $p_1 = \left(\frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{3}{8}, \frac{1}{8}\right)$, $p_2 = \left(\frac{6}{18}, \frac{1}{18}, \frac{7}{18}, \frac{2}{18}, \frac{1}{18}, \frac{1}{18}\right)$, and $p_3 = \left(\frac{3}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}\right)$. It is easy to check that the expectation of X , conditioning p_1 on Trader 1's information, is 0.5 at all states. The same is true for Trader 2 with p_2 and Trader 3 with p_3 . Using the third claim of Lemma 1, the myopic announcement is 0.5. This is true for all states and all traders. As X is not constant on Ω , it is not strongly separable, and there is no information aggregation at any state.

To show that a security that assigns different values to each state, and therefore can predict all events, is not always strongly separable, consider the following counter example. Let security X' that pays $\{1, 3.6, 6, 3.1, 3, 4\}$ in states ω_1 to ω_6 , respectively. We can easily check that, when $v = 3.5$ and $E = \Omega$, no trader knows at some state $\omega \in E$ whether the value of the security is always above or below 3.5, hence the condition of Proposition 2 is violated.

The previous example, together with that of Section 2, shows that information aggregation can fail for separable securities, when there are multiple priors. However, in both cases the failure occurs for a (potentially) unique announcement of the market maker. An interesting question is whether there are examples where the failure occurs for several different announcements from the market maker. We show here how such examples can easily be constructed.

Consider two examples, A and B, with the same set of traders I , state spaces Ω^A, Ω^B , prior beliefs $\mathcal{P}^A, \mathcal{P}^B$, securities X^A, X^B which are separable, information structures $\Pi^A = \{\Pi_i^A\}_{i \in I}, \Pi^B = \{\Pi_i^B\}_{i \in I}$ and suppose there is failure of information aggregation for initial announcements $x^A \neq x^B$ at states ω^A, ω^B , respectively. We can then create a new example, C, which is just the concatenation of the previous two, where the information aggregation failure occurs at both x^A and x^B . In particular, let $\Omega^C = \Omega^A \cup \Omega^B$ and $\Pi_i^C(\omega) = \Pi_i^A(\omega)$ if $\omega \in \Omega^A$, otherwise $\Pi_i^C(\omega) = \Pi_i^B(\omega)$. The set of priors \mathcal{P}^C consists of all priors p^C constructed as follows. For each $p^A \in \mathcal{P}^A, p^B \in \mathcal{P}^B$, construct $p^C = 1/2 p^A + 1/2 p^B$. Note that \mathcal{P}^C is compact, convex and regular with respect to Π^C .

Construct security X^C such that $X^C(\omega) = X^A(\omega)$ if $\omega \in \Omega^A$, otherwise $X^C(\omega) = X^B(\omega)$. From Proposition 1 and using the same λ_i , if X^A and X^B are separable, then so is X^C . Moreover, since Ω^C consists of two disjoint common knowledge events, Ω^A and Ω^B , there is no information aggregation for initial announcements $x^A \neq x^B$ at states $\omega^A, \omega^B \in \Omega^C$, respectively. By concatenating more examples like that, one can construct examples with multiple announcements where information aggregation fails at some state.

Finally, following [Ostrovsky \(2012\)](#), Example 2 illustrates how the MSR model can be reinterpreted as a basic model of trading with an automatic inventory-based market maker who offers to buy or sell at a price, which is a function of the (possibly negative) net inventory that he holds. In addition, we show that, in the inventory-based interpretation, information does not always aggregate in the presence of ambiguity averse traders.

Example 2. Suppose there are two traders, the state space is $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ and the information structure is $\Pi_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ and $\Pi_2 = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$. The security is given by $X(\omega_1) = 2, X(\omega_2) = X(\omega_3) = X(\omega_4) = 1$, and the price function is $q(z) = e^{-z}$, where z is the market maker's net inventory. The common set of priors is $\mathcal{P} = \text{conv} \left\{ \left(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}\right) \right\}$. Suppose that initially the market maker holds zero inventory of the security so that $z = 0$.

Suppose that the true state is ω_1 . First, Trader 1 makes a decision about how many shares of the security to buy or sell. We assume, for consistency, that the amount of shares belong to $Z = q^{-1}([y, \bar{y}])$, which is compact. Thus, the trader solves $\max_{z \in Z} \min_{p \in \mathcal{P}} E_p[\int_0^z q(\bar{z}) - X(\omega)d\bar{z}] = \min_{p \in \mathcal{P}} \max_{z \in Z} E_p[\int_0^z q(\bar{z}) - X(\omega)d\bar{z}]$. We have the equality by applying the same argument as in the proof of Lemma 1.⁴⁹

As in [Ostrovsky \(2012\)](#), given the price function, we can define the strictly proper scoring rule $s(X(\omega), y) = \int_0^{q^{-1}(y)} q(\bar{z}) - X(\omega)d\bar{z}$. We have that the price function is one-to-one continuous with a continuous inverse function. Therefore, we can conclude that in the MSR market, based on that strictly proper scoring rule, the trader solves $\max_{y \in [y, \bar{y}]} \min_{p \in \mathcal{P}} E_p[\int_0^{q^{-1}(y)} q(\bar{z}) - X(\omega)d\bar{z}] = \min_{p \in \mathcal{P}} \max_{y \in [y, \bar{y}]} E_p[\int_0^{q^{-1}(y)} q(\bar{z}) - X(\omega)d\bar{z}]$.⁵⁰ We shall show that if z^* solves the first optimization problem and y^* the second one, then it is $q(z^*) = y^*$ and that the revenues or losses are the same (i.e. $\max_{z \in Z} \min_{p \in \mathcal{P}} E_p[\int_0^z q(\bar{z}) - X(\omega)d\bar{z}] = \max_{y \in [y, \bar{y}]} \min_{p \in \mathcal{P}} E_p[\int_0^{q^{-1}(y)} q(\bar{z}) - X(\omega)d\bar{z}]$). The conclusion is that the purchase of the optimal amount of shares and the announcement of the myopic prediction are related with a one-to-one relation using the pricing function and that the two markets are equivalent in terms of revenues and losses.

We can observe that for every $p \in \mathcal{P}$, the amount z'_p that solves $\max_{z \in Z} E_p[\int_0^z q(\bar{z}) - X(\omega)d\bar{z}]$ is unique and that $q(z'_p) = E_p[X]$. Similarly, for every $p \in \mathcal{P}$, the prediction y'_p that solves $\max_{y \in [y, \bar{y}]} E_p[\int_0^{q^{-1}(y)} q(\bar{z}) - X(\omega)d\bar{z}]$ is $y'_p = E_p[X]$, hence $q^{-1}(y'_p) = z'_p$.

Therefore, for every $p \in \mathcal{P}$, we have that $E_p[\int_0^{z'_p} q(\bar{z}) - X(\omega)d\bar{z}] = E_p[\int_0^{q^{-1}(y'_p)} q(\bar{z}) - X(\omega)d\bar{z}]$. We can conclude that $\min_{p \in \mathcal{P}} E_p[\int_0^{z'_p} q(\bar{z}) - X(\omega)d\bar{z}] = \min_{p \in \mathcal{P}} E_p[\int_0^{q^{-1}(y'_p)} q(\bar{z}) - X(\omega)d\bar{z}]$ and it is achieved for the same p^* .

49. We use that $F(z) = \int_0^z q(\bar{z}) - X(\omega)d\bar{z}$ is continuous and we follow the arguments of Lemma 1.

50. Similarly, we follow the arguments of Lemma 1 with the continuous function $F(y) = \int_0^{q^{-1}(y)} q(\bar{z}) - X(\omega)d\bar{z}$.

The optimal quantity of shares z^* for the ambiguity averse trader is such that $q(z^*) = E_{p^*}[X]$ and the optimal prediction y^* is such that $y^* = E_{p^*}[X]$, and thus we get the conclusion.⁵¹

Finally, the first trader finds that the belief that achieves the minimum gives at state ω_1 zero probability. From the previous paragraph, we conclude that the optimal amount to purchase z^* is such that $q(z^*) = 0 \cdot 2 + 1 \cdot 1 = 1$ or equivalently (as long as q is 1-1) $z^* = 0$. Hence, she neither buys nor sells any shares (equivalently she would have announced 1 as her prediction). From symmetry, it is straightforward that the same would happen for every state in the partition cell $\{\omega_3, \omega_4\}$ and for Trader 2. The conclusion is that both traders do not purchase shares from the market maker and there is no information aggregation.

Acknowledgments. We would like to thank Laurence Ales, Sarah Auster, Giacomo Bonanno, Andrés Carvajal, Antonella Ianni, David Johnson, Max Kwiek, Miltos Makris, Alexander Michaelides, Michael Ostrovsky, Herakles Polemarchakis, Burkhard C. Schipper, Dominik Schmidt, and seminar participants at the University of California, Davis, FUR 2018 at York, Université Paris 1 Panthéon-Sorbonne, City, University of London, St. Andrews, Kent, WEHIA 2019, CRETE 2019, the Kent-Bristol-City Workshop in Economic Theory, and EWMES22 for helpful discussions and suggestions. Galanis would like to acknowledge financial support from the Economic and Social Research Council (ES/V004425/1). Ioannou would like to acknowledge financial support from the Agence Nationale de la Recherche (ANR-21-CE26-0001). Finally, we are indebted to the editor, Andrea Galeotti, and five anonymous referees for their insightful comments, which significantly improved the article. The usual disclaimer applies.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability Statement

The data and code underlying this article are available at <https://doi.org/10.5281/zenodo.10033669>.

REFERENCES

- ALCHOURRÓN, C. E., GÄRDENFORS, P. and MAKINSON, D. (1985), "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions", *The Journal of Symbolic Logic*, **50**, 510–530.
- ALES, L. and SLEET, C. (2014), "Revision Proofness", *Journal of Economic Theory*, **152**, 324–355.
- ANANTANASUWONG, K., KOUWENBERG, R., MITCHELL, O. S., *et al.* (2019), "Ambiguity Attitudes about Investments: Evidence from the Field" (Technical Report, National Bureau of Economic Research).
- ASHEIM, G. B. (1997), "Individual and Collective Time-Consistency", *The Review of Economic Studies*, **64**, 427–443.
- ATANASOV, P., RESCOBER, P., STONE, E., *et al.* (2017), "Distilling the Wisdom of Crowds: Prediction Markets vs. Prediction Polls", *Management Science*, **63**, 691–706.
- AUMANN, R. J. (1976), "Agreeing to Disagree", *Annals of Statistics*, **4**, 1236–1239.
- BACHARACH, M. (1985), "Some Extensions of a Claim of Aumann in an Axiomatic Model of Knowledge", *Journal of Economic Theory*, **37**, 167–190.
- BARELLI, P., BHATTACHARYA, S. and SIGA, L. (2020), "Full Information Equivalence in Large Elections" (Mimeo).
- BATES, J. M. and GRANGER, C. W. J. (1969), "The Combination of Forecasts", *Journal of the Operational Research Society*, **20**, 451–468.
- BATTIGALLI, P., CATONINI, E., LANZANI, G., *et al.* (2019), "Ambiguity Attitudes and Self-Confirming Equilibrium in Sequential Games", *Games and Economic Behavior*, **115**, 1–29.
- BEAUCHÊNE, D., LI, J. and LI, M. (2019), "Ambiguous Persuasion", *Journal of Economic Theory*, **179**, 312–365.
- BERG, J. E., NELSON, F. D. and RIETZ, T. A. (2008), "Prediction Market Accuracy in the Long Run", *International Journal of Forecasting*, **24**, 285–300.
- BERGE, C. (1963), *Topological Spaces: Including a Treatment of Multi-Valued Functions, Vector Spaces, and Convexity* (United Kingdom: Courier Dover Publications).
- BONANNO, G. (2013), "AGM-Consistency and Perfect Bayesian Equilibrium. Part I: Definition and Properties", *International Journal of Game Theory*, **42**, 567–592.
- BONANNO, G. (2016), "AGM-Consistency and Perfect Bayesian Equilibrium. Part II: From PBE to Sequential Equilibrium", *International Journal of Game Theory*, **45**, 1071–1094.
- BOSE, S. and DARIPA, A. (2009), "A Dynamic Mechanism and Surplus Extraction Under Ambiguity", *Journal of Economic Theory*, **144**, 2084–2114.
- BOSE, S. and RENOU, L. (2014), "Mechanism Design with Ambiguous Communication Devices", *Econometrica*, **82**, 1853–1872.

51. This is true because of the saddle point inequality and the uniqueness of the optimal quantity and prediction (given the belief p^*).

- BRIER, G. W. (1950), "Verification of Forecasts Expressed in Terms of Probability", *Monthly Weather Review*, **78**, 1–3.
- CAMERER, C. F. (1998), "Can Asset Markets be Manipulated? A Field Experiment with Racetrack Betting", *Journal of Political Economy*, **106**, 457–482.
- CARVAJAL, A. and CORREIA-DA SILVA, J. (2010), "Agreeing to Disagree with Multiple Priors" (Technical Report, Universidade do Porto, Faculdade de Economia do Porto).
- CAVE, J. A. K. (1983), "Learning to Agree", *Economics Letters*, **12**, 147–152.
- CERREIA-VIOGLIO, S., MACCHERONI, F., MARINACCI, M., et al. (2011), "Uncertainty Averse Preferences", *Journal of Economic Theory*, **146**, 1275–1330.
- CHAMBERS, C. P. (2008), "Proper Scoring Rules for General Decision Models", *Games and Economic Behavior*, **63**, 32–40.
- CHEN, J. Y. (2022), "Sequential Learning under Informational Ambiguity" (Mimeo).
- CHEN, Y., DIMITROV, S., SAMI, R., et al. (2010), "Gaming Prediction Markets: Equilibrium Strategies with a Market Maker", *Algorithmica*, **58**, 930–969.
- CHEN, Y., RUBERRY, M. and VAUGHAN, J. W. (2012), "Designing Informative Securities", in *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*.
- CHOO, L., KAPLAN, T. R. and ZULTAN, R. (2019), "Information Aggregation in Arrow–Debreu Markets: An Experiment", *Experimental Economics*, **22**, 625–652.
- CHOO, L., KAPLAN, T. R. and ZULTAN, R. (2022), "Manipulation and (Mis)trust in Prediction Markets", *Management Science*, **68**, 6716–6732.
- CONDIE, S. and GANGULI, J. V. (2011), "Ambiguity and Rational Expectations Equilibria", *The Review of Economic Studies*, **78**, 821–845.
- COWGILL, B. and ZITZEWITZ, E. (2015), "Corporate Prediction Markets: Evidence from Google, Ford, and Firm X", *The Review of Economic Studies*, **82**, 1309–1341.
- Cultivate Forecasts (2016), "Will the UK Vote to Leave the EU in the June 2016 Referendum?" <https://alphacast.cultivateforecasts.com/questions/1311-will-the-uk-vote-to-leave-the-eu-in-the-june-2016-referendum>, Accessed: 04-11-2018.
- Cultivate Labs (2021), "How Does the Logarithmic Market Scoring Rule (LMSR) Work?" <https://www.cultivatelabs.com/prediction-markets-guide/how-does-logarithmic-market-scoring-rule-lmsr-work>, Accessed: 04-17-2021.
- DEMARZO, P. and SKIADAS, C. (1998), "Aggregation, Determinacy, and Informational Efficiency for a Class of Economies with Asymmetric Information", *Journal of Economic Theory*, **80**, 123–152.
- DEMARZO, P. and SKIADAS, C. (1999), "On the Uniqueness of Fully Informative Rational Expectations Equilibria", *Economic Theory*, **13**, 1–24.
- DIMITROV, S. and SAMI, R. (2008), "Non-Myopic Strategies in Prediction Markets", in *Proceedings of the 9th ACM Conference on Electronic Commerce* 200–209.
- DOMINIAK, A. and LEFORT, J.-P. (2013), "Agreement Theorem for Neo-Additive Beliefs", *Economic Theory*, **52**, 1–13.
- DOMINIAK, A. and LEFORT, J.-P. (2015), "“Agreeing to Disagree” Type Results under Ambiguity", *Journal of Mathematical Economics*, **61**, 119–129.
- DUBEY, P., GEANAKOPOLOS, J. and SHUBIK, M. (1987), "The Revelation of Information in Strategic Market Games: A Critique of Rational Expectations Equilibrium", *Journal of Mathematical Economics*, **16**, 105–137.
- ECKEL, C. C. and GROSSMAN, P. J. (2002), "Sex Differences and Statistical Stereotyping in Attitudes Toward Financial Risk", *Evolution and Human Behavior*, **23**, 281–295.
- ECKEL, C. C. and GROSSMAN, P. J. (2008), "Forecasting Risk Attitudes: An Experimental Study Using Actual and Forecast Gamble Choices", *Journal of Economic Behavior and Organization*, **68**, 1–17.
- EICHBERGER, J., KELSEY, D. and GRANT, S. (2019), "Strategic Ambiguity in Multi-Stage Games with Almost Perfect Information" (Mimeo).
- EKMEKCI, M. and LAUERMANN, S. (2020), "Manipulated Electorates and Information Aggregation", *The Review of Economic Studies*, **87**, 997–1033.
- ELLIS, A. (2016), "Condorcet Meets Ellsberg", *Theoretical Economics*, **11**, 865–895.
- ELLIS, A. (2018), "On Dynamic Consistency in Ambiguous Games", *Games and Economic Behavior*, **111**, 241–249.
- EPSTEIN, L. G. and SCHNEIDER, M. (2003), "Recursive Multiple-Priors", *Journal of Economic Theory*, **113**, 1–31.
- FISCHBACHER, U. (2007), "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments", *Experimental Economics*, **10**, 171–178.
- FRÉCHETTE, G. R. and YUKSEL, S. (2017), "Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination", *Experimental Economics*, **20**, 279–308.
- FUDENBERG, D. and TIROLE, J. (1991), "Perfect Bayesian Equilibrium and Sequential Equilibrium", *Journal of Economic Theory*, **53**, 236–260.
- GALANIS, S. (2011), "Syntactic Foundations for Unawareness of Theorems", *Theory and Decision*, **71**, 593–614.
- GALANIS, S. (2013), "Unawareness of Theorems", *Economic Theory*, **52**, 41–73.
- GALANIS, S. (2021), "Dynamic Consistency, Valuable Information and Subjective Beliefs", *Economic Theory*, **71**, 1467–1497.
- GALANIS, S. and KOTRONIS, S. (2021), "Updating Awareness and Information Aggregation", *The B.E. Journal of Theoretical Economics*, **21**, 613–635.
- GEANAKOPOLOS, J. and POLEMARCHAKIS, H. (1982), "We Can't Disagree Forever", *Journal of Economic Theory*, **28**, 192–200.

- GILBOA, I. and SCHMEIDLER, D. (1989), "Maxmin Expected Utility with Non-Unique Prior", *Journal of Mathematical Economics*, **18**, 141–153.
- GOLDMAN, S. M. (1980), "Consistent Plans", *The Review of Economic Studies*, **47**, 533–537.
- GOLOSOV, M., LORENZONI, G. and TSYVINSKI, A. (2014), "Decentralized Trading with Private Information", *Econometrica*, **82**, 1055–1091.
- GROSSMAN, S. (1976), "On the Efficiency of Competitive Stock Markets Where Trades Have Diverse Information", *The Journal of Finance*, **31**, 573–585.
- HANANY, E. and KLIBANOFF, P. (2007), "Updating Preferences with Multiple Priors", *Theoretical Economics*, **2**, 261–298.
- HANANY, E. and KLIBANOFF, P. (2009), "Updating Ambiguity Averse Preferences", *The B.E. Journal of Theoretical Economics*, **9**, 1–53.
- HANANY, E., KLIBANOFF, P. and MUKERJI, S. (2020), "Incomplete Information Games with Ambiguity Averse Players", *American Economic Journal: Microeconomics*, **12**, 135187.
- HANSEN, L. P. and SARGENT, T. J. (2001), "Robust Control and Model Uncertainty", *American Economic Review: Papers and Proceedings*, **91**, 60–66.
- HANSON, R. (2003), "Combinatorial Information Market Design", *Information Systems Frontiers*, **5**, 107–119.
- HANSON, R. (2007), "Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation", *The Journal of Prediction Markets*, **1**, 3–15.
- HANSON, R. and OPREA, R. (2009), "A Manipulator Can Aid Prediction Market Accuracy", *Economica*, **76**, 304–314.
- HANSON, R., OPREA, R. and PORTER, D. (2006), "Information Aggregation and Manipulation in an Experimental Market", *Journal of Economic Behavior & Organization*, **60**, 449–459.
- HARRISON, J. M. and KREPS, D. M. (1978), "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations", *The Quarterly Journal of Economics*, **92**, 323–36.
- HAYEK, F. A. (1945), "The Use of Knowledge in Society", *The American Economic Review*, **35**, 519–530.
- HEALY, P. J., LINARDI, S., LOWERY, J. R., *et al.* (2010), "Prediction Markets: Alternative Mechanisms for Complex Environments with Few Traders", *Management Science*, **56**, 1977–1996.
- HEATH, C. and TVERSKY, A. (1991), "Preference and Belief: Ambiguity and Competence in Choice under Uncertainty", *Journal of Risk and Uncertainty*, **4**, 5–28.
- HELLWIG, M. F. (1982), "Rational Expectations Equilibrium with Conditioning on Past Prices: A Mean-Variance Example", *Journal of Economic Theory*, **26**, 279–312.
- IOANNOU, C. A., MATHEVET, L., ROMERO, J., *et al.* (2023), "Data Mining in Repeated Games" (Mimeo).
- IOANNOU, C. A. and ROMERO, J. (2014), "A Generalized Approach to Belief Learning in Repeated Games", *Games and Economic Behavior*, **87**, 178–203.
- JIAN, L. and SAMI, R. (2012), "Aggregation and Manipulation in Prediction Markets: Effects of Trading Mechanism and Information Distribution", *Management Science*, **58**, 123–140.
- KAJII, A. and UI, T. (2005), "Incomplete Information Games with Multiple Priors", *The Japanese Economic Review*, **56**, 332–351.
- KAJII, A. and UI, T. (2009), "Interim Efficient Allocations under Uncertainty", *Journal of Economic Theory*, **144**, 337–353.
- KELLNER, C. and LE QUEMENT, M. T. (2017), "Modes of Ambiguous Communication", *Games and Economic Behavior*, **104**, 271–292.
- KELLNER, C. and LE QUEMENT, M. T. (2018), "Endogenous Ambiguity in Cheap Talk", *Journal of Economic Theory*, **173**, 1–17.
- KILKA, M. and WEBER, M. (2001), "What Determines the Shape of the Probability Weighting Function under Uncertainty?", *Management Science*, **47**, 1712–1726.
- KLIBANOFF, P., MARINACCI, M. and MUKERJI, S. (2005), "A Smooth Model of Decision Making Under Ambiguity", *Econometrica*, **73**, 1849–1892.
- KYLE, A. S. (1985), "Continuous Auctions and Insider Trading", *Econometrica*, **53**, 1315–35.
- LAMBERT, N. S., OSTROVSKY, M. and PANOV, M. (2018), "Strategic Trading in Informationally Complex Environments", *Econometrica*, **86**, 1119–1157.
- MACCHERONI, F., MARINACCI, M. and RUSTICHINI, A. (2006a), "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences", *Econometrica*, **74**, 1447–1498.
- MACCHERONI, F., MARINACCI, M. and RUSTICHINI, A. (2006b), "Dynamic Variational Preferences", *Journal of Economic Theory*, **128**, 4–44.
- MAILATH, G. J. and SAMUELSON, L. (2020), "Learning under Diverse World Views: Model-Based Inference", *American Economic Review*, **110**, 1464–1501.
- MCKELVEY, R. D. and PAGE, T. (1986), "Common Knowledge, Consensus, and Aggregate Information", *Econometrica*, **54**, 109–27.
- MCKELVEY, R. D. and PAGE, T. (1990), "Public and Private Information: An Experimental Study of Information Pooling", *Econometrica*, **58**, 1321–39.
- MELE, A. and SANGIORGI, F. (2015), "Uncertainty, Information Acquisition and Price Swings in Asset Markets", *The Review of Economic Studies*, **82**, 1533–1567.
- NIELSEN, L. T. (1984), "Common Knowledge, Communication, and Convergence of Beliefs", *Mathematical Social Sciences*, **8**, 1–14.

- NIELSEN, L. T., BRANDENBURGER, A., GEANAKOPOLOS, J., *et al.* (1990), "Common Knowledge of an Aggregate of Expectations", *Econometrica*, **58**, 1235–1239.
- OSTROVSKY, M. (2012), "Information Aggregation in Dynamic Markets With Strategic Traders", *Econometrica*, **80**, 2595–2647.
- OTTAVIANI, M. and SØRENSEN, P. N. (2007), "Outcome Manipulation in Corporate Prediction Markets", *Journal of the European Economic Association*, **5**, 554–563.
- PAGE, L. and SIEMROTH, C. (2017), "An Experimental Analysis of Information Acquisition in Prediction Markets", *Games and Economic Behavior*, **101**, 354–378.
- PAGE, L. and SIEMROTH, C. (2021), "How Much Information is Incorporated Into Financial Asset Prices? Experimental Evidence", *The Review of Financial Studies*, **34**, 4412–4449.
- PAHLKE, M. (2022), "Dynamic Consistency in Incomplete Information Games with Multiple Priors", *Games and Economic Behavior*, **133**, 85–108.
- PELEG, B. and YAARI, M. E. (1973), "On the Existence of a Consistent Course of Action When Tastes are Changing", *The Review of Economic Studies*, **40**, 391–401.
- PIRES, P. (2002), "A Rule for Updating Ambiguous Beliefs", *Theory and Decision*, **53**, 137–152.
- RADNER, R. (1979), "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices", *Econometrica*, **47**, 655–678.
- RHODE, P. W. and STRUMPF, K. S. (2004), "Historical Presidential Betting Markets", *Journal of Economic Perspectives*, **18**, 127–142.
- SCHLEGEL, J. C., KWAŚNICKI, M. and MAMAGEISHVILI, A. (2022), "Axioms for Constant Function Market Makers" Available at SSRN.
- SEBENIUS, J. and GEANAKOPOLOS, J. (1983), "Don't Bet on It: Contingent Agreements with Asymmetric Information", *Journal of the American Statistical Association*, **78**, 424–426.
- SIGA, L. and MIHM, M. (2021), "Information Aggregation in Competitive Markets", *Theoretical Economics*, **16**, 161–196.
- SINISCALCHI, M. (2011), "Dynamic Choice under Ambiguity", *Theoretical Economics*, **6**, 379–421.
- SNOWBERG, E., WOLFERS, J. and ZITZEWITZ, E. (2013), "Prediction Markets for Economic Forecasting", in *Handbook of Economic Forecasting* (Vol. 2) (United Kingdom: Elsevier) 657–687.
- STROTZ, R. H. (1955), "Myopia and Inconsistency in Dynamic Utility Maximization", *The Review of Economic Studies*, **23**, 165–180.
- The Economist (2021), "How Spooks are Turning to Superforecasting in the Cosmic Bazaar" <https://www-economist-com.ezphost.dur.ac.uk/science-and-technology/2021/04/15/how-spooks-are-turning-to-superforecasting-in-the-cosmic-bazaar>. Accessed: 04-17-2021.
- The Financial Times (2016a), "Brexit Poll Tracker" <https://ig.ft.com/sites/brexit-polling>. Accessed: 04-6-2017.
- The Financial Times (2016b), "How Accurate are the Brexit Polls?" <https://www.ft.com/content/6a63c2ca-2d80-11e6-bf8d-26294ad519fc>. Accessed: 01-24-2023.
- VEIGA, H. and VORSATZ, M. (2010), "Information Aggregation in Experimental Asset Markets in the Presence of a Manipulator", *Experimental Economics*, **13**, 379–398.
- VESPA, E. (2019), "An Experimental Investigation of Strategies in the Dynamic Common Pool Game", *International Economic Review*, **61**, 417–440.
- WOLFERS, J. and LEIGH, A. (2002), "Three Tools for Forecasting Federal Elections: Lessons from 2001", *Australian Journal of Political Science*, **37**, 223–240.
- WOLFERS, J. and ZITZEWITZ, E. (2004), "Prediction Markets", *Journal of Economic Perspectives*, **18**, 107–126.
- WOLINSKY, A. (1990), "Information Revelation in a Market with Pairwise Meetings", *Econometrica*, **58**, 1–23.
- ZITZEWITZ, E. (2007), "Manipulation Can Affect Prices" <http://www.midasoracle.org/2007/05/30/manipulation-can-affect-prices/>. Accessed: 04-6-2017.