Department of Economics

Information Aggregation Under Ambiguity: Theory and Experimental Evidence

Spyros Galanis[1]
City, University of London

Christos A. Ioannou
Université Paris 1 Panthéon - Sorbonne

Stelios Kotronis
City, University of London

Department of Economics
Discussion Paper Series
No. 20/05

[1] Corresponding author: Spyros Galanis, Department of Economics, City, University of London, Northampton Square, London EC1V 0HB, UK. E-mail: Spyros.Galanis@city.ac.uk

Information Aggregation Under Ambiguity: Theory and Experimental Evidence*

Spyros Galanis[†] Christos A. Ioannou[‡] Stelios Kotronis[§]

November 25, 2019

Abstract

We study information aggregation in a dynamic trading model with partially informed and ambiguity averse traders. We show theoretically that separable securities, introduced by Ostrovsky (2012) in the context of Subjective 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 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 theoretical predictions, which we are able to confirm in the lab.

JEL Classification Numbers: D82, D83, D84, G14, G41.

Keywords: Information Aggregation, Ambiguity Aversion, Financial Markets, Information Markets, Prediction Markets.

1 Introduction

When do financial markets aggregate information, which is dispersed among individual traders? The mechanism, through prices, is intuitive. If the price is low (high) and some

*We would like to thank Laurence Ales, Sarah Auster, Giacomo Bonanno, Andres Carvajal, Antonella Ianni, Max Kwiek, Alexander Michaelides, Michael Ostrovsky, Herakles Polemarchakis, Burkhard C. Schipper and seminar participants at the University of California, Davis, FUR 2018 at York, City, University of London, St. Andrews, Kent, WEHIA 2019 and CRETE 2019 for helpful discussions and suggestions.

[†]Department of Economics, City, University of London, spyros.galanis@city.ac.uk.

[‡]Centre d'Économie de la Sorbonne, Université Paris 1 Panthéon - Sorbonne, christos.ioannou@univ-paris1.fr.

[§]Department of Economics, City, University of London, stelioskotronis@gmail.com.

traders have private information that the real value of the security is high (low), they will increase (decrease) their demand and the price. Moreover, these price movements could reveal to a trader information that others might have, prompting her to update her beliefs and either buy or sell, thus further revealing some of her own private information.

This intuition is largely correct. Starting from Hayek (1945), there is an extensive literature showing that under various settings, information gets aggregated. For instance, Ostrovsky (2012) shows that even if there are few large and strategic traders, information gets aggregated for a large class of securities, called *separable*, which includes the Arrow-Debreu securities.

In recent years, many firms and institutions have leveraged this property of information aggregation by designing *prediction markets*, as a tool to better forecast political events, sales of products and box office successes, among others (O'Leary (2011)). Google, Microsoft, Ford, General Electric and HP, to name a few, run internal prediction markets as a corporate governance and prediction's tool, whereas Cultivate Labs, Inkling Markets, Consensus Point, Crowdcast and Iowa Electronic Markets are examples of Internet-based prediction markets.[1] These markets are usually implemented using a Market Scoring Rule (MSR) (Hanson (2003, 2007)), which, in turn, is based on proper scoring rules (e.g. Brier (1950)).[2]

In several cases, prediction markets perform significantly better than other conventional forecasting methods, such as polls. Berg et al. (2008) compared the predictions, for 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 in advance it outperformed the polls at every election.

In other cases, however, where the question is less common, the performance of prediction markets is mixed. On the one hand, Dreber et al. (2015) show that a prediction market, populated by researchers, was better at predicting the reproducibility of 44 studies published in prominent psychological journals, as compared to the pre-trade average of the market participants' individual forecasts. On the other hand, Camerer et al. (2016) show that the two methods are equally capable of predicting the reproducibility of economics studies.

Interestingly, in the case of a "once in a lifetime" event, prediction markets fared 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 but failed spectacularly, as the closing prediction was a 20% probability of voting for Brexit.[3] On the other hand, 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. The actual result was 51.9% and 48.1%, respectively.[4]

Our motivation for this paper stems from trying to understand when prediction (and more generally financial) markets are efficient at aggregating information. In particular, are they efficient at aggregating information for events that are rare or uncommon and for

[1] Such an implementation is described at https://www.cultivatelabs.com/prediction-markets-guide/how-does-logarithmic-market-scoring-rule-lmsr-work.

[2] McKelvey and Page (1990) were the first to use a proper scoring rule to aggregate information in a market.

[3] The market can be found at https://alphacast.cultivateforecasts.com/questions/1311-will-the-uk-vote-to-leave-the-eu-in-the-june-2016-referendum.

[4] The details can be found at https://ig.ft.com/sites/brexit-polling/.

which beliefs are imprecise? Up to our awareness, the literature has so far only focused on traders who have precise probabilities about events and Subjective Expected Utility (SEU) preferences. But is this always a valid assumption, especially for events, like the outcome of the Brexit referendum or the recent financial crisis, which occur once in a generation?

Using a simple theoretical example, we show that separable securities may fail to aggregate information if traders are ambiguity averse and have imprecise beliefs. Moreover, these securities are prone to manipulation, as the degree of information aggregation is greatly influenced by the initial price, set by an uninformed market maker. These observations are also confirmed in the lab, where we run an experiment using prediction markets.

Our main contribution is to propose a new class of *strongly separable* securities, which are robust to the above considerations. In particular, we show that they are necessary and sufficient for information aggregation, in both strategic and non-strategic environments with ambiguity averse traders, in a prediction market which implements the MSR. Moreover, they are not prone to manipulation, as the initial price does not influence the degree of information aggregation. We derive several theoretical predictions, which we are able to confirm in the lab.

The economic value of getting better predictions is difficult to estimate, however a benchmark is the revenues of the forecasting industry, at around $300 billion (Atanasov et al. (2016)). Our results, however, are not restricted to the case of prediction markets. We further show that there does not exist a security that is strongly separable for all information structures.[5] This is in contrast to the case of separable securities where, for example, Arrow-Debreu securities are separable for all information structures. Because strongly separable securities are both sufficient and necessary for information aggregation, we have the following negative result. With ambiguity aversion and imprecise beliefs, there is no security that aggregates information for *all* information structures. If we cannot find such a security in the special case of prediction markets, we cannot hope to find one in the case of general financial markets. In other words, imprecise beliefs can severely constrain the ability of markets to generically aggregate information.

To build some intuition, consider an example with a unique common prior and SEU, which shows how prices aggregate information in the simple case of two myopic, or non-strategic, traders. The formal treatment of this example, in the case of imprecise beliefs, is presented in Section 2. In Section 6, we use the same example to run an experiment with prediction markets and compare information aggregation across separable and strongly separable securities.

The prediction market concerns the 2016 Brexit referendum in the UK. Suppose there are three possible states: the referendum takes place and the result is in favour of Brexit, it takes place but the outcome is against Brexit, or the referendum is cancelled. Let Arrow-Debreu security X, which pays 1 if Brexit occurs and 0 otherwise.

After an initial announcement by the market maker, the two traders take turns in announcing a value for the security X. If Trader i announces a low value whereas j announces in the next round a high value, it is as if j buys the security from i through the market maker. Payoffs are determined using the MSR. For non-strategic traders, who only care

[5]Strongly separable securities, just like separable ones, only depend on the information structure and not on the specific scoring rule.

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

Figure 1: Private signals

about their current payoff, MSR ensures that the trader's optimal strategy is to announce the expected value of the security given her posterior beliefs.

The private information of each trader is depicted in Figure 1. In particular, Trader 1 is informed whether the referendum is cancelled or not, whereas if the referendum does take place Trader 2 then knows its outcome. As a result, their pooled information always reveals the true state.

Suppose that Brexit is the true state and the two traders have a common prior which assigns strictly positive probability to that state. Then, by announcing sequentially their expectations about the value of security X, information gets aggregated. To see this, note that, in the first round, the announcement of Trader 1 about X is strictly positive, hence Trader 2 realizes that Trader 1 assigns probability 0 to the state that the referendum is cancelled, as this would imply an announcement of 0. The public information revealed is that the referendum will take place. In the second round, Trader 2, by combining this extra information with her own private signal, realizes that Brexit will happen and announces 1. In the third round, Trader 1 realizes as well that Brexit is the true state, hence she also announces 1 and the prediction market aggregates all information.

This result of information aggregation relies heavily on the assumption that each trader's belief is a unique prior.[6] However, Brexit is a once in a lifetime event for which no historical data of similar events exist. How can we be sure that the traders have precise probabilities for such a hard-to-quantify event? If we cannot maintain the hypothesis of a unique prior and SEU, it is no longer the case that markets aggregate information, even if the traders' (multiple) priors are common. More importantly, even a slight departure from a unique prior could result in the traders agreeing on a value of the security that is very far from the true one.

To show this, consider the ambiguity aversion model of Gilboa and Schmeidler (1989), where a decision maker acts as if having multiple priors over the states and chooses the prediction that maximizes the minimum expected utility over these priors. An important insight, which we prove in Lemma 1 and use heavily is that, with multiple priors, MSR ensures that the optimal announcement of a myopic trader is still unique and the expectation of security X according to *one* of her beliefs. The choice of the belief depends on the previous announcement, thus introducing path-dependence (which is absent if the prior is unique). However, if the previous announcement happens to be the expectation of X according to some of i's beliefs, then i's optimal myopic strategy is to repeat it.

Trader 1's private information is such that she does not distinguish between states Brexit and No Brexit, because her signal only informs her whether the referendum is cancelled or

[6] More generally, the assumption we need is that of a common prior.

not. What is the announcement that she makes at these two states? Suppose that the market maker's initial announcement is 0. Unlike the SEU framework, this initial announcement is important for the result because of path-dependence. For instance, suppose that at least one prior (but not all) assigns zero probability to Brexit. When Trader 1 learns that either Brexit or No Brexit are true, she updates her beliefs so that one of her posteriors assigns 0 probability to Brexit. Because the expectation of X according to that belief is also 0, from the argument of the previous paragraph we conclude that Trader 1 will also announce 0.[7]

If the referendum was cancelled, Trader 1 would know this and announce 0. Because the same announcement of 0 would be made in all possible states, no public information is revealed from her announcement. As a result, Trader 2 does not learn anything from 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 no one learns that Brexit will happen, hence no one learns that the true value of the security is 1. The traders agree on their announcements, but on a value for the security that is very far from the truth.

Section 2 provides the formal treatment and we confirm the following observations. First, the result of no information aggregation with separable securities and imprecise beliefs does not depend on having a prior which assigns probability 0 to the true state. As we show in Example 1 of Appendix C, separable securities may not aggregate information also in the case of full support priors. Moreover, these results hold for all proper scoring rules. Second, the initial announcement by the market maker matters. If the announcement was 1, there would be information aggregation, so there is path-dependence and the possibility of manipulation by the uninformed market-maker. Although in the current example information aggregation fails for a unique initial announcement, it is easy to construct examples where this is not the case, as we argue in Appendix C. Finally, even if the imprecision of beliefs is vanishingly small, there is still no information aggregation.

To accommodate the case of imprecise probabilities, consider security Y that pays 3 if Brexit occurs, 2 if there is No Brexit and 1 if the referendum is cancelled. In states Brexit and No Brexit, Trader 1 considers only these states to be possible. Whatever are her beliefs, she announces an expected value of Y between 3 and 2. If the referendum is cancelled, she knows this and she announces 1. In other words, the announcement about Y differs with respect to whether the referendum takes place, because 1 is not a convex combination of 2 and 3. Since the possible announcements are different, in the next round Trader 2 can infer that Trader 1's signal cannot be that the referendum is cancelled. Combining this piece of information with her own signal, Trader 2 concludes that Brexit is the true state, hence announces 3. In the next round, Trader 1 can infer as well that Trader 2 knows that Brexit will happen, therefore information gets aggregated.

Note that security Y aggregates information for any initial announcement of the market maker and 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 Y 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. Theorem 1 characterizes information aggregation in

[7]In Section 2, we provide the formal details for the case of a quadratic scoring rule. However, the argument works for any proper scoring rule.

terms of strongly separable securities, for the case of myopic or non-strategic players.

For the strategic case, 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 will not be optimal for the same trader at a later time. This feature is absent in the SEU framework of Ostrovsky (2012) and further complicates our analysis. Theorem 2 shows that strongly separable securities are both necessary and sufficient for information aggregation for all interim equilibria. Theorem 3 provides the same result using the revision-proof equilibrium, which has been studied by Asheim (1997) and Ales and Sleet (2014), in the context of infinite horizon and complete information games with time-inconsistent preferences.

We conclude by emphasizing that even though we restrict attention to prediction markets, our results are more general. In particular, prediction markets can be reinterpreted in order to correspond to the classic approach with an inventory-based 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.

1.1 Related literature

There is a large literature related to information aggregation and information revelation in dynamic markets, starting with Hayek (1945). 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), Nielsen (1984), McKelvey and Page (1986), Dubey et al. (1987), Wolinsky (1990), Nielsen et al. (1990) and Golosov et al. (2014), among others. Siga and Mihm (2018) provide microfoundations for REE, using common-value auctions, and study when prices aggregate information.

The *no trade theorems* stem from Aumann (1976) and Milgrom and Stokey (1982). Geanakoplos and Polemarchakis (1982), Cave (1983), Sebenius and Geanakoplos (1983), Nielsen (1984), Bacharach (1985) and Nielsen et al. (1990) study information communication in a non-strategic setting, where agents announce posterior beliefs or other aggregate statistics. However, they do not fully characterize under what conditions the consensus yields the true posterior or expectation of the security, a gap which is filled by DeMarzo and Skiadas (1998, 1999).

Ostrovsky (2012) and Chen et al. (2012) show that in a market with either myopic or strategic traders, separable securities are both necessary and sufficient for information aggregation. Their models are based on MSR and hence their results are directly applicable to prediction markets. Similar approaches can be found in Chen et al. (2010) and Dimitrov and Sami (2008), where the focus is on considering whether information gets aggregated when traders are strategic, under various assumptions regarding the signal structure. Galanis and Kotronis (2019) study the information aggregation properties of separable securities in an environment with unawareness.

Ostrovsky (2012) provides similar results for separable securities, using the model of Kyle (1985) which includes noise traders and competitive market makers. However, in this 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 revealed eventually. As pointed out by Ostrovsky (2012), the MSR focuses on the issue of information aggregation. Lambert et al. (2018) study trading in informationally complex environments, using the single-period version of Kyle (1985). They show that there is always a unique linear equilibrium and, under some conditions, prices in large markets aggregate all available information. Information aggregation has been studied in several other settings, for example in elections (e.g. Barelli et al. (2017), Ekmekci and Lauermann (2019)).

Dominiak and Lefort (2013, 2015), Carvajal and Correia-da Silva (2010) and Kajii and Ui (2005, 2009) extend the result of Aumann (1976) in an environment with ambiguity aversion. Finally, within a REE setting with ambiguity averse preferences, the existence and robustness of partially-revealing rational expectations equilibria is shown in Condie and Ganguli (2011).

In dynamic choice problems under uncertainty, agents who violate the sure-thing principle of Savage (1954) may be dynamically inconsistent. One way of dealing with dynamic inconsistency is the concept of consistent planning, which was introduced by Strotz (1955) and further developed by Peleg and Yaari (1973) and Goldman (1980). Siniscalchi (2011) provides behavioural foundations in a single-agent setting. In an environment with Maxmin Expected Utility (MEU) preferences, Epstein and Schneider (2003) show that prior by prior updating of a rectangular set of priors preserves dynamic consistency.

Few papers study equilibrium notions in general dynamic games under ambiguity, such as Hanany et al. (2018), Eichberger et al. (2017), Pahlke (2018) and Battigalli et al. (2019). 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).

Prediction markets have been studied extensively, both experimentally and with real data. An overview of the literature is provided by Wolfers and Zitzewitz (2004) and Snowberg et al. (2013). One of the main questions is whether prediction markets can be manipulated. Most of the literature finds very little evidence of price manipulation, both in actual markets (Camerer (1998), Rhode and Strumpf (2004), Wolfers and Leigh (2002)) and in the lab (Hanson et al. (2006), Hanson and Oprea (2009)). However, Zitzewitz (2007) and Snowberg et al. (2013) describe one case where a manipulator was able to sustain an artificially high price on the contract that paid if Hillary Clinton became the next U.S. President. Veiga and Vorsatz (2010) show experimentally that, under some conditions, prices can be manipulated by an uninformed trader. We also find evidence of price manipulation, but through the completely different channel of imprecise beliefs and the actions of the market maker, that has not been studied before in the literature. Ottaviani and Sørensen (2007) provide the first formal analysis of outcome manipulation in corporate prediction markets, where traders are able to influence the outcome.

Atanasov et al. (2016) compare prediction markets with prediction polls, in one of the first large-scale experimental tests, using the forecasts on 261 events, of more than 2400 participants. They find that prediction markets outperform the simple means of prediction polls, but they fare worse when compared with forecasts from prediction polls that are statistically aggregated, using criteria such as past performance. The experiment of Jian and Sami (2012) on prediction markets is the closest to our experimental design. However, there are several differences. First, it only specifies an environment with SEU, as all other papers that we are aware of. Second, it does not directly test the effect of the initial price on the degree of information aggregation.

Healy et al. (2010) compare the effectiveness of prediction markets with other trading mechanisms, such as double actions, finding that no mechanism performs best under all conditions. In such a setting, Alfarano et al. (2019) study the effect of the connectivity of the market on trading. They find experimentally that information aggregation does not depend on whether a market is more connected or not.

The paper is organized as follows. Section 3 describes the model, whose components are the ambiguity averse preferences, the MSR trading environment, the decision function for the myopic case and the properties for the strategic case. In Section 4, we characterize information aggregation for the case of myopic traders, whereas in Section 5 we examine the strategic case. In Section 6, we describe our experiment and discuss the support for our theory. We conclude in Section 7. All proofs are included in the appendices. The experimental instructions are included in the Online Appendix.

2 Example

We illustrate our approach by providing the formal details of the example discussed in the Introduction. The same example is used in our experiment, which we discuss in Section 6.

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. 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. This can be interpreted as if, 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 2's is $\Pi_2 = \{\{\omega_1, \omega_3\}, \{\omega_2\}\}$. They trade an Arrow-Debreu security X that pays 1 at ω_1 and 0 otherwise.

The two traders share a common set of priors $\mathcal{P}$, which is the convex hull of $p^1 = (0, \frac{1}{2}, \frac{1}{2})$ and $p^2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Their preferences are represented by Maxmin Expected Utility (MEU) with beliefs $\mathcal{P}$ and $u(x) = x$, $x \in \mathbb{R}$. If trader i's announcement is y, the true value of the security is $x^* = X(\omega)$ and the announcement of the previous trader (or the market maker) is y_{-1}, i's utility is $s(y, x^*) - s(y_{-1}, x^*)$, where $s(y, x^*) = -(y - x^*)^2$ is the quadratic scoring rule (or more generally a proper scoring rule).

Because i does not know what the true state is, her utility (in the current period) from announcing y is

$$U^i(y) = \min_{q \in \mathcal{P}} \sum_{\omega \in \Omega} [s(y, X(\omega)) - s(y_{-1}, X(\omega))] q(\omega).$$

Proper scoring rules have the attractive property in the SEU framework that a myopic trader, who only cares about her current period utility, will truthfully announce her expectation of X. Formally, if $\mathcal{P} = \{p\} \in \Delta(\Omega)$, then $\underset{y \in Y}{argmax}\ U^i(y) = E_p[X]$. Note that this announcement is independent of the previous announcement, y_{-1}. With MEU preferences, Lemma 1 establishes that $\underset{y \in Y}{argmax}\ U^i(y) = E_p[X]$ for *some* $p \in \mathcal{P}$, which depends on y_{-1}, thus introducing path-dependence. Moreover, if $y_{-1} = E_p[X]$ for some $p \in \mathcal{P}$, then $\underset{y \in Y}{argmax}\ U^i(y) = y_{-1}$.

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 maximises her utility myopically. Using Lemma 1 and letting p_{E_1} be the conditional of p given E_1, she solves $\min_{p \in \mathcal{P}} E_{p_{E_1}}[s(E_{p_{E_1}}[X], X(\omega)) - s(0, X(\omega))] = \min_{p \in \mathcal{P}}[p_{E_1}(\omega_1)^2(2 - p_{E_1}(\omega_1) - p_{E_1}(\omega_2))] = \min_{p \in \mathcal{P}} p_{E_1}(\omega_1)^2$. 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 true value of X was 0 and she would announce 0.

The above imply that Trader 2 cannot learn anything from 1's announcement, hence can only rely on her private signal $E_2 = \{\omega_1, \omega_3\}$. Maximising 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. Because the true value of the security is 1, there is no information aggregation through the announcements, even though their pooled information would reveal state ω_1 and the true value, 1.

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.

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 information to Trader 2. 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. However, this is due to the simplicity of the model. 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 1 will also announce 0.

3 Model

In this section, we first describe the ambiguity averse preferences of the traders and the market scoring rule (MSR) trading environment. We then distinguish between two cases. First, all traders are myopic, so that they only care about the current period's payoff. Second, all traders act strategically and care about the future.

3.1 Preferences and updating

Consider a finite state space $\Omega = \{\omega_1, ..., \omega_l\}$ and let the powerset $\mathcal{P}(\Omega)$ be the σ-algebra over Ω. Traders are ambiguity averse and have Maxmin Expected Utility (MEU) preferences (Gilboa and Schmeidler (1989)). In particular, each trader evaluates act $f : \Omega \to \mathbb{R}$ as

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

where $\mathcal{P}$ is a convex and closed subset of $\Delta(\Omega)$, endowed with the weak* topology. We assume that $\mathcal{P}$ is common among all traders and, without loss of generality, $\bigcup_{p \in \mathcal{P}} Supp(p) = \Omega$, 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 $I = \{1, \ldots, n\}$. Trader i's initial private information is represented by partition Π_i of Ω. Without loss of generality, we assume that the join (the coarsest common refinement) of partitions $\Pi = \{\Pi_1, \ldots \Pi_n\}$ consists of singleton sets. This implies that, for any two states $\omega_1 \neq \omega_2$, there exists Trader i such that $\Pi_i(\omega_1) \neq \Pi_i(\omega_2)$, so that the traders' pooled information always reveals the true state.[8]

When a trader learns event E, her beliefs are $\mathcal{P}_E$, the prior by prior updating of $\mathcal{P}$.[9] 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 $\mathcal{E}$.

3.2 Trading environment

Trading is organized as follows. At time $t_0 = 0$, nature selects a state $\omega^* \in \Omega$ and the uninformed market maker makes a prediction y_0 about the value of security $X : \Omega \to \mathbb{R}$. At time $t_1 > t_0$, Trader 1 makes a revised prediction y_1, at $t_2 > t_1$ trader t_2 makes her prediction, and so on. At time $t_{n+1} > t_n$, Trader 1 makes another prediction y_{n+1}. Let $a(t)$ be the trader that makes a prediction at time t. All predictions are observed by all traders. Each prediction y_k is required to be within the set $Y = [\min_{\omega \in \Omega} X(\omega), \max_{\omega \in \Omega} X(\omega)]$.

The process repeats until time $t_\infty = \lim_{k \to \infty} t_k$. At time $t^* > t_\infty$ the true value $x^* = X(\omega^*)$ is revealed. The traders' payoffs are computed using a scoring rule, $s(y, x^*)$, where

[8] This is without loss of generality because, if the conjunction of the traders' private information does not reveal the state, we cannot expect that trading a security will reveal its true value.

[9] This rule is axiomatised by Pires (2002).

x^* is the true 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 maximised 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 $a < \min\limits_{\omega \in \Omega} X(\omega), b > \max\limits_{\omega \in \Omega} X(\omega)$.

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 at t, is $s(y_t, x^*) - s(y_{t-1}, x^*)$, where y_{t-1} is the previous announcement and x^* is the true value of the security. We then say that the trader "buys out" the previous trader's prediction.[10]

We examine trading in two settings. The myopic or non-strategic is analyzed in Section 4, where each trader does not care about future payoffs when making an announcement. We denote this setting by $\Gamma^M(\Omega, I, \Pi, X, \mathcal{P}, y_0, Y, s)$.

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 $\beta^k(s(y_t, x^*) - s(y_{t-1}, x^*))$, where $0 \leq \beta \leq 1$. The total payoff of each trader is the sum of all payments for revisions. Denote this setting by $\Gamma^S(\Omega, I, \Pi, X, \mathcal{P}, y_0, Y, s, \beta)$.

3.3 Properties of scoring rules

In the SEU framework, the optimal (myopic) choice of y_t that maximises $E_p[s(y_t, x^*) - s(y_{t-1}, x^*)]$ does not depend on the previous announcement y_{t-1}, because p is fixed. This is no 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.[11] 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 $Y = [a,b]$, $a, b \in \mathbb{R}$, and $z \in Y$ be an announcement. Then,*

- $y^* \equiv \underset{y \in Y}{argmax}\ \underset{p \in \mathcal{P}}{min} E_p\big[s(y,X) - s(z,X)\big]$ *is unique,*
- $y^* = E_p[X]$ *for some (not necessarily unique)* $p \in \underset{p \in \mathcal{P}}{argmin}\ \underset{y \in Y}{max} E_p\big[s(y,X) - s(z,X)\big]$,
- *if* $z = E_p[X]$ *for some* $p \in \mathcal{P}$, *then* $y^* = z$.

3.4 Information aggregation

We say that information gets aggregated if the traders' predictions converge to the true value of the security, $X(\omega)$.

[10] A trader can guarantee 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.

[11] Lemma 1 is related to a result in Chambers (2008). The proofs are closely related, too.

Definition 1 *Under a profile of strategies in Γ^M or Γ^S, information gets aggregated if sequence $\{y_k\}_{k=1}^{\infty}$ converges in probability to random variable $X(\omega)$, for all $\omega \in \Omega$.*

Since Ω is finite, this is equivalent to requiring that, for any $\epsilon > 0$ and $\delta > 0$, there exists K such that, for any $k > K$, for all states $\omega \in \Omega$, the probability that $|y_k - X(\omega)| > \epsilon$ is less than δ. Note that the finiteness of Ω does not necessarily imply that y_k will deterministically converge to some value, because non-myopic players might use mixed strategies. As we do not model strategic ambiguity, however, each trader mixes with a unique probability.

3.5 Strong separability

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

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 $\omega \in Supp(p)$,

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

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.[12] 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, with ambiguity aversion even separable securities may not aggregate information, as shown in Section 2. The result of no information aggregation does not rely on p^1 assigning 0 to the true state. Example 1 in Appendix C obtains the same result by assuming that all priors have full support.

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

$$d_{\mathcal{P}}(E, v) = \underset{y \in Y}{argmax} \min_{p \in \mathcal{P}_E} E_p\big[s(y, X) - s(v, X)\big]$$

be the (unique from Lemma 1) myopic announcement that maximises 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 we are back to the SEU case, $d_{\mathcal{P}}(E, v) = E_p[X|E]$ for any v and proper scoring rule s. Hence, the myopic announcement $d_{\mathcal{P}}(E, v)$ under ambiguity is a direct

[12] 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 true value of X, which is never 0.

generalization of the myopic announcement under SEU, $E_p[X|E]$. Below, we generalise the notion of separability by substituting $E_p[X|E]$ with $d_{\mathcal{P}}(E, v)$. To save on notation and since security X is fixed throughout the paper, 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 Π_i, $i = 1, ..., n$, and $v \in \mathbb{R}$ such that:*

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

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

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 $E_p[X|\Pi_i(\omega)] = v$ under SEU 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, 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 $d_{\mathcal{P}}(E, v) = \underset{y \in Y}{argmax} \min_{p \in \mathcal{P}_E} E_p\big[s(y, X) - s(v, X)\big]$. This is not the case for separability, which only depends on the information structure. Proposition 2 later in this section establishes that strong separability is also independent of the particular continuous strictly proper scoring rule.

In Section 2, the Arrow-Debreu security is not strongly separable given the information structure and quadratic scoring rule. To see this, note that condition (ii) in the definition is satisfied for all states with $v = 0$. Since some priors put 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. Finally, the class of strongly separable securities is not empty, in general. 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. This is a direct consequence of Proposition 2.

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 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 $\lambda_i : \Pi_i \to \mathbb{R}$ for $i = 1, \ldots, n$ 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, which 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.

Proposition 2 *Security X is strongly separable under partition structure Π if and only if for any $v \in \mathbb{R}$, for any $E \subseteq \{\omega \in \Omega : X(\omega) \neq v\}$, there exists Trader i, state $\omega \in E$ and $\lambda \in \mathbb{R}$ such that for all $\omega' \in \Pi_i(\omega) \cap E$,*

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

The following Lemma shows that there is no security which is strongly separable for all information structures. This is in contrast to the case of separable securities, where, for example, Arrow-Debreu securities are separable for all information structures.

Lemma 2 *If Ω has at least three states, there is no (non-constant) security X which is strongly separable under all partition structures $\Pi = \{\Pi_1, \ldots \Pi_n\}$, where the join of Π consists of singleton sets.*

As we show in subsequent sections (Theorems 1, 2 and 3), strong separability is not only sufficient but also necessary for information aggregation under ambiguity. This suggests a negative result, that there is no security that aggregates information for all information structures. In other words, if the analyst 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 aggregate information 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.

4 Myopic traders

Let $\Gamma^M(\Omega, I, \Pi, X, \mathcal{P}, y_0, Y, s)$ 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. Then, Trader 1 announces her prediction $y_1 \in Y$, where $y_1 \in d_{\mathcal{P}}(\Pi_1(\omega^*), y_0) = arg\max_{y \in Y} \min_{p \in \mathcal{P}_{\Pi_1(\omega^*)}} E_p\big[s(y, X) - s(y_0, X)\big]$. As mentioned above, y_1 depends on the market maker's announcement y_0, which is not the case with SEU.

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 $a(t_k) = i$ at time t_k has private information $F = \mathcal{F}^{k-1}(\omega^*) \cap \Pi_i(\omega^*)$ and announces $y_k = d_{\mathcal{P}}(F, y_{k-1}) = \underset{y \in Y}{argmax} \min_{p \in \mathcal{P}_F} E_p\big[s(y, X) - s(y_{k-1}, X)\big]$.

Let $\mathcal{E} = \{\mathcal{F}^k(\omega) \cap \Pi_{a(t_k)}(\omega)\}_{k \geq 0, \omega \in \Omega}$ 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 terms of strongly separable securities, 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 gets aggregated for any regular $\Gamma^M(\Omega, I, \Pi, X, \mathcal{P}, y_0, Y, s)$ if and only if X is strongly separable.*

To provide some intuition about the result, consider the following Lemma, which is a generalisation of "reaching a consensus", first studied in Geanakoplos and Polemarchakis (1982).

Lemma 3 *Let regular $\Gamma^M(\Omega, I, \Pi, X, \mathcal{P}, y_0, Y, s)$. At any state ω, there exists time t_k such that*

(i) Public information is no longer updated, so that $\mathcal{F}^{k'}(\omega) = \mathcal{F}^k(\omega)$ for every $t_{k'} \geq t_k$,

(ii) No trader i changes her prediction y_i after time t_{k+2n},

(iii) The traders reach an agreement, so that $y = y_i$ for all $i \in I$.

The first result, that the public (and therefore private) information is no longer updated after some time t_k, is a direct consequence of the finiteness of the state space. The second result specifies that, at most two rounds of predictions after t_k, no trader changes her prediction any more. This result is not straightforward, because the myopic prediction depends not only on the private information, as in the SEU case, but also on the previous 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. This does not occur, partly because of a monotonicity property of the scoring rule, adjusted for the case of multiple priors, that the further away the prediction is from the true expected value, the lower is the expectation of the score. The third result states that traders eventually agree on the prediction. Again, this is not straightforward because, with MEU, 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.[13] Nevertheless, separability is indirectly imposed by the scoring rule and the fact that each trader's (constant across partition cells) prediction is tied to another trader's (constant across partition cells) prediction, hence allowing us to derive the result.

5 Strategic traders

Consider a game $\Gamma^S(\Omega, I, \Pi, X, \{\mathcal{P}_i\}_{i\in I}, y_0, Y, s, \beta)$, 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, $Y = [\underline{y}, \overline{y}]$ is the set of possible announcements that a player can make and β is the common discount rate. For simplicity, we assume that each trader has the same set $\mathcal{P}_i = \mathcal{P}$ of priors. Alternatively, we could assume that each i has beliefs $\mathcal{P}_i$ and there is a common prior, so that $\bigcap_{i\in I} \mathcal{P}_i \neq \emptyset$.

Let $H^k = (y_1, \ldots, y_k)$ be a history of announcements up to time t_k and H^0 be the empty history. Given any two histories $H^k = (y_1, ..., y_k)$ and $H^l = (y'_1, ..., y'_l)$, let (H^k, H^l) be their concatenation. An important element of the game is that there is no strategic ambiguity. Although traders have multiple priors over Ω, a mixed strategy consists of randomising using a unique probability distribution. Player i trades at periods t_{i+nk}, $k \in \mathbb{N}$, hence $a(t_{i+nk}) = i$. Her mixed strategy at time t_k is a measurable function $\sigma_{i,k} : \Pi_i \times [\underline{y}, \overline{y}]^{k-1} \times [0, 1] \longrightarrow [\underline{y}, \overline{y}]$. It specifies an announcement y_k, given the element of her partition, the history of announcements $(y_1, \ldots, y_{k-1})$ up to time t_k and the realisation 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 $\phi = (\omega, \iota_1, \iota_2, \ldots)$, describing the initial uncertainty and the randomisations 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 $\sigma = (\sigma_1, \ldots, \sigma_n)$ 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 $y_1(\sigma, \phi), y_2(\sigma, \phi), \ldots$. Let $H^k(\sigma, \phi) = (y_1(\sigma, \phi), \ldots, y_k(\sigma, \phi))$ 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 $H^{k-1+m}(\sigma, \phi|H^{k-1}) = \big(H^{k-1}, y_k(\sigma, \phi|H^{k-1}), \ldots, y_{k-1+m}(\sigma, \phi|H^{k-1})\big)$ the history that is generated by these announcements.

Let $\omega(\phi)$ and $\iota_k(\phi)$ be the first and $k + 1$ components of full state $\phi = (\omega, \iota_1, \ldots)$, respectively. At time t_k, Trader i knows component $\iota_l(\phi)$, which denotes the realisation of the random variable at t_l, if $a(t_l) = i$ and $l \leq k$. Her private information at time t_k and state ϕ is $\Pi_i^k(\phi) = \Pi_i(\omega(\phi)) \times [0, 1]^k \bigcap [\phi' : \iota_l(\phi') = \iota_l(\phi)$ for all $l \leq k$ with $a(t_l) = i]$. 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 $\mathscr{I}_i^k$ be the collection of all information sets for i at time t_k and $\mathscr{I}$ be the collection of all information

[13] Equivalently, dynamic consistency is violated. See Galanis (2019) for a discussion of dynamic consistency in a general framework with multiple beliefs and convex preferences.

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

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

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 $a(t_{k+m}) = i$, 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

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

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 $\mathscr{P} = \{\mathcal{P}(\mathcal{I})\}_{\mathcal{I} \in \mathscr{I}}$ such that each $\mathcal{P}(\mathcal{I})$ is compact and convex.*

To save on notation, we denote the beliefs $\mathcal{P}(\mathcal{I}(H^k, \phi))$ at an information set as $\mathcal{P}(H^k, \phi)$.

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

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

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

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 denote V_i with V, as it is clear in each time t_k who is making the announcement. There is one exception at period 0, after each player has received her private information but before player 1 has made the first announcement, so the history H^0 is empty. In that case, we denote i's ex-ante payoff as $V_i(H^0, \phi, \sigma, \mathscr{P})$.

5.1 Ex-ante and interim equilibria

One of the main issues 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. Ostrovsky (2012) shows that with SEU preferences and for any Nash equilibrium, separable securities characterize information aggregation. Due to dynamic inconsistency, a similar result with MEU preferences is not true. However, we show that strongly separable securities characterize information aggregation with interim equilibria. In an interim equilibrium, the strategy profile is optimal at each time t_k, when everyone plays on-path and there is prior by prior updating. We are able to provide a stronger result, as we only require that optimality holds for each $t_k \geq t_{k'}$, for some k'. We call this *an interim equilibrium at the limit*. We also argue that if we impose dynamic consistency, for example using rectangular priors (Epstein and Schneider (2003)), then an ex-ante equilibrium is also interim. In such a case, strongly separable securities characterize information aggregation in an ex-ante equilibrium, thus generalizing the result of Ostrovsky (2012).

Definition 6 *A strategy profile σ^* is an ex-ante equilibrium if $V_i(H^0, \phi, \sigma^*, \mathscr{P}) \geq V_i(H^0, \phi, \sigma_i', \sigma_{-i}^*, \mathscr{P})$, for all $\phi \in \Phi, \sigma_i'$ and $i \in I$.*

Before defining the notion of an interim equilibrium, we specify that pair $(\sigma, \mathscr{P})$ is consistent if there is prior by prior updating on-path, given the information that is revealed in each period.

Definition 7 *Pair $(\sigma, \mathscr{P})$ is consistent if, for any $\phi \in \Phi$, $k \geq 0$ and player $a(t_k) = i$,*

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

(ii) $\mathcal{P}(H^{k+n}(\phi, \sigma), \phi)$ is the prior by prior updating of $\mathcal{P}(H^k(\phi, \sigma), \phi)$ given F.[14]

Consistent pair $(\sigma, \mathscr{P})$ is an interim equilibrium at the limit if, after some time $t_{k'}$, σ is optimal at each subsequent period. It is an interim equilibrium if $t_{k'}$ is the initial period. Note that we require that at t_k, player $a(t_k)$ best responds given the equilibrium strategy σ^*.

Definition 8 *Consistent pair $(\sigma, \mathscr{P})$ is an interim equilibrium at the limit if, for some $k' \geq 0$ and all $k \geq k'$, $V(H^k(\sigma, \phi), \phi, \sigma^*, \mathscr{P}) \geq V(H^k(\sigma, \phi), \phi, \sigma'_{a(t_k)}, \sigma^*_{-a(t_k)}, \mathscr{P})$, for all $\sigma'_{a(t_k)}$ and ϕ. It is an interim equilibrium if $k' = 0$.*

Our first result is that if X is strongly separable, then for any interim equilibrium at the limit, there is information aggregation. Conversely, if X is not strongly separable, then for some interim equilibrium at the limit there is no information aggregation.

Theorem 2 *Fix information structure Π and bounds Y.*

(i) If security X is strongly separable under Π, then for any Γ^S and any interim equilibrium at the limit, information gets aggregated.

(ii) If security X is not strongly separable under Π, then there exist game Γ^S and interim equilibrium at the limit, such that information does not get aggregated.

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

We now present an informal discussion on the issue of existence of equilibrium. As with Ostrovsky (2012), standard existence results do not apply, because the action spaces are infinite. He resolves this problem by considering a discrete version of the game, so that each player can only choose from a finite set of predictions and then shows that, as the grid becomes sufficiently fine, information gets approximately aggregated. An additional problem we encounter in our setting is that we employ an interim, instead of an ex-ante, equilibrium, due to dynamic inconsistency. However, we argue that in the special case where beliefs are rectangular and therefore there is dynamic consistency, an ex-ante equilibrium is also interim. We could then extend the result of Pahlke (2018), who shows that there always exists an ex-ante equilibrium with rectangular beliefs, in a setting with finite actions and finitely many periods.

Rectangularity, examined by Epstein and Schneider (2003), is a generalisation of the law of iterated expectations, which specifies that a (full support) prior can be decomposed by the marginals and the Bayesian updates given a partition Π of Ω, so that $p(\omega) = \sum_{E \in \Pi} p(E) p_E(\omega)$ for all $\omega \in \Omega$, where $p_E = \frac{p(\cdot)}{p(E)}$ is the Bayesian update of p given E.

Let $\mathcal{P}$ be a set of priors and Π be a partition of Ω. For each partition element $E \in \Pi$, let $p_E = \frac{p(\cdot)}{p(E)}$, $p \in \mathcal{P}$, be the Bayesian update of some prior $p \in \mathcal{P}$, given E. We say that $\mathcal{P}$ is rectangular with respect to Π if the following condition holds. Consider any collection of Bayesian updates p_E, one for each $E \in \Pi$, noting that p_E and $p_{E'}$ may not be the Bayesian updates of the same prior $p' \in \mathcal{P}$. Then, for any prior $p' \in \mathcal{P}$, the reconstructed prior $p(\cdot) = \sum_{E \in \Pi} p'(E) p_E(\cdot)$ is also an element of $\mathcal{P}$. Although we define rectangularity in the case of finite state space Ω, this is easily extended in the case of Φ.

Epstein and Schneider (2003) show that rectangular priors and prior by prior updating imply dynamic consistency.[15] Suppose that σ^* is an ex-ante equilibrium and beliefs $\mathscr{P}$ are such that each $\mathcal{P}_i$ is rectangular given the sequence of partitions $\{\mathcal{F}_i^k(\sigma, \phi) : \phi \in \Phi\}$, that are generated by the revealing of information, at each time t_k where i makes an announcement. Then, dynamic consistency implies that i plays a best response at each time t_k, so that $(\sigma, \mathscr{P})$ is also an interim equilibrium. We therefore have the following result.[16]

Corollary 1 *Fix information structure Π and bounds Y. Consider a consistent pair $(\sigma^*, \mathscr{P})$ and suppose that each $\mathcal{P}_i$ is rectangular with respect to i's partition $\{\mathcal{F}_i^k(\sigma, \phi) : \phi \in \Phi\}$, for each t_k where i makes an announcement.*

- (i) *If security X is strongly separable under Π and σ^* is an ex-ante equilibrium, then information gets aggregated.*
- (ii) *If security X is not strongly separable under Π, then there exist game Γ^S and an ex-ante equilibrium such that information does not get aggregated.*

[15] Other rules that imply dynamic consistency are updating only a subset of priors, as in Hanany and Klibanoff (2007) for MEU preferences, a generalisation for uncertainty averse preferences in Hanany and Klibanoff (2009), including the smooth rule for the smooth model of Klibanoff et al. (2005), and the Bayesian updating of subjective beliefs in Galanis (2019) for general convex preferences.

[16] We do not prove this Corollary formally, as (i) is straightforward, whereas the proof of part (ii) is the same as the proof for Theorem 2.

Pahlke (2018) shows that an ex-ante equilibrium is also an interim equilibrium with rectangular priors, but her result can be extended in our setting with infinite periods, as the proof is inductive.[17] She also shows that an ex-ante equilibrium with rectangular beliefs always exists. 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 SEU preferences. We avoid such a criticism in the case of rectangular priors, because the proof of Theorem 2 allows for different priors $\mathcal{P}_i$, with a non-empty intersection. Bade (2016) argues that the behavior of dynamically consistent players, who follow an ex ante optimal plan, cannot be distinguished from the behavior of players with SEU preferences.[18] However, our result of information aggregation only requires an interim equilibrium at the limit, not an ex ante equilibrium.

5.2 Revision-proof equilibrium

The interim equilibrium at the limit imposes dynamic consistency "eventually". However, another 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. 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 generalisation would be to check for one-shot deviations.

Before defining the notion of consistent planning, we strengthen Definition 8 to off-path consistency, by additionally imposing prior by prior updating at all decision nodes, whenever possible.[19]

Definition 9 *Tuple $(\sigma, \mathscr{P})$ is consistent off-path if, for any full state $\phi \in \Phi$, history H^k, $k \geq 0$ and player $a(t_k) = i$,*

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

(ii) If $\bigcup_{p \in \mathcal{P}(H^k, \phi)} Supp(p) \bigcap F \neq \emptyset$, then $\mathcal{P}(H^{k+n}(\sigma, \phi | H^k), \phi)$ is the prior by prior updating of $\mathcal{P}(H^k, \phi)$ given F.[20]

At decision node (H^k, ϕ), the beliefs of player $a(t_k) = i$ 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 $F = \mathcal{F}_i^{k+n}(\sigma, \phi | H^k)$. 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.

[17] Using the smooth rule and the smooth model, Hanany et al. (2018) show the existence of a Sequential Equilibrium with Ambiguity, in a setting with finite actions and periods.

[18] She proposes the notion of semi-consistency, which allows for different sets of equilibria and predictions.

[19] Off-path 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).

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

Definition 10 *Consistent off-path tuple* $(\sigma^*, \mathscr{P})$ *is a consistent-planning equilibrium if there is no decision node* (H^{k-1}, ϕ)*, player* $a(t_k) = i$ *and alternative strategy* $\sigma = (\sigma_i, \sigma^*_{-i})$*, with* $\sigma_{i,k'} = \sigma^*_{i,k'}$ *for all* $k' \neq k$*, such that*

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

This 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.[21]

A consistent off-path tuple $(\sigma^*, \mathscr{P})$ 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}. Note that we only check initial deviations from each on-path decision node $(H^{k-1}(\phi, \sigma^*), \phi)$, not from any history.

Definition 11 *Consistent off-path pair* $(\sigma^*, \mathscr{P})$ *is a revision-proof equilibrium if there is no decision node* $(H^{k-1}(\phi, \sigma^*), \phi)$*, player* $a(t_k) = i$ *and alternative strategy* $\sigma = (\sigma_i, \sigma^*_{-i})$ *such that for all* $r \geq 0$ *and* H^{nr}*,*

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

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.

Our second main result in a strategic environment shows that strongly separable securities aggregate information in all revision-proof equilibria.

Theorem 3 *Fix information structure* Π *and bounds* Y*.*

(*i*) *If security* X *is strongly separable under* Π*, then for any* Γ^S *and any revision-proof equilibrium, information gets aggregated.*

(*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 get aggregated.*

[21] Note that, as is the case with complete information games, revision-proof equilibria may not always exist.

6 Experiment

Our experimental design focuses on three dimensions. The first is whether beliefs about events are precise or not. In particular, we either inform subjects about the exact composition of the urn, so that they formulate precise beliefs and have Subjective Expected Utility (SEU) preferences, or we provide partial information about the composition, so that they formulate multiple priors that give rise to Maxmin Expected Utilty (MEU) preferences. The second dimension relates to the type of security that is traded: separable securities, such as an Arrow-Debreu security, or strongly separable securities.

The third dimension relates to the initial value or price of the security, which is provided by the uninformed market maker. We allow for two initial values: 0 and 50 (more details on the choices are provided below). We therefore examine the impact on information aggregation when the market maker announces an initial value of 0, and when he announces an initial value of 50. In summary, we apply a $2 \times 2 \times 2$ experimental design to examine the impact on information aggregation of the market type, security type and initial value.

6.1 Experimental design

Our theory generates several testable implications, which we test in the lab. We also compare the information aggregation properties of separable and strongly separable securities. In this section we describe our experimental design.

Initially, subjects were endowed with 6,000 Experimental Currency Units (ECUs) as a show-up fee. The conversion was 2,000 ECUs for €1. There were 3 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 value. 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.

In the game play stage, subjects were recruited to play the role of traders forecasting the value of stock, which is either high or low. The stock value was randomly determined, based on the color of a ball drawn from an urn with 90 colored balls. The information on the color composition of the urn as well as the mapping of colors to high or low stock values were treated variables, reflecting the market type (unique versus multiple priors) and security type (separable versus strongly separable), respectively. In addition, before making a decision, subjects would receive a signal about the color of the drawn ball. The information about the composition of the urn, the structure of the signals, the payoff functions as well as the initial value (also a treated variable) were explicitly explained in the experimental instructions. The experimental instructions are reported in the Appendix.

To ensure that the subjects understood the environment, before the actual game play, they had to complete a 15-question quiz. After the quiz, subjects were asked to take part in 12 rounds of prediction markets. In each round, traders made sequential predictions about the value of the security (this was an integer from 0 to 100). 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.[22] Whether or not there would be an extra period was thus determined by a random draw. Furthermore, trading pairs were fixed for the duration of the round, but new pairs were formed in every new round. This information was also common knowledge.

At the beginning of each round, traders were given an endowment of 1,500 ECUs (recall the conversion was 2,000 ECUs for €1). Payoffs were calculated based on the Market Scoring Rule (MSR) *at the end of each trading period*. Thus, the trader's payoff was a function of (a) the stock value (high or low), (b) the trader's own prediction, and (c) the previous trader's reported prediction. When the value of the stock was high, the trader's payoff was given by the formula:

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

When the value of the stock was low, the trader's payoff was calculated by the formula:

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

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 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.[23] 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 payoffs were the summation of all the round payoffs of the trader in the 12 rounds played. We describe next the treatments.

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 Subjective Expected Utility (SEU) preferences. Henceforth, this market is referred to as *SEU*.

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 mirrors the example of 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

[22]This assumption is similar to that made in Roth and Murnighan (1978), Fréchette and Yuksel (2017), Cabral et al. (2014) and Vespa (2011). It is necessary in order to simulate the infinitely-many-periods assumption of the theoretical setting and avoiding subjects implementing backward induction reasoning.

[23]In the actual experiments, no subject lost the entire endowment given in the beginning of the round.

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 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 given that red is revealed, and therefore all beliefs must assign 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 Maxmin Expected Utilty (MEU) preferences. Henceforth, this market is referred to as *Amb*.

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). In the case of strongly separable securities, we informed subjects that if the red ball 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).

The initial value of the security was also a treated variable. The two security types exhibit the same information aggregation, in every single state, for all initial values with the exception of 0. At the 0 initial value, the information aggregation 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.[24] We thus chose to investigate experimentally information aggregation at the initial value of 0, as well as at an initial value where the two security types perform the same. We chose 50, as the midpoint between 0 and 100.

Finally, all treatments included identical information structure about traders' signals. The structure was common knowledge and was presented to subjects in a tabular form as shown in Table 1, though the signal was private. This information was also explicitly discussed in the instructions. For instance, subjects were told that "if the drawn ball is red (hence the value of the stock is high), 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 colors as well.

Table 1: Information Structure

	Private Information	
Ball Drawn	Trader 1	Trader 2
Red	Not Blue	Not Green
Green	Not Blue	Green
Blue	Blue	Not Green

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

[24]Note that the failure of information aggregation in the *Amb* setting with a separable security is special to the particular example we use. In general, information aggregation can fail at multiple initial values.

Recall that there were 12 rounds of predictions where the probability of having another trading period within a given round was 95%. The draws for the number of trading periods within each round was done *ex ante* to ensure that all treatments had 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 {*Red,Blue,Blue,Blue,Red,Blue,Red,Green,Red,Green,Blue,Blue*}.[25] The realized color 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 color could reflect a low stock value (in the case of separable securities) or a high stock value (in the case of strongly separable securities).

The experimental sessions took place in February of 2019 at the Laboratoire d'Économie Expérimentale de Paris (LEEP). We conducted two sessions per treatment. The 288 subjects were recruited from the database of the Université Paris 1 Panthéon - Sorbonne. We sent emails publicizing the experiment, and 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 2. Note that each treatment is denoted by an acronym. In particular, the acronym *(market type, security type, initial value)* consists of the *market type* (*SEU* for the market with SEU 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 value* (0 or 50).

6.2 Theoretical predictions

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

To measure the degree of information aggregation in a market, we use the true value of the security as a benchmark. This is the most natural candidate to serve as a benchmark for several reasons. First, by construction, the true 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 SEU preferences, the predictions of Bayesian traders always converge to the true value for separable securities. The same holds true in environments with

[25] The respective signals (i, j), where i is the signal of Trader 1 and j is the signal of Trader 2, were {*(Not Blue, Not Green),(Blue, Not Green),(Blue, Not Green),(Blue, Not Green),(Not Blue, Not Green),(Blue, Not Green),(Not Blue, Not Green),(Not Blue, Green),(Not Blue, Not Green),(Not Blue, Green),(Blue, Not Green),(Blue, Not Green)*}.

Table 2: Characteristics of the Experimental Sessions

# of Subj.	# of Ses.	Market Type	Security Type	Acronym
Initial Value is 0				
36	2	SEU	Separable	SEUS0
36	2	Amb	Separable	AmbS0
36	2	SEU	Str. Separable	SEUStS0
36	2	Amb	Str. Separable	AmbStS0
Initial Value is 50				
# of Subj.	# of Ses.	Market Type	Security Type	Acronym
36	2	SEU	Separable	SEUS50
36	2	Amb	Separable	AmbS50
36	2	SEU	Str. Separable	SEUStS50
36	2	Amb	Str. Separable	AmbStS50

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 value. The acronyms in the last column consist of the *market type* (*SEU* for the market with SEU 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 value* (0 or 50).

MEU preferences and strongly separable securities (Theorems 1, 2 and 3). We therefore use the true value of the security as our baseline and measure its absolute distance from the *final prediction*.[26] We call this measure, for brevity, AD (absolute difference). We say that the information aggregation in market B is not as good as that in market A, if the AD in market B is significantly larger than that in market A.

We now formulate our conjectures, which refer to separable securities in an environment with ambiguity, and theoretical predictions, which refer to the testable implications of our theory of strongly separable securities. Conjecture 1 interprets the main result of Ostrovsky (2012) in an environment with ambiguity aversion. It specifies that separable securities aggregate information irrespective of whether traders have a unique belief and SEU, or imprecise beliefs and ambiguity aversion.

[26] Our criterion is one of many. For example, we could have used the last predictions of both traders, instead of the final prediction, in our distance measure. The results are almost identical. However, to maintain consistency between the theory and the statistical analysis, we chose to measure the distance between the true value of the security and the final trader's prediction. Another example would be the Euclidean distance. Our results are not sensitive to similar distance measures. For our purposes, the simplest suffices.

Conjecture 1 *Assuming an initial value of 0 and separable securities, information aggregation across the SEU and Amb markets is the same, regardless of the color of the drawn ball.*

Prediction 2 is a direct implication of Theorems 1, 2 and 3, which show that strongly separable securities always aggregate information, in both *SEU* and *Amb* markets, with myopic or strategic traders.

Prediction 2 *Assuming an initial value of 0 and strongly separable securities, information aggregation across the SEU and Amb markets is the same, regardless of the color of the drawn ball.*

We now test whether the market maker's initial announcement has an impact on the degree of information aggregation. Ostrovsky (2012) specifies that the initial price does not influence the information aggregation in an environment with SEU. Since his theory does not extend to an environment with ambiguity aversion, we again state this as a conjecture. In the case of strongly separable securities, we state this as a testable prediction. In both cases, we examine whether information aggregation is influenced when the initial value is 50.

Conjecture 3 *For separable securities and assuming an initial value of 50, information aggregation across the SEU and Amb markets is the same regardless of the color of the drawn ball.*

Prediction 4 *For strongly separable securities and assuming an initial value of 50, information aggregation across the SEU and Amb markets is the same regardless of the color of the drawn ball.*

Finally, we test whether, holding the *Amb* market constant, changing the initial value from 0 to 50 has any impact on the degree of information aggregation.

Conjecture 5 *In the* Amb *market with separable securities and for any color of the drawn ball, the information aggregation under an initial value of 0 is not significantly different than under an initial value of 50.*

Prediction 6 *In the* Amb *market with strongly separable securities and for any color of the drawn ball, the information aggregation under an initial value of 0 is not significantly different than under an initial value of 50.*

6.3 Results

Each hypothesis is matched with the corresponding result; that is, result i is a report on the test of conjecture or prediction i.

6.3.1 Descriptive statistics

In this section, we report some descriptive statistics about the absolute difference (AD) in distance of the trader's final prediction from the true 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.

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 SEUS0, 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, SEUStS0, 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 color of the drawn ball to at least one trader, in contrast to the green and blue states. For better visualisation of the results, we display in Figure 2 the box plots of the ADs across the market types when the initial value is 0, and in Figure 3, we display the box plots when the initial value is 50. It is evident from the box plots that there was a lot of variability in the reports of the subjects. Again, this is not surprising as the nature of the game allows for strategic behavior, which results in noisier predictions.

6.3.2 Information aggregation

In this section, we perform statistical analysis to investigate the impact on information aggregation of the treated variables. For the analysis, we use the Mann-Whitney test, where the H_0 states that the AD in the SEU market is greater or equal to the AD in the Amb market when fixing the realized state. The p-values are displayed in Table 3.

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

Result 1 *For an initial value of 0, information aggregation across the SEU and Amb markets for separable securities is the same when the drawn balls are green or blue. When the drawn ball is red, information aggregation in the Amb market is not as good.*

Support. Contrary to our conjecture, we find that in the red state, information aggregation in the *Amb* market is not as good (p-value is 0.001) as that in the *SEU* market. Therefore, the H_0 can be rejected at the conventional 5% level of statistical significance.

Figure 2: Box Plots for Initial Value of 0

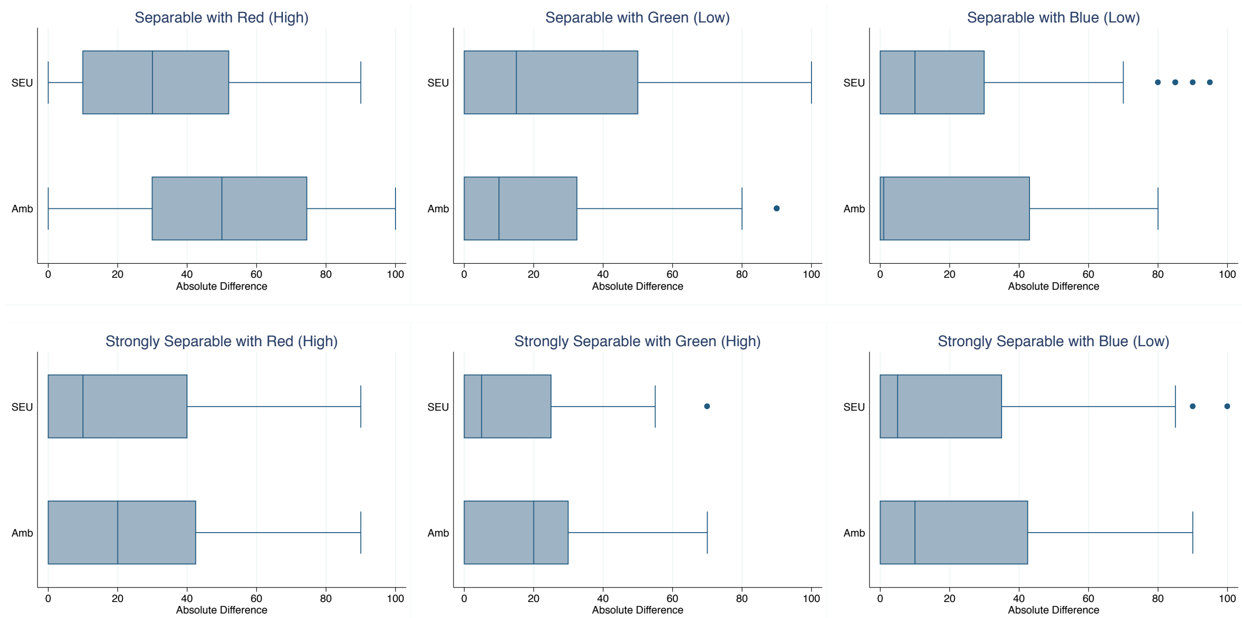

Notes: We display the box plots of the absolute difference across the market types conditional on the realized state (red, green, blue), when the initial value is 0.

The second prediction also dealt with an initial value of 0, though this time, the information aggregation in the strongly separable securities is investigated. Our second result sheds light to the strength of the strong separability condition.

Result 2 *For an initial value of 0, information aggregation across the SEU and Amb markets for strongly separable securities is the same regardless of the color of the drawn ball.*

Support. The prediction is confirmed in the red, green and blue states, where the p-values are 0.107, 0.140 and 0.195, respectively.

Conjecture 3 and Prediction 4 test the effect on information aggregation of the market maker's announcement of the focal value of 50, for separable and strongly separable securities, respectively.

Result 3 *For an initial value of 50 and given a separable security, information aggregation across the SEU and Amb markets is the same regardless of the color of the drawn ball.*

Support. The conjecture is confirmed for all three states (p-values are 0.392, 0.342 and 0.265 in the red, green and blue states, respectively).

Result 4 *For an initial value of 50 and given a strongly separable security, information aggregation across the SEU and Amb markets is the same regardless of the color of the drawn ball.*

Figure 3: Box Plots for Initial Value of 50

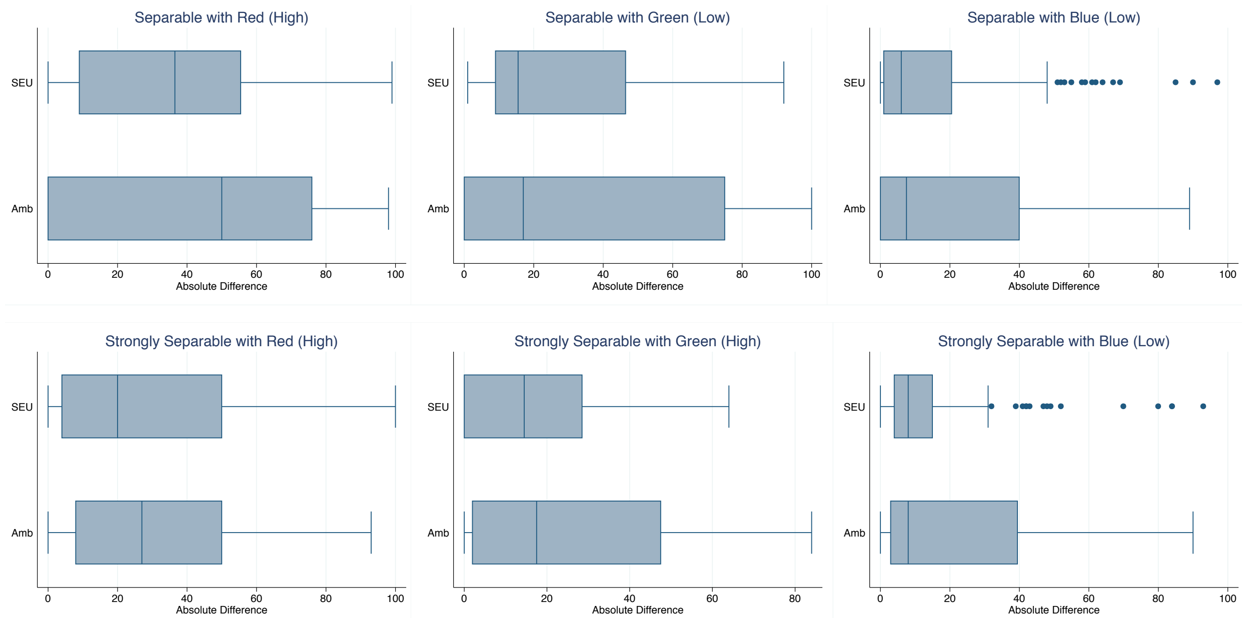

Notes: We display the box plots of the absolute difference conditional on the realized state (red, green, blue) when the initial value is 50.

Support. The prediction is confirmed for all three states (p-values are 0.316, 0.168 and 0.262 in the red, green and blue states, respectively). Therefore, our results, here, are entirely in line with the theoretical prediction in the strongly separable securities.

Conjecture 5 and Prediction 6 test the degree of information aggregation in an environment with ambiguity, when the initial value changes from 0 to 50. For the analysis, we again use the Mann-Whitney test, where the H_0 states that the AD in the Amb market is the same across 0 and 50 initial values when fixing the realized state.

Result 5 *In the* Amb *market with a separable security, information aggregation across the 0 and 50 initial values is the same in the red and green states, but not in the blue state.*

Support. None of the p-values is statistically significant in the red and green states (the p-values are 0.143 and 0.195, 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.

Result 6 *In the* Amb *market with a strongly separable security, information aggregation across the 0 and 50 initial values is the same, regardless of the color of the drawn ball.*

Support. The prediction is confirmed all three states (red, green and blue states, where the p-values are 0.111, 0.184 and 0.231, respectively). Therefore, our results, here, are again entirely in line with the theoretical prediction in the strongly separable securities.

Table 3: Mann-Whitney Tests on Information Aggregation

Panel A	Initial Value is 0	
	Separable	Strongly Separable
Alternative hypothesis:	$AD_i < AD_j$	
	p-values	
Red State		
SEU vs. Amb	**0.001**	0.107
Green State		
SEU vs. Amb	0.479	0.140
Blue State		
SEU vs. Amb	0.447	0.195
Panel B	Initial Value is 50	
	Separable	Strongly Separable
Alternative hypothesis:	$AD_i < AD_j$	
	p-values	
Red State		
SEU vs. Amb	0.392	0.316
Green State		
SEU vs. Amb	0.342	0.168
Blue State		
SEU vs. Amb	0.265	0.262

Notes: We utilize the Mann-Whitney tests to determine whether the absolute difference (AD) in the SEU market is greater or equal to the 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 value is 0. In Panel B, we report the p-values of the comparisons in the ADs when the initial value is 50.

7 Concluding remarks

The main purpose of the paper is to study the information aggregation properties of markets, and in particular prediction markets with ambiguity averse traders. We find that separable securities, which aggregate information in environments with precise probabilities and SEU, are no longer sufficient when probabilities are imprecise. We confirm this result also in an experimental setting with subjects trading in a prediction market. 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. We provide several testable implications of our theory, which we are able to confirm in the lab. However, we also show that there is no security that is strongly separable for all information structures. This is a negative result, because strong separability is necessary for information aggregation, hence we find that there is no security that aggregates information for all information structures. This is not only a negative result for the ability of prediction markets to aggregate information with ambiguity, but of financial markets in general.

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 $s(y)(.) \equiv s(y, X(.))$. We first show that $arg\max_{y\in Y}\min_{p\in\mathcal{P}} E_p\big[s(y) - s(y_{-1})\big]$ does, in fact, exist. This is true because s is continuous function, therefore $\min_{p\in\mathcal{P}} E_p\big[s(y) - s(y_{-1})\big]$ is upper semi continuous (as infimum of continuous functions) as a function of y. Since Y is compact, a maximum exists and the set $arg\max_{y\in Y}\min_{p\in\mathcal{P}} E_p\big[s(y) - s(y_{-1})\big]$ is not empty.

Next, we define Z to be the convex hull of $\{s(y)\}_{y\in Y}$. The set $\{s(y)\}_{y\in Y}$ is compact in $\mathbb{R}^l$ because s is continuous in y and Y is compact, hence Z is compact. Consider the function $G : \mathcal{P} \times Z \longrightarrow \mathbb{R}$ defined by $G(p, z) = E_p[z - s(y_{-1})]$. The function is linear in p and affine in z. Moreover, it is continuous both in p and in z. The first is because of the definition of weak* convergence and the second applying Lebesgue's dominated convergence theorem.

By Sion's minimax Theorem (Berge (1963), p. 210), there exists $p^* \in \mathcal{P}$ and $z^* \in Z$ such that for all $(p, z) \in \mathcal{P} \times Z$ it is $E_{p^*}\big[z - s(y_{-1})\big] \le E_{p^*}\big[z^* - s(y_{-1})\big] \le E_p\big[z^* - s(y_{-1})\big]$. Then we get that $\min_{p\in\mathcal{P}}\max_{z\in Z} E_p\big[z - s(y_{-1})\big] = \max_{z\in Z}\min_{p\in\mathcal{P}} E_p\big[z - s(y_{-1})\big]$ and it is achieved at $p = p^*$, $z = z^*$.

For a fixed p, and because $G(p, z)$ is affine in z, the unique maximiser of $E_p\big[z-s(y_{-1})\big]$ over z is $s(E_p[X])$ (since s is a proper scoring rule, by definition of Z), so that $z^* = s(E_{p^*}[X])$. Hence we may conclude $\min_{p\in\mathcal{P}}\max_{y\in Y} E_p\big[s(y) - s(y_{-1})\big] = \max_{y\in Y}\min_{p\in\mathcal{P}} E_p\big[s(y) - s(y_{-1})\big]$ and it is achieved at $p = p^*$, $y = E_{p^*}[X]$.

We claim that $y = E_{p^*}[X]$ is a unique element of $arg\max_{y\in Y}\min_{p\in\mathcal{P}} E_p\big[s(y, X(\omega))-s(y_{-1}, X(\omega))\big]$. To see that, let $y' \neq E_{p^*}[X]$. Then,

$$\min_{p\in\mathcal{P}} E_p\big[s(y', X(\omega)) - s(y_{-1}, X(\omega))\big] \le E_{p^*}\big[s(y', X(\omega)) - s(y_{-1}, X(\omega))big\big] <$$

$$E_{p^*}\big[s(E_{p^*}[X], X(\omega)) - s(y_{-1}, X(\omega))\big] = \max_{y\in Y}\min_{p\in\mathcal{P}} E_p\big[s(y, X(\omega)) - s(y_{-1}, X(\omega))\big].$$

Hence, the maximiser is unique.

For the third claim, note that $E_p\big[s(z,X)-s(z,X)\big] = 0$ for all $p \in \mathcal{P}$, hence $\underset{y\in Y}{max}\ \underset{p\in\mathcal{P}}{min} E_p\big[s(y,X)-s(z,X)\big] \geq 0$. Because $z = E_p[X]$ for some $p \in \mathcal{P}$, we have that $p \in \underset{p\in\mathcal{P}}{argmin}\ \underset{y\in Y}{max} E_p\big[s(y,X)-s(z,X)\big]$ and $y^* = z$.

■

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 $\omega \in \bigcup_{p\in\mathcal{P}} Supp(p) = E$, for each $i \in I$, we have $E_p[X(\omega) - v|\Pi_i(\omega)] = 0$, for some $p \in \mathcal{P}$, ignoring without loss of generality states ω' for which $X(\omega') = v$. Because $Supp(p) \subseteq E$, 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 $Supp(p'') = E$ such that $E_{p''}[X(\omega) - v|\Pi_i(\omega)] = 0$. 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.

■

Proof of Lemma 2. 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.

■

Proof of Lemma 3.

For (i), by construction, $\mathcal{F}^0(\omega) \supseteq \mathcal{F}^1(\omega) \supseteq ... \supseteq \mathcal{F}^k(\omega)$. Because Ω is finite, there exists t_k such that $\mathcal{F}^{k'}(\omega) = \mathcal{F}^k(\omega)$ for every $t_{k'} \geq t_k$.

For (ii), observe that the function $\Phi(p) = E_p[s(E_p[X], X) - s(z, X)]$, with $p \in \Delta(\Omega)$ is convex in p, for any $z \in [a, b]$, with $a = \min\{E_p[X] : p \in \Delta(\Omega)\}$ and $b = \max\{E_p[X] : p \in \Delta(\Omega)\}$. Define the function $g(E_p[X]) = \Phi(p)$. Note that g is convex in $\{E_p[X] : p \in \Delta(\Omega)\}$ and because its unique minimiser is at z we get that g is decreasing at $[a, z]$ and increasing at $[z, b]$.[27] From Lemma 1, the myopic announcement of Trader i at time t_k, when the previous announcement is z, is given by $d_{\mathcal{P}}(\mathcal{F}^{k-1}(\omega) \cap \Pi_i(\omega), z) = E_{p^*}[X]$ for some $p^* \in \mathcal{P}_{\mathcal{F}^{k-1}(\omega)\cap\Pi_i(\omega)}$, hence $d_{\mathcal{P}}(\mathcal{F}^{k-1}(\omega) \cap \Pi_i(\omega), z) = \underset{x\in\{E_p[X]:p\in\mathcal{P}_{\mathcal{F}^{k-1}(\omega)\cap(\Pi_i(\omega))}\}}{\arg\min} g(x)$.

If z (the unique minimiser of $E_p[X]$ for all $p \in \Delta(\Omega)$) is on the left hand side of $A = \{E_p[X] : p \in \mathcal{P}_{\mathcal{F}^{k-1}(\omega)\cap\Pi_i(\omega)}\}$, then the left hand side extreme point of A is the minimising value, and similarly if z is on the right hand side of A. This is due to the convexity of Φ, therefore of g, and the fact that z is the global minimum.

Define $A^i_{\omega'} = \{E_p[X|\Pi_i(\omega')] : p \in \mathcal{P}_{\mathcal{F}^k(\omega)}\}$ for every $i = 1, ..., n$ and $\omega' \in \mathcal{F}^k(\omega) = \{\omega_1, \ldots, \omega_l\}$.

Step 1: If there is no information revelation after t_k, $A^i = \bigcap_{\omega'\in\mathcal{F}^k(\omega)} A^i_{\omega'} \neq \emptyset$ for every $i = 1, ..., n$.

If $A^i = \emptyset$, then $A^i_{\omega'} \cap A^i_{\omega''} = \emptyset$, for two states $\omega', \omega'' \in \mathcal{F}^k(\omega)$. The second property of Lemma 1 shows that at ω', the trader can only make an announcement in $A^i_{\omega'}$, and similarly for ω''. Since $A^i_{\omega'} \cap A^i_{\omega''} = \emptyset$, either ω' or ω'' is revealed not to be the true state, hence there is further information revelation, a contradiction.

Step 2: If $\bigcap_{j\in\{1,...,n\}} A^j = \emptyset$, then no trader changes her prediction after t_{k+2n}.

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

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 $\bigcap_{j\in\{1,...,i_0-1\}} A^j$. For the same reason, any subsequent announcement up to $i_0 - 1$ also belongs to $\bigcap_{j\in\{1,...,i_0-1\}} A^j$. 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_{i_0}.

[27]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.

For the next round, the announcement of the last trader potentially triggers different announcements for traders $j = 1, ..., i_0 - 1$. However, the same argument as before shows that the announcement of $i_0 - 1$ belongs to $\bigcap_{j \in \{1,...,i_0-1\}} A^j$ and i_0 announces v_{i_0}. Hence, the announcements of all subsequent traders after i_0 remain the same, implying that for the next round, the announcements of every trader do not change anymore.

Step 3: If $\bigcap_{j \in \{1,...,n\}} A^j \neq \emptyset$, then no trader changes her prediction after t_{k+2n}.

There are A^{i_1} and A^{i_2} such that one defines the left hand side extreme point of the interval and the other defines the right hand side extreme point. Using similar arguments as before, for any y_{k-1}, trader $i_3 = \max\{i_1, i_2\}$ gives a prediction belonging in the set $\bigcap_{j \in \{1,...,n\}} A^j$. We denote the corresponding announcement with v_{i_3}. From the second and third properties of Lemma 1, we conclude that for $j = i_3, ..., n$ their announcements are v_{i_3}. Because $v_{i_3} \in \bigcap_{j \in \{1,...,n\}} A^j$ we conclude that at the next round the announcement of each trader $1, ..., i_3 - 1$ is v_{i_3}, too. Hence we get $v_1 = ... = v_n = v_{i_3}$.

For (iii), denote, for simplicity, $\mathcal{F}^k(\omega) = \mathcal{F}$. Let v_i be trader i's permanent (from (ii)) prediction. Since v_i is the myopically optimal answer and i can always get 0 by repeating the previous announcement, we have that $\min_{p \in \mathcal{P}} E_{p|\mathcal{F} \cap \Pi_i(\omega')} \big[s(v_i, X) - s(v_{i-1}, X)\big] \geq 0$ for all $\omega' \in \mathcal{F}$ and $i \in I$.[28] This implies that $p(\mathcal{F} \cap \Pi_i(\omega')) E_{p|\mathcal{F} \cap \Pi_i(\omega')} \big[s(v_i, X) - s(v_{i-1}, X)\big] \geq 0$ and $p(\mathcal{F} \cap \Pi_i(\omega')) > 0$, for every $\omega' \in \mathcal{F}$ and $p \in \mathcal{P}$.[29] Summing over $\mathcal{C}_i = \{\Pi_i(\omega) : \omega \in \mathcal{F}\}$ we get $p(\mathcal{F}) E_{p|\mathcal{F}} \big[s(v_i, X) - s(v_{i-1}, X)\big] \geq 0$. By summing over i and ignoring $p(\mathcal{F})$, we have

$$E_{p|\mathcal{F}}\big[s(v_1, X) - s(v_n, X)\big] + E_{p|\mathcal{F}}\big[s(v_2, X) - s(v_1, X)\big] + \ldots +$$

$$+ E_{p|\mathcal{F}}\big[s(v_n, X) - s(v_{n-1}, X)\big] = 0.$$

For all $i \in I$, because each term is non negative, we have $E_{p|\mathcal{F}}\big[s(v_i, X) - s(v_{i-1}, X)\big] = 0$ for every $p \in \mathcal{P}$. For the same reason, $E_{p|\mathcal{F} \cap \Pi_i(\omega')}\big[s(v_i, X) - s(v_{i-1}, X)\big] = 0$ for all $\omega' \in \mathcal{F}$ and $p \in \mathcal{P}$. One solution to this equation is $v_i = v_{i-1}$. However, it is also the unique solution. The reason is that the left hand side is a strictly concave function of v_i, so it achieves the maximum at unique v_i. From Lemma 1, this maximum is achieved when v_i is the myopically optimal announcement. By assumption, v_i is the myopically optimal announcement, hence the only v_i that solves this equation is $v_i = v_{i-1}$. As this is true for all $i \in I$, we have $v_i = v_j$ for all i, j and agreement is reached.

It is important to note that this result is true for any $p \in \mathcal{P}$. In fact, even if we no longer have common priors, so that each player j has a set $\mathcal{P}_j$ of priors, but there is a nonempty intersection $\bigcap_{i \in I} \mathcal{P}_i \neq \emptyset$, we still get the result that there is agreement.

■

Proof of Theorem 1. ($\Leftarrow$) Suppose X is strongly separable. By Lemma 3 (i), there exists time t_k such that $\mathcal{F}^{k'}(\omega) = \mathcal{F}^k(\omega)$ for every $t_{k'} \geq t_k$. We denote this set by $\mathcal{F}^T \equiv \mathcal{F}^k$.

From Lemma 3 (iii), traders reach an agreement, hence there exists $v \in \mathbb{R}$ such that

[28] By v_0 we denote, when appropriate, the v_n.

[29] It is $p(\mathcal{F} \cap \Pi_i(\omega')) > 0$ for every $p \in \mathcal{P}$. This is because for every $\omega' \in \mathcal{F}$ there exists $p \in \mathcal{P}$ with $p(\omega') > 0$, by its definition. Regularity then implies $p(\mathcal{F} \cap \Pi_i(\omega')) > 0$ for every $p \in \mathcal{P}$.

for every $i = 1, ..., n$ it is $d_{\mathcal{P}}(\Pi_i(\omega) \cap \mathcal{F}^T, v) = v$ for every $\omega \in \mathcal{F}^T$, with $p(\omega|\mathcal{F}^T) > 0$ for some $p \in \mathcal{P}$ (this last property is trivially satisfied by the construction of $\mathcal{F}$). By defining $\mathcal{P}_{\mathcal{F}^T} = \{p(\cdot|\mathcal{F}^T) : p \in \mathcal{P}\}$, we can observe that for every $i = 1, ..., n$ it is $d_{\mathcal{P}_{\mathcal{F}^T}}(\Pi_i(\omega), v) = v$ for every $\omega \in \Omega$, with $q(\omega) > 0$ for some $q \in \mathcal{P}_{\mathcal{F}^T}$.

In Definition 3 of not strong separability, we observe that (ii) is satisfied for v and $\mathcal{P}_{\mathcal{F}}$. Because X is strongly separable, (i) should be violated, so that $X(\omega) = v$ for all $\omega \in \bigcup_{p \in \mathcal{P}_{\mathcal{F}}} Supp(p)$. This implies that information gets aggregated.

$(\Rightarrow)$ Suppose that for any regular Γ^M, information gets aggregated, so that $y_k(\omega) = d_{\mathcal{P}}(\Pi_{a(t_k)}(\omega) \cap \mathcal{F}^{k-1}(\omega), y_{k-1}) \longrightarrow X(\omega)$, for every $\omega \in \bigcup_{p \in \mathcal{P}} Supp(p)$. 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 $i = 1, ..., n$ and $\omega \in \bigcup_{p \in \mathcal{P}} Supp(p)$. Consider regular $\Gamma^M(\Omega, I, \Pi, X, \mathcal{P}, y_0, Y, s)$ with initial announcement $y_0 = v$. Then, the predictions $y_{t_k}(\omega)$, $k = 0, 1, ...$, are equal to v, for all $\omega \in \bigcup_{p \in \mathcal{P}} Supp(p)$. If we have $X(\omega) \neq v$ for some $\omega \in \bigcup_{p \in \mathcal{P}} Supp(p)$, then at ω all traders agree on v, which is the wrong value for 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.

■

B Proofs for the strategic environment

Proof of Theorem 2. For (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 first 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 $\mathcal{D}$ be the collection of regular sets of beliefs $\mathcal{P}$ that describe some uncertainty about the value of the security. That is, for each $\mathcal{P} \in \mathcal{D}$, there exist $\omega, \omega' \in \bigcup_{p \in \mathcal{P}} Supp(p)$ such that $X(\omega) \neq X(\omega')$.

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

Generalizing the notion of Ostrovsky (2012), we define the instant opportunity of Trader

i given regular beliefs $\mathcal{P}$ and previous announcement z to be

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

Note that at each partition cell $\Pi_i(\omega)$, the agent chooses a possibly different $p\in\mathcal{P}_{\Pi_i(\omega)}$ that minimises her expected utility. The instant opportunity is the ex ante (minimal over $\mathcal{P}$) expected utility, aggregating 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 agent 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 $i\in\{1,\ldots,n\}$ 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,

$$\min_{p\in\mathcal{P}_{\Pi_i(\omega)}}\sum_{\omega'\in\Pi_i(\omega)}p(\omega')\Big(s\big(d_{\mathcal{P}}(\Pi_i(\omega),z),X(\omega')\big)-s\big(z,X(\omega')\big)\Big),\tag{1}$$

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 $\omega\in\bigcup_{p\in\mathcal{P}}Supp(p)$. Moreover, because $\mathcal{P}$ is regular, each $p\in\mathcal{P}$ assigns positive probability to each $\Pi_i(\omega)$, where $\omega\in\bigcup_{p\in\mathcal{P}}Supp(p)=E$. 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, define $A^i_\omega=\{E_p[X|\Pi_i(\omega)]:p\in\mathcal{P}\}$. From the second point of Lemma 1, this is the set of possible myopically optimal announcements by i at ω, for any previous announcement z. Let $A^i=\bigcap_{\omega\in E}A^i_\omega$.

We now show that for some i, $A^i=\emptyset$. Suppose not, so that $A^i\neq\emptyset$ for all i. Since $\mathcal{P}\in\mathcal{D}$ describes some uncertainty about security X, so that condition (i) of Definition 3 is satisfied, the definition of strong separability implies that for each $x_r\in\mathbb{R}$, there exists Trader i, such that $d_{\mathcal{P}}(\Pi_i(\omega),x_r)\neq x_r$ for some $\omega\in\bigcup_{p\in\mathcal{P}}Supp(p)=E$. This implies that there does not exist x_r such that $x_r\in A^i$ for all i, hence $\bigcap_{i\in I}A^i=\emptyset$.

In part (ii), Step 2 of the proof of Lemma 3, we show that if each $A^i\neq\emptyset$ and $\bigcap_{i\in I}A^i=\emptyset$, then we can find a list of announcements v_j, one for each $j\in I$, such that v_k is agent k's myopic best response given a previous announcement of v_{k-1}, for $k\geq 0$, where $v_0=v_n$. Hence, each agent j makes the same announcement at all states in E. In part (iii) of the proof of Lemma 3, we show that $v_1=\ldots=v_n$. Because each $v_i\in A^i$, we have $\bigcap_{i\in I}A^i\neq\emptyset$, a contradiction. Note that, as we note in part (iii), the proof also works when each i has beliefs $\mathcal{P}_i$, with a non-empty intersection.

Let i be such that $A^i = \emptyset$. Because each A^i_ω is a convex set, there exist states $a, b \in E$ with $A^i_a = [c, d]$, $A^i_b = [c', d']$ such that $c' > d$. Let $k = (c' - d)/2$ and z be the previous announcement. If $z > k$ then $\min_{p \in \mathcal{P}_{\Pi_i(a)}} \max_{y \in Y} E_p\big[s(y, X) - s(z, X)\big] \geq \min_{p \in \mathcal{P}_{\Pi_i(a)}} \max_{y \in Y} E_p\big[s(y, X) - s(k, X)\big] \equiv \chi_1 > 0$, whereas if $z \leq k$ then $\min_{p \in \mathcal{P}_{\Pi_i(b)}} \max_{y \in Y} E_p\big[s(y, X) - s(z, X)\big] \geq \min_{p \in \mathcal{P}_{\Pi_i(b)}} \max_{y \in Y} E_p\big[s(y, X) - s(k, X)\big] \equiv \chi_2 > 0$. The lower bound $\chi > 0$ is just the minimum of χ_1 and χ_2. Moreover, it is independent of the previous announcement z.

■

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 $\phi \in \bigcup_{p \in \mathcal{P}} Supp(p)$ 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. Because 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 $\mathcal{Q}_\infty$.

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 expected instant opportunity of player i^* exceeds η^* at infinitely many trading times t_k.

Step 3, Case 1: Suppose that for some $\phi \in \bigcup_{p \in \mathcal{P}} Supp(p)$, there is positive probability that some random variable $q \in \mathcal{Q}_\infty$ 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 is in the ε-neighbourhood of r is greater than ζ. This can be done for every $q \in \mathcal{Q}_\infty$, and in that case the K can be selected uniformly because it is affected only by the uncertainty due to mixed strategies.[30]

[30]Indeed, the beliefs about Ω will be updated until some t, and subsequently the only change in them is because they are weighted by the belief about the mixed strategy. Therefore, because the convergence to q

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.

Because the definition of instant opportunity minimises 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 definition of instant opportunity) and hence we get that the instant opportunity at t (for t big enough) is greater than $\chi > 0$ for any previous announcement for some $\zeta > 0$ probability.[31]

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), $\bigcup_{p \in \mathcal{Q}_\infty} Supp(p)$ 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 true value.

Suppose that y_k does not converge in probability to the true 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 $|y_k - X(\omega')| > \epsilon$ is greater than δ. Because all players have more information than the outside observer, also their beliefs about the value of the security converges to the true 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 $t_{n\kappa+i^*}$ 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 $\mathcal{P}(H^k)$ 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. Note that with mixed strategies, H^k occurs with some probability. Moreover, because the equilibrium profile may consist of mixed strategies, $\mathcal{P}(H^{k+1})$ may not be the same as $\mathcal{P}(H^k)$, 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_0}.

Fix t_k, history H^k and suppose i makes an announcement. Her continuation payoff given

is dictated only by that belief, it is uniform for any belief about Ω.

[31] A proper scoring rule may not be continuous. However, Ostrovsky (2012) shows, in footnote 19 of page 2620, that his instant opportunity is continuous. This implies that our instant opportunity is also continuous.

history H^k and state ϕ, divided by β_k, is $V(H^k, \phi) = \min_{p \in \mathcal{P}_i(H^k, \phi)} E_p \sum_{\kappa=0}^{\infty} \beta^{\kappa} \left(s_{n\kappa}(\phi') - s_{n\kappa-1}(\phi')\right)$, where $s_{n\kappa}(\phi')$ is the score at state ϕ' and time $t_{n\kappa}$.

We now argue that her continuation payoff $V(H^k, \phi)$ is greater than the one-period payoff from playing the myopic strategy at t_k. The trader can guarantee such a payoff by playing the myopic strategy at t_k and then repeating the previous announcement at each subsequent period where she makes an announcement. Such a strategy guarantees the one-period myopic payoff, irrespective of how the other traders play. Because the strategy profile is sequentially rational after some period t_k, i's strategy is a best response at each information set and in particular given the beliefs $\mathcal{P}_i(H^k, \phi)$. Hence, it must provide a weakly better payoff than the one-period payoff of playing the myopic strategy at t_k.[32]

Because $V(H^k, \phi)$ is greater than the one-period payoff of the myopic strategy at t_k, we have that $\min_{p \in \mathcal{P}(H^k)} E_p V(H^k, \phi)$ is greater than i's instant opportunity given $\mathcal{P}(H^k)$ and the previous announcement at t_{k-1}, determined by history H^k.

A similar argument shows that the continuation payoff at t_k of each agent $j \neq i$, who announces after t_k, is weakly positive at each state ω and history H^k. As before, the reason is that she can guarantee a zero payoff by instructing all future selves to repeat the previous announcement. Since this is true for all states $\phi \in \bigcup_{p \in \mathcal{P}(H^k)} Supp(p)$, we have that $\min_{p \in \mathcal{P}(H^k)} E_p V(H^k, \phi) \geq 0$.

Since $\min_{p \in \mathcal{P}(H^k)} E_p V(H^k, \phi)$ is weakly positive for each $i \in I$, we have that $\sum_{i \in I} E_p V(H^k, \phi)$ is weakly positive, for any $q \in \mathcal{P}(H^k)$. Moreover, it is strictly positive if i's instant opportunity is strictly positive given $\mathcal{P}(H^k)$ and the previous announcement at t_{k-1}. Since this is true for all $q \in \mathcal{P}(H^k)$ and any previous announcement, by fixing $p \in \mathcal{P}$ and considering the (unique) probability over histories H^k 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

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

The $\overline{s}_k$ is the expected score of prediction y_k, where the expectation is over all ϕ, given the fixed $p \in \mathcal{P}$ and the moves of players according to the mixed equilibrium. We keep $p \in \mathcal{P}$ 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 occurs and i's instant opportunity is strictly positive.

The last step is identical to that of Ostrovsky (2012), because all Ψ_k are calculated using the same $p \in \mathcal{P}$. Consider $\lim_{K \to \infty} \sum_{k=1}^{K} \Psi_k$. 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^*. However, for any

[32]For the case of a revision-proof equilibrium, the analogous argument is presented in Proposition 3.

K, we have

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

where $M = \max_{y\in[\underline{y},\overline{y}],\omega\in\Omega}|s(y,X(\omega))|$. Hence, both cases of Step 3 are impossible and y_k must converge to the true 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 $\omega\in\bigcup_{p\in\mathcal{P}}Supp(p)$ and (b) $d_{\mathcal{P}}(\Pi_i(\omega),v)=v$ for all $i=1,...,n$ and $\omega\in\bigcup_{p\in\mathcal{P}}Supp(p)$.

Consider game $\Gamma^S(\Omega,I,\Pi,X,\mathcal{P},y_0,Y,s,\beta)$, where the initial announcement of the market maker is $y_0=v$. We will show that for some interim (and also ex-ante) equilibrium $(\sigma^*,\mathscr{P})$, information does not get aggregated. Define tuple $(\sigma^*,\mathscr{P})$, where σ^* specifies that each Trader i announces v after any history. At each information set $\mathcal{I}$ of Trader i, set $\mathscr{P}(\mathcal{I})=\mathcal{P}_{\Pi_i(\omega)}$. Any player, if she wants to deviate, she will deviate at some period t_k by not announcing the myopic best response, which is v. But all other players continue announcing v, hence in any subsequent period she will not gain anything, because no information is revealed, the beliefs are the same and her best response would be the myopic announcement, v. Note that this is both an ex-ante and interim equilibrium.

■

Proof of Theorem 3. The proof for (i) is identical to the proof of Theorem 2 (i). The only difference is that we need to show (in Step 4) that the continuation value of the revision-proof equilibrium is weakly greater than the player's instant opportunity, for each time $t_{k'}\geq t_k$, for some k. We show that below.

Proposition 3 *In a revision-proof equilibrium, there is time t_k such that for all $k'\geq k$, the continuation value for player i who plays at $t_{k'}$ is at least as much as her one-period payoff from playing the myopic strategy.*

Proof. We construct a deviation strategy that guarantees for the continuation game at least as much as the one-period payoff from playing the myopic strategy. Since the payoff relevant state space Ω is finite, there exists time T such that no more public information about Ω is revealed. We will prove the claim for $t>T$. We will show that for each $t_0>T$, the continuation payoff is weakly more than η, which is i's one-period payoff from playing the myopic strategy.

Suppose that at t_0, player i makes an announcement. We define a deviation strategy $\sigma=(\sigma_i,\sigma^*_{-i})$, where all $j\neq i$ follow the equilibrium strategy σ^* and σ_i is identical to σ^*_i up to time $t_0 - 1$. At t_0, 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 $H^1, ..., H^m$ be the possible paths of announcements by all other traders $j \neq i$, from t_0 to $t_0 + n - 1$, together with the common history of announcements up to $t_0 - 1$. They are finitely many, because we consider mixing over finite actions. At $t_0 + n$, σ_i specifies that:

(a) If $V(H^m, \phi, \sigma, \mathscr{P}) \geq 0$, then σ_i coincides with σ^*,

(b) If $V(H^m, \phi, \sigma, \mathscr{P}) < 0$, then σ_i repeats the previous trader's prediction.

If (a) occurs, then σ_i coincides with σ^* in every succeeding information set. If (b) occurs, then 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^*.

The deviation strategy σ_i' is feasible for i because each future self will either be indifferent or strictly prefer it over σ_i^*. Moreover, the continuation value of σ_i' is weakly above the one-period myopic payoff. The reason is that when player i deviates, every one of her future selves will get a which is weakly positive continuation payoff. The future self of i evaluates her continuation payoff using some posterior q, whereas the current i evaluates her own continuation payoff using p, and it may not be that the Bayesian update of p is q. However, because there is prior by prior updating and the future i minimises given q, her continuation payoff given the Bayesian update of p is also weakly positive. Since this is true for all possible paths, the current i has a weakly positive continuation payoff, plus the myopic payoff she gets at t_0. Since σ_i' is a feasible deviation which was not chosen by i, it must be that the continuation value of σ_i^* is weakly greater than the one-period myopic payoff.

■

For part (ii), consider the same pair $(\sigma^*, \mathscr{P})$ that was described in part (ii) of Theorem 2, where everyone announces the myopically optimal v. Since v is announced also in the case a player deviates, there is never any information revealed, hence $(\sigma^*, \mathscr{P})$ is off-path consistent. To show that $(\sigma^*, \mathscr{P})$ is revision-proof, we need to 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. Once player 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 provide two examples and an argument in order to illustrate 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.

Example 1 *Consider state space* $\Omega = \{\omega_1, ..., \omega_6\}$ *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* $\lambda_i : \Pi_i \to \mathbb{R}$ *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* $p_1 = (\frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{3}{8}, \frac{1}{8})$, $p_2 = (\frac{6}{18}, \frac{1}{18}, \frac{7}{18}, \frac{2}{18}, \frac{1}{18}, \frac{1}{18})$ *and* $p_3 = (\frac{3}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8})$. *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. Because* X *is not constant on* Ω, *it is not strongly separable and there is no information aggregation at any state.*

The previous example, together with that of Section 2, show 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 $p^C = 1/2p^A + 1/2p^B$. 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, Example 2 illustrates how the MSR model can be re-interpreted as an inventory-based market. In addition, we show that, in the inventory-based interpretation, information does not get aggregated always in the presence of ambiguity averse traders. The example is interesting because, in practice, prediction markets might not implement sequential announcements but an interface of selling and buying securities, as in Inkling Markets.

Example 2 *Consider the state space* $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, *the price function to be the* $q(z) = e^{-z}$ *where* z *is the market maker's net inventory. The security is given by* $X(\omega_1) = 2$, $X(\omega_2) = X(\omega_3) = X(\omega_4) = 1$ *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 set of priors is the* $\mathcal{P} = conv\{(0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}), (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})\}$. *Consider that initially the market maker holds zero inventory of the security (i.e. z=0).*

Firstly, Trader 1 makes a myopic decision about how much shares of the security to buy or sell. We assume, for consistency, that the amount of shares belong to $Z = p^{-1}(Y)$, *which is compact. Thus it is implied that trader solves (for the true state to be either* ω_1 *or* ω_2*)* $\max_{z \in Z} \min_{p \in \mathcal{P}} E_p[\int_0^z q(\overline{z}) - X(\omega) d\overline{z}] = \min_{p \in \mathcal{P}} \max_{z \in Z} E_p[\int_0^z q(\overline{z}) - X(\omega) d\overline{z}]$. *We have the equality by applying the same argument as in the proof of Lemma 1.*[33]

As in Ostrovsky (2012), given the price function we can define the strictly proper scoring rule $s(X(\omega), y) = \int_0^{q^{-1}(y)} q(z) - X(\omega) dz$. *We have that the price function p is 1-1 continuous with continuous inverse function. Therefore we can conclude that in the MSR market, based on that strictly proper scoring rule, the trader solves* $\max_{y \in Y} \min_{p \in \mathcal{P}} E_p[\int_0^{q^{-1}(y)} q(\overline{z}) - X(\omega) d\overline{z}] = \min_{p \in \mathcal{P}} \max_{y \in Y} E_p[\int_0^{q^{-1}(y)} q(\overline{z}) - X(\omega) d\overline{z}]$.[34] *We shall show that if* z^* *solves the first optimisation problem and* y^* *the second one, then it is* $p(z^*) = y^*$ *and that the revenue or losses are the same, i.e.* $\max_{z \in Z} \min_{p \in \mathcal{P}} E_p[\int_0^z q(\overline{z}) - X(\omega) d\overline{z}] = \max_{y \in Y} \min_{p \in \mathcal{P}} E_p[\int_0^{q^{-1}(y)} q(\overline{z}) - X(\omega) d\overline{z}]$. *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 the* $\max_{z \in Z} E_p[\int_0^z q(\overline{z}) - X(\omega) d\overline{z}]$ *is unique and such that* $p(z'_p) = E_p[X]$. *Similarly, for every* $p \in \mathcal{P}$ *the prediction* y'_p *that solves the* $\max_{y \in Y} E_p[\int_0^{q^{-1}(y)} q(\overline{z}) - X(\omega) d\overline{z}]$ *is the* $y'_p = E_p[X]$, *hence* $q^{-1}(y'_p) = z'_p$.

Therefore, for every $p \in \mathcal{P}$ *we have that* $E_p[\int_0^{z'_p} q(\overline{z}) - X(\omega) d\overline{z}] = E_p[\int_0^{q^{-1}(y'_p)} q(\overline{z}) -$

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

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

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

We conclude that 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.[35]

Finally, the first trader finds 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 $p(z^*) = 0*2+1*1 = 1$ or equivalently (as long as p is 1-1) $z^* = 0$. Hence she neither buy or sell any shares (equivalently she would have announced 1 as her prediction, i.e. the price). It is easy to see that the same would happen for every state in the partition $\{\omega_3, \omega_4\}$ and for the Trader 2 for symmetry reasons. The conclusion is that both traders does not purchase shares from the market maker and no one can infer the true state, even if that would be the case if they pooled their information.

References

Carlos E Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, 1985.

Laurence Ales and Christopher Sleet. Revision proofness. *Journal of Economic Theory*, 152: 324–355, 2014.

Simone Alfarano, Albert Banal-Estanol, Eva Camacho, Giulia Iori, and Burcu Kapar. Trade frictions in decentralised markets: An experimental study. *Mimeo*, 2019.

Geir B Asheim. Individual and collective time-consistency. *The Review of Economic Studies*, 64(3):427–443, 1997.

Pavel Atanasov, Phillip Rescober, Eric Stone, Samuel A Swift, Emile Servan-Schreiber, Philip Tetlock, Lyle Ungar, and Barbara Mellers. Distilling the wisdom of crowds: Prediction markets vs. prediction polls. *Management Science*, 63(3):691–706, 2016.

Robert Aumann. Agreeing to disagree. *Annals of Statistics*, 4:1236–1239, 1976.

Michael Bacharach. Some extensions of a claim of Aumann in an axiomatic model of knowledge. *Journal of Economic Theory*, 37(1):167–190, 1985.

Sophie Bade. Dynamic semi-consistency. *Working paper*, 2016.

Paulo Barelli, Sourav Bhattacharya, and Lucas Siga. On the possibility of information aggregation in large elections. *Working paper*, 2017.

[35] By using the saddle point inequality and the uniqueness of the optimal quantity and prediction (given the belief p^*).

Pierpaolo Battigalli, Emiliano Catonini, Giacomo Lanzani, and Massimo Marinacci. Ambiguity attitudes and self-confirming equilibrium in sequential games. *Games and Economic Behavior*, 2019.

Dorian Beauchêne, Jian Li, and Ming Li. Ambiguous persuasion. *Journal of Economic Theory*, 179:312–365, 2019.

Joyce E Berg, Forrest D Nelson, and Thomas A Rietz. Prediction market accuracy in the long run. *International Journal of Forecasting*, 24(2):285–300, 2008.

Claude Berge. *Topological Spaces: including a treatment of multi-valued functions, vector spaces, and convexity.* Courier Dover Publications, 1963.

Giacomo Bonanno. AGM-consistency and perfect Bayesian equilibrium. Part I: definition and properties. *International Journal of Game Theory*, 42(3):567–592, 2013.

Giacomo Bonanno. AGM-consistency and perfect Bayesian equilibrium. Part II: from PBE to sequential equilibrium. *International Journal of Game Theory*, 45(4):1071–1094, 2016.

Subir Bose and Arup Daripa. A dynamic mechanism and surplus extraction under ambiguity. *Journal of Economic theory*, 144(5):2084–2114, 2009.

Subir Bose and Ludovic Renou. Mechanism design with ambiguous communication devices. *Econometrica*, 82(5):1853–1872, 2014.

Glenn W Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950.

Luis Cabral, Erkut Y Ozbay, and Andrew Schotter. Intrinsic and instrumental reciprocity: An experimental study. *Games and Economic Behavior*, 87:100–121, 2014.

Colin F Camerer. Can asset markets be manipulated? a field experiment with racetrack betting. *Journal of Political Economy*, 106(3):457–482, 1998.

Colin F Camerer, Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Johan Almenberg, Adam Altmejd, Taizan Chan, et al. Evaluating replicability of laboratory experiments in economics. *Science*, 351(6280):1433–1436, 2016.

Andrés Carvajal and Joao Correia-da Silva. Agreeing to disagree with multiple priors. Technical report, Universidade do Porto, Faculdade de Economia do Porto, 2010.

Jonathan AK Cave. Learning to agree. *Economics Letters*, 12(2):147–152, 1983.

Christopher P Chambers. Proper scoring rules for general decision models. *Games and Economic Behavior*, 63(1):32–40, 2008.

Yiling Chen, Stanko Dimitrov, Rahul Sami, Daniel M Reeves, David M Pennock, Robin D Hanson, Lance Fortnow, and Rica Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, 58(4):930–969, 2010.

Yiling Chen, Mike Ruberry, and Jennifer Wortman Vaughan. Designing informative securities. *arXiv preprint arXiv:1210.4837*, 2012.

Scott Condie and Jayant V Ganguli. Ambiguity and rational expectations equilibria. *The Review of Economic Studies*, 78(3):821–845, 2011.

Peter DeMarzo and Costis Skiadas. Aggregation, determinacy, and informational efficiency for a class of economies with asymmetric information. *Journal of Economic Theory*, 80 (1):123–152, 1998.

Peter DeMarzo and Costis Skiadas. On the uniqueness of fully informative rational expectations equilibria. *Economic Theory*, 13(1):1–24, 1999.

Stanko Dimitrov and Rahul Sami. Non-myopic strategies in prediction markets. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 200–209, 2008.

Adam Dominiak and Jean-Philippe Lefort. Agreement theorem for neo-additive beliefs. *Economic Theory*, 52(1):1–13, 2013.

Adam Dominiak and Jean-Philippe Lefort. "agreeing to disagree" type results under ambiguity. *Journal of Mathematical Economics*, 61:119–129, 2015.

Anna Dreber, Thomas Pfeiffer, Johan Almenberg, Siri Isaksson, Brad Wilson, Yiling Chen, Brian A Nosek, and Magnus Johannesson. Using prediction markets to estimate the reproducibility of scientific research. *Proceedings of the National Academy of Sciences*, 112(50):15343–15347, 2015.

Pradeep Dubey, John Geanakoplos, and Martin Shubik. The revelation of information in strategic market games: A critique of rational expectations equilibrium. *Journal of Mathematical Economics*, 16(2):105–137, 1987.

Catherine C. Eckel and Philip J. Grossman. Sex differences and statistical stereotyping in attitudes toward financial risk. *Evolution and Human Behavior*, 23(4):281–95, 2002.

Catherine C. Eckel and Philip J. Grossman. Forecasting risk attitudes: An experimental study using actual and forecast gamble choices. *Journal of Economic Behavior and Organization*, 68(1):1–17, 2008.

Jürgen Eichberger, David Kelsey, and Simon Grant. Ambiguity and the centipede game: Strategic uncertainty in multi-stage games. *Working Paper*, 2017.

Mehmet Ekmekci and Stephan Lauermann. Manipulated electorates and information aggregation. *Working paper*, 2019.

Andrew Ellis. On dynamic consistency in ambiguous games. *Games and Economic Behavior*, 111:241–249, 2018.

Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, November 2003.

Urs Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10(2):171–178, 2007.

Guillaume R Fréchette and Sevgi Yuksel. Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Experimental Economics*, 20(2): 279–308, 2017.

Drew Fudenberg and Jean Tirole. Perfect bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260, 1991.

Spyros Galanis. Dynamic consistency, valuable information and subjective beliefs. *Mimeo*, 2019.

Spyros Galanis and Stelios Kotronis. Updating awareness and information aggregation. *mimeo*, 2019.

John Geanakoplos and Heraklis Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28:192–200, 1982.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

Steven M Goldman. Consistent plans. *The Review of Economic Studies*, 47(3):533–537, 1980.

Mikhail Golosov, Guido Lorenzoni, and Aleh Tsyvinski. Decentralized trading with private information. *Econometrica*, 82(3):1055–1091, 2014.

Sanford Grossman. On the efficiency of competitive stock markets where trades have diverse information. *The Journal of Finance*, 31(2):573–585, 1976.

Eran Hanany and Peter Klibanoff. Updating preferences with multiple priors. *Theoretical Economics*, 2:261–298, 2007.

Eran Hanany and Peter Klibanoff. Updating Ambiguity Averse Preferences. *The B.E. Journal of Theoretical Economics*, 9(1):1–53, November 2009.

Eran Hanany, Peter Klibanoff, and Sujoy Mukerji. Incomplete information games with ambiguity averse players. *Working paper*, 2018.

Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5 (1):107–119, 2003.

Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.

Robin Hanson and Ryan Oprea. A manipulator can aid prediction market accuracy. *Economica*, 76(302):304–314, 2009.

Robin Hanson, Ryan Oprea, and David Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*, 60(4):449–459, 2006.

Friedrich August Hayek. The use of knowledge in society. *The American Economic Review*, pages 519–530, 1945.

Paul J Healy, Sera Linardi, J Richard Lowery, and John O Ledyard. Prediction markets: alternative mechanisms for complex environments with few traders. *Management Science*, 56(11):1977–1996, 2010.

Martin F Hellwig. Rational expectations equilibrium with conditioning on past prices: A mean-variance example. *Journal of Economic Theory*, 26(2):279–312, 1982.

Lian Jian and Rahul Sami. Aggregation and manipulation in prediction markets: effects of trading mechanism and information distribution. *Management Science*, 58(1):123–140, 2012.

Atsushi Kajii and Takashi Ui. Incomplete information games with multiple priors. *Japanese Economic Review*, 56(3):332–351, 2005.

Atsushi Kajii and Takashi Ui. Interim efficient allocations under uncertainty. *Journal of Economic Theory*, 144(1):337–353, January 2009.

Christian Kellner and Mark T Le Quement. Modes of ambiguous communication. *Games and Economic Behavior*, 104:271–292, 2017.

Christian Kellner and Mark T Le Quement. Endogenous ambiguity in cheap talk. *Journal of Economic Theory*, 173:1–17, 2018.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, November 2005.

Albert S Kyle. Continuous auctions and insider trading. *Econometrica*, pages 1315 – 1335, 1985.

Nicolas S Lambert, Michael Ostrovsky, and Mikhail Panov. Strategic trading in informationally complex environments. *Econometrica*, 86(4):1119–1157, 2018.

Richard D McKelvey and Talbot Page. Common knowledge, consensus, and aggregate information. *Econometrica*, pages 109–127, 1986.

Richard D McKelvey and Talbot Page. Public and private information: An experimental study of information pooling. *Econometrica*, pages 1321–1339, 1990.

Paul Milgrom and Nancy Stokey. Information, trade, and common knowledge. *Journal of Economic Theory*, 26:17–27, 1982.

Lars Tyge Nielsen. Common knowledge, communication, and convergence of beliefs. *Mathematical Social Sciences*, 8(1):1–14, 1984.

Lars Tyge Nielsen, Adam Brandenburger, John Geanakoplos, Richard McKelvey, and Talbot Page. Common knowledge of an aggregate of expectations. *Econometrica*, pages 1235–1239, 1990.

Daniel E O'Leary. Prediction markets as a forecasting tool. In *Advances in business and management forecasting*, pages 169–184. Emerald Group Publishing Limited, 2011.

Michael Ostrovsky. Information aggregation in dynamic markets with strategic traders. *Econometrica*, 80(6):2595–2647, November 2012.

Marco Ottaviani and Peter Norman Sørensen. Outcome manipulation in corporate prediction markets. *Journal of the European Economic Association*, 5(2-3):554–563, 2007.

Marieke Pahlke. Dynamic consistency in incomplete information games with multiple priors. Technical report, Center for Mathematical Economics, Bielefeld University, 2018.

Bezalel Peleg and Menahem E Yaari. On the existence of a consistent course of action when tastes are changing. *The Review of Economic Studies*, 40(3):391–401, 1973.

Pacheco Pires. A rule for updating ambiguous beliefs. *Theory and Decision*, 53(2):137–152, 2002.

Roy Radner. Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica*, pages 655–678, 1979.

Paul W Rhode and Koleman S Strumpf. Historical presidential betting markets. *Journal of Economic Perspectives*, 18(2):127–141, 2004.

Alvin E Roth and J. Keith Murnighan. Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical Psychology*, 17:189–198, 1978.

Leonard J. Savage. *The Foundations of Statistics*. New York: John Wiley, 1954.

James Sebenius and John Geanakoplos. Don't bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78:424–426, 1983.

Luca Siga and Maximilian Mihm. Information aggregation in competitive markets. *Mimeo*, 2018.

Marciano Siniscalchi. Dynamic choice under ambiguity. *Theoretical Economics*, 6(3), September 2011.

Erik Snowberg, Justin Wolfers, and Eric Zitzewitz. Prediction markets for economic forecasting. In *Handbook of Economic Forecasting*, volume 2, pages 657–687. Elsevier, 2013.

Robert Henry Strotz. Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 23(3):165–180, 1955.

Helena Veiga and Marc Vorsatz. Information aggregation in experimental asset markets in the presence of a manipulator. *Experimental economics*, 13(4):379–398, 2010.

Emanuel Vespa. An experimental investigation of strategies in the dynamic common pool game. *Available at SSRN 1961450*, 2011.

Justin Wolfers and Andrew Leigh. Three tools for forecasting federal elections: Lessons from 2001. *Australian Journal of Political Science*, 37(2):223–240, 2002.

Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.

Asher Wolinsky. Information revelation in a market with pairwise meetings. *Econometrica*, pages 1–23, 1990.

Eric Zitzewitz. Manipulation can affect prices. *midasoracle.org*, 2007. URL http://www.midasoracle.org/2007/05/30/manipulation-can-affect-prices/.