Journal of Banking and Finance 169 (2024) 107300

Contents lists available at ScienceDirect

Journal of Banking and Finance

journal homepage: www.elsevier.com/locate/jbf

Conflicting versus reinforcing private information, information aggregation, and the time series properties of asset prices

Charles Schnitzlein [a], Patricia Chelley-Steeley [b], James M Steeley [c,*]

[a] Grossman School of Business at the University of Vermont, United States
[b] University of Birmingham, UK
[c] Brunel University London, UK

A R T I C L E I N F O

Keywords:
Information aggregation
Conflicting versus reinforcing information
Continuous double auction
Prediction market
Information efficiency

A B S T R A C T

We study how the relationship between independent private information signals affects information aggregation in laboratory asset markets. We employ two mechanisms, a continuous double auction and a prediction market. Under both mechanisms, when information is reinforcing, partial information aggregation occurs. When information is in conflict, information aggregation lessens and attempts to profit from private information frequently harm informational efficiency. In both mechanisms, results become stronger with experience in previous experimental sessions, and provide a private information benchmark for studies of the implications of conflicting public information. Under reasonable assumptions, our results are consistent with both momentum effects and weak reversals.

Understanding how heterogeneous information is incorporated into price has been described as one of the most fundamental issues in finance (Carlin, Longstaff, and Matoba (2014)). This question has motivated a large theoretical literature, a prominent branch of which studies how the behavior of strategic informed traders determines the time series properties of asset prices. A clear result from this literature is that the way in which the information structure is modeled determines liquidity patterns and the rate of information aggregation. For example, Holden and Subrahmanyam (1992) establish a baseline result by extending the Kyle (1985) monopolistic insider framework to multiple insiders with identical information. They find insiders compete aggressively and even with only two insiders, their information is reflected in price almost immediately. In sharp contrast, Foster and Viswanathan (1996) show the lower the conditional correlation between information signals, the slower the convergence of price to intrinsic value. With low positive correlation, non-monotonic patterns in liquidity arise endogenously and information aggregation is incomplete at the end of trade.

In these models, the correlation between signals is a parameter choice necessary to solve the model. In this paper, we use the control over information sets afforded by experimental asset markets to permit the realized relationship between signals to arise randomly. We do this by studying information aggregation in an asset market in which there are multiple pieces of unconditionally independent private information, which are held by different groups of traders, and together additively determine the intrinsic value of the asset. Combinations of positive and negative private information will often imply that the intrinsic value is located between the highest and lowest signals. In this case, we define the signals to be *conflicting*, and convergence of price to intrinsic value requires that the interaction of informed agents through the aggregation mechanism will induce revisions of the initial conditional expectations of at least some of the informed traders towards the unconditional expectation. In contrast, we define uniformly positive or negative private information to be *reinforcing* signals. In this case, successful aggregation requires market activity to move price farther from the unconditional expectation than each piece of private information.[1]

The information structure we employ has a natural economic

September 3, 2024We thank the Associate Editor and two anonymous reviewers for their helpful and constructive comments that have led to an improved paper, and we thank seminar participants at Hebrew University in Jerusalem, the University of Central Florida, the University of Aix-Marseille and the Ivey Business School for comments on partial results, and the University of Vermont. We also thank Philip Bond, Burton Holifield, and Andrey Ukhov for their helpful comments and suggestions. We thank Aston Business School and the College of Business Administration, UCF, for research support.

* Corresponding author.
E-mail address: james.steeley@brunel.ac.uk (J.M. Steeley).

[1] The precise definition of this ex-post signal realization classification is contained in the next section where we explain the experimental design. Appendix A Table A-1 shows the classification of all the ex-post signal realizations.

https://doi.org/10.1016/j.jbankfin.2024.107300
Received 2 May 2023; Accepted 5 September 2024
Available online 7 September 2024

interpretation. In a modern market economy, many types of events may lead to significant changes in fundamental value in large, complex corporations. Examples include legal and liability developments, merger and acquisition activity, key personnel changes, developments pertaining to product lines or competitors, technological innovation and patents, and changes in the regulatory environment: both in home markets and abroad. These types of events may engender material private information, leading to multiple independent sources of informational advantage, with each held by different groups of traders.[2]

Our experimental design employs two different mechanisms that establish benchmarks on the dimension of strategic complexity. The first is the continuous double auction (CDA), variants of which are the predominant mechanism for trading equities around the world. In our markets, traders are free to submit market or limit orders over the trading interval and there are no constraints on short sales, borrowing, the timing or the number trades. As is standard in asset market experiments, subjects participate in a sequence of independent periods. While the operation of the mechanism is intuitive and easy for subjects to understand, it is strategically very complex because traders choose the timing, frequency, and types of orders. In order to allay external validity concerns, we recruit subjects from trading and microstructure courses and have subjects return for up to three sessions of experience. The second mechanism we employ is based on a variant of a prediction market that employs a market scoring rule (MSR). Market scoring rules are currently used in a number of prediction markets, with participants compensated based on the accuracy of predictions of future events. They are easy to implement and require less trading activity than CDAs to aggregate information, as shown in (Healy, Linardi, Lowery, and Ledyard, (2010)).

Our primary motivation is to study how the relationship between information signals affects aggregation in a CDA market, a widely employed trading mechanism. Although prediction markets are not widely used for asset trading, we include this treatment in a strategically simple version consistent with a unique RNBNE to provide evidence for how strategic complexity and risk preferences (which may be masked by complex strategic interactions) affect the aggregation of reinforcing and conflicting information. The simplification of the prediction market experimental design (e.g. fewer participants and a single prediction opportunity) precludes a controlled comparison of the two mechanisms, and we do not conduct this analysis.

We find that under both mechanisms, information aggregation depends on whether private information signals are reinforcing or conflicting, and while experience in a previous session improves efficiency when signals are reinforcing, it damages efficiency when signals are conflicting. In detail, when signals are reinforcing, in the CDA markets, end-of-period prices are more efficient than the unconditional expectation, but price errors are high and price never reaches intrinsic value by the end of the trading period. With experience in a previous session, aggregation improves significantly, but even in a third session of experience, less than half of private information is reflected in price. In the prediction market (PM), aggregation of reinforcing signals also improves significantly with experience, but convergence to intrinsic value is also incomplete in this mechanism and end-of-period price errors are large and statistically significant.

When signals are conflicting, under both mechanisms, market activity does not improve informational efficiency relative to the unconditional expectation. In sharp contrast to the markets with reinforcing signals, information aggregation deteriorates with experience in a previous session. Under both the CDA and PM, with experienced subjects, more than one-third of ending prices are on the "wrong" side of the unconditional expectation implying that attempts to profit from private information damage informational efficiency. Although not part of the experimental design, these patterns coupled with the eventual disclosure of the information in private signals and natural variation in the realized distribution of information signals provide a potential explanation for both weak negative serial correlation when signals are in conflict and much stronger momentum effects when private information is reinforcing.[3]

The control over information sets afforded by the experimental design allows us to identify explanations for the contrasting aggregation results for conflicting and reinforcing signals. In the CDA, subjects with signals that imply a large expected revision in fundamental value, with experience, use market orders more aggressively. This trades off small losses when information is conflicting for large profits when information is reinforcing, and leads to the aggregation patterns we observe. In the PM, predictions that deviate from the risk-neutral Bayesian Nash equilibrium by trading off reduced expected profit for reduced profit variability are increasing in both the extremeness of the private information signal and the extremeness of the previous prediction relative to the unconditional expectation.

Our analysis of how aggregation is affected by the realized (ex-post) relationship between signals complements and extends earlier research that shows under some information structures, attempts to profit from conditionally independent private information can lead to reductions in informational efficiency relative to a public information benchmark. In a setting with conditionally independent signals, Healy et al (2010) find incomplete aggregation in both continuous double auction and prediction market mechanisms. We also find incomplete aggregation under both mechanisms, but also show how the realized relationship between signals and experience interact to determine the degree of information aggregation, with experience improving aggregation when signals are reinforcing but reducing aggregation when signals are conflicting, and there being more aggregation when signals are reinforcing.

Jian and Sami (2012) compare different implementations of a prediction market under information structures with unconditionally independent or conditionally independent private signals in a very simple market with two traders, binary signals, and a binary security. They find incomplete aggregation in all treatments, evidence that markets perform better when the two agents receive identical signals (rather than opposite signals), some evidence for slower convergence with unconditionally independent information, and no evidence for learning effects as subjects gain experience in a session in this single-session experiment. Our finding of incomplete aggregation in a much larger market in which subjects gain experience in a previous session (larger markets promote aggregation, and experience in a previous session is a necessary although not sufficient for successful aggregation e.g. Forsythe and Lundholm (1990)) is additional evidence for the unlikelihood of achieving aggregation with independent information.[4]

Linardi (2017) studies a prediction market with two privately informed predictors and varies the state space and dispersion of conditionally independent signals. The experimental design implies full aggregation assuming perfect rationality, however aggregation is incomplete. Linardi finds in almost all cases, subjects overweight their priors, and predictions that improve aggregation are more frequent when private information signals are similar than when they are

[2] The cross-border activities of many firms may be of particular importance in generating these types of information events. For example, in 2017, 43.6% of sales in the S&P 500 represented products and services produced and sold outside the U.S. (Compiled by S&P Dow Jones Indices LLC from data provided by S&P Global Market Intelligence).

[3] Thus providing an alternative explanation to certain behavioural phenomena such as conservatism and representativeness bias (Barberis et al, 1998), overconfidence and self-attribution bias (Daniel et al, 1998) of under and over-reaction in asset price movements.

[4] Our finding that neither the CDA nor a PM mechanism is able to successfully aggregate multiple independent information signals is consistent with the theoretical prediction of Chen et al (2010) that the aggregation of independent signals may be impossible within a finite trading interval.

different. We find similar results in a setting with unconditionally independent signals and a much larger state space. Specifically, we find more aggregation relative to the public information set when signals are reinforcing. In addition, we study the effect of experience in a previous session, and show that its effect depends on whether signals are reinforcing or conflicting. Our result that attempts to profit from private information can lead to less efficient outcomes than the unconditional expectation when signals are conflicting is related to Linardi's finding that aggregation attempts result in less efficient outcomes than the unconditional expectation under information structures where it is difficult to discern the information content of signals.

Our results further complement a literature that shows public disagreement impedes efficient pricing (Diether, Malloy and Scherbina (2002) and subsequent papers) wherein stocks with higher dispersion in analysts' forecasts earn lower returns than otherwise similar stocks. They show that this evidence is consistent with the hypothesis that prices will reflect the optimistic view whenever investors with the lowest valuations do not trade. This result is explained in subsequent work to be due to short-sales constraints (Boehme et al, 2006) and high transaction costs which hinder the efficient aggregation of information (Sadka and Scherbina (2007)).

An issue with using analyst forecasts as proxies for beliefs is that they are generated through combinations of public and private information, and they are likely influenced by how forecasts relative to consensus affect analyst credibility (Brown, Call, Clement, and Sharp (2015)). In this paper, we show that dispersion in private information is associated with inaccurate pricing without appealing to short-sales constraints or other asset specific transactions costs. Periods with conflicting signals have greater information dispersion and experience less aggregation, and experience harms this further. Given the strong evidence that managers delay the disclosure of bad news relative to good news (Kothari, Shu, and Wysocki (2009)), and delay the disclosure of bad news due to career concerns (Baginski et al., 2018), the aggregation patterns we observe when signals are in conflict combined with the relatively slow release of bad news documented in the literature are sufficient to generate the pattern observed in Diether et al (2002).

Our study also builds on previous experimental research that shows traders do not accurately infer what market activity conveys about others' private information. Early studies show that with a simple information structure and perfectly informed insiders, information is fully aggregated by the end of trade (e.g. Plott and Sunder (1982), Copeland and Friedman (1987), Forsythe, et al (1982)). However, the seminal paper by Plott and Sunder (1988) and especially subsequent papers that incorporate imperfect, diverse information find failure to converge to a rational expectations equilibrium that implies complete aggregation of private information. Other examples include Biais et al (2005) and Corgnet et al (2015). Especially relevant is the study by Page and Siemroth (2020), who in a meta-analysis of 664 markets from five different studies and additional new experiments, find markets successfully incorporate public information into prices, but typically impound less than 30% of private information by the end of trade. This growing body of experimental results and empirical evidence from a variety of field market settings that investors do not sufficiently account for the informational content of market activity has motivated theoretical research that models the implications for volume and informational efficiency (Eyster, Rabin, and Vayanos (2019)).[5] They show in a static model that when traders do not fully appreciate what prices convey about others' private information, prices underreact to private signals, leading to incomplete aggregation. We add to this literature by examining how the aggregation of diverse information depends on whether pieces of private information are reinforcing or in conflict.

The plan of the paper is as follows. We next discuss the experimental design. In Section 2, we develop hypotheses and explain our methods of analysis. In Section 3, we report our results and in Section 4, we conclude.

1. Experimental design

Our information structure is represented as follows. Let PI represent the public information set, and assume intrinsic asset value (V) is determined by three unconditionally independent private information signals, with each signal (S_i) a draw from an i.i.d. mean zero distribution that indicates the revision in intrinsic value due to one piece of information;

$$V = E[V|PI] + \sum_{i=1}^{3} S_i \tag{1}$$

A trader observing signal S_i therefore enters the market with $E[V|S_i] = E[V|PI] + S_i$, because by assumption additional information is unpredictable and mean zero. The key feature of the setting we study is that since independent events generate new pieces of information, they combine additively to determine value.

We employ two information aggregation mechanisms: a continuous double auction (CDA) and a prediction market (PM). Both experimental designs have multiple independent periods in each experimental session. In each period the intrinsic value of the single risky asset (denominated in lab dollars (L$)) is determined as follows. We take three independent draws from a uniform distribution with support on the integers on the interval [-6,+6], each representing an information signal. The end-of-period value of the risky asset is the sum of the three draws plus L $100. This results in a symmetric, approximately bell-shaped distribution for intrinsic value with end-points L$82 and L$118, an unconditional expected value ($E[V|PI]$) of L$100, and support on the integers. The correlation between each signal and intrinsic value is 57.7% ($1/3^{1/2}$). In the CDA mechanism, each information signal is held by two informed traders in order to induce competition.

As it is possible for all three ex-post signal realizations to be of the same sign, or for one of them to be of a different sign to the other two, we formally distinguish the ex-post signal sets between those that are reinforcing and those that are conflicting. Specifically,

$$ConflictingSignals : E[V|PI] + S_H \geq V \geq E[V|PI] + S_L \tag{2a}$$

$$ReinforcingSignals : V > E[V|PI] + S_{H;} or V < E[V|PI] + S_L \tag{2b}$$

where $E[V|PI]$ is the unconditional (conditional on the public information, PI) mean of the distribution of the intrinsic value, V, and the highest and lowest signals in the ex-post signal realization are S_H and S_L, respectively. If the realized signals are all of the same sign, then the signal set is defined to be reinforcing. If the realized signals are of differing signs, then the signal set is defined to be conflicting, unless the absolute size of the conflicting signal (within the set of three signals) is smaller than the absolute size of the smallest of the two reinforcing signals (within the set of three signals). For example, if the unconditional expected asset value is L$100, and the information signals are {4, 5, -1}, the asset value is L$108 and the signals are classified as reinforcing, because $|-1| < |4|$. If the information signals are {4, 5, -6}, the asset value is L$103 and the information signals are classified as

[5] Eyster et al (2019) review experimental and empirical evidence that traders do not sufficiently appreciate what prices convey about others' private information. Avery and Zemsky (1998) show that herding can affect price when there is "composition uncertainty" (uncertainty as to the average accuracy of the signals). The different aggregation process that is required in each of the conflicting versus reinforcing cases in this study can be interpreted as generating a form of "composition uncertainty".

conflicting, because |-6|≥|4|. This classification rule means that the intrinsic value is outside of the range of the signals when the signals are reinforcing but not exceeding the range of the signals when the signals are conflicting. This makes the information aggregation task distinctly different in each case. With reinforcing signals, successful aggregation requires market activity to move price further from the unconditional mean than the most extreme signal, while for conflicting signals, successful aggregation requires revisions of the initial conditional expectations of at least two sets of informed traders towards the unconditional expectation.

Although two signals would have been sufficient to distinguish conflicting and reinforcing signals, we use three signals as we wanted the focus to be on the ability to infer diverse information from market activity. If we had had just two signals, the extreme signal of the pair will always be on the same side of the unconditional expectation, $E[V|PI]$, as the intrinsic value, V. With three signals this may be the case, for example {5, 3, -3} where V=L$105, or may not be the case, for example {5, -3, -3} where V=L$99. Given the potential importance of the unconditional expectation as a focal point, the three-signal design enables us to study the interactions between two traders with high value signals (recall that each signal is received by two traders in the CDA to induce competition), and four traders who have received moderate signals, but who together hold more information. This would not have been possible in a 2-signal design.

As well as the aggregation task being distinctly different depending on whether the periods have conflicting or reinforcing signals, the classification also reflects different levels of disagreement between the informed traders in each situation. Disagreement can be measured by the standard deviation of the 3 signals values. Ex-ante, for periods with reinforcing signals the standard deviation of the signal values is on average 2.48, while for conflicting signals it is on average 4.16. Thus, periods with conflicting signals reflect more disagreement than those with reinforcing signals. Our definition of reinforcing and conflicting signals also generates a difference in the potential extremeness of the realized asset value since, by construction, the sum of conflicting signals is bounded in [-6,6], whereas for reinforcing signals it is bounded in [-18,18]. To ensure that our analysis reflects the differing type of aggregation tasks, rather than just the differing levels of asset value extremeness, we control for the different levels of extremeness in our measures of aggregation and the associated explanatory regressions.

The experimental instructions explain how the intrinsic asset value is determined, and the distribution of intrinsic values conditional on each possible signal. Given the information structure, each informed agent's initial conditional expectation is equal to their draw plus the unconditional expectation. More extreme draws therefore imply a larger updating of beliefs and contain more information.

The participants in the experiments were either students from the Business School at Aston University (University A) or the College of Business Administration at the University of Central Florida (University B). The subjects from University A were MSc Finance students and drawn primarily from a class that studied microstructure and trading. The subjects from University B were either undergraduates or MBA students recruited from an advanced elective in trading and market microstructure, or undergraduate and graduate students in finance that had previously participated in common value multi-unit auction experiments. Each treatment was conducted at both universities but each subject participated in only one mechanism, even if they participated in two or three sessions.

Our use of subjects that have participated in a previous session is an important feature of the experimental design. Although this complicates the statistical analysis of the results, studies that incorporate experimental asset markets show experience in a previous session promotes convergence to equilibrium outcomes.[6]

1.1. Continuous double auction (CDA)

We conduct 19 CDA experimental sessions (13 at University A and 6 at University B), each lasting on average 2.5 hours with inexperienced subjects and two hours when subjects have experience in a previous session. There are ten initial sessions, seven sessions in which subjects participate for the second time, and two sessions where subjects participate for the third time. Each session involves a group of eight subjects comprised of six informed traders and two liquidity traders, interacting over 12 independent trading periods. We randomly assigned subjects to roles that were maintained over the entire session. In sessions with experienced subjects, we assigned the role of liquidity traders to subjects who had previously been informed traders. At the beginning of each trading period there are new random draws that determine the three independent information signals that together determine the single risky asset's intrinsic value as in Eq. (1).

We impose liquidity shocks that are exogenous to intrinsic value, such that each of the liquidity traders is required to finish the trading period with a predetermined position in the asset.[7] Including traders with both informational and non-informational trading motives is a natural feature of equity markets, which affects strategic interactions and market dynamics. In addition, although we do not replicate the specific technical conditions for a no-trade theorem to apply, the presence of liquidity traders gives informed traders a clear motive to trade, as traders in this group should earn profits on average.

One liquidity trader is a buyer while the other is a seller. Prior to the start of each period, the required end-of-period position of each liquidity trader (which is private information) is determined by an independent draw from a publicly known discrete uniform distribution ranging from one to five units. The net liquidity trade therefore forms a triangular distribution with mean of zero, standard deviation of two, and support on the integers between -4 and 4, endpoints included, where a negative sign indicates liquidity traders are net sellers. If a liquidity trader does not meet the required position, a penalty is imposed at the end of the period equal to L$36 times the absolute value of the deviation between the required position and the actual end-of-period position. The penalty ensures that demand is perfectly inelastic at the required position. If a liquidity trader goes bankrupt (this occurs if the end-of-period cash balance drops below zero) they are removed from subsequent market periods, but this is not disclosed to the market.[8]

Traders interact through a transparent double-auction (open limit order book) conducted on a series of networked personal computers with custom software. Each subject's screen provides continuously updated market information, including bids and asks, transaction prices, whether trades are buyer or seller initiated, net market order imbalance (buyer initiated trades less seller initiated trades) and the time

[6] For example, Forsythe and Lundholm (1990) show that in a market with diverse information, experience in a previous session is a necessary (although not sufficient) condition for successful information aggregation. In a more complicated setting with a single insider, but two markets that are open simultaneously, de Jong et al (2006) find two previous sessions of experience are necessary before subjects employ strategies consistent with equilibrium behavior. Even in settings with a simple information structure and symmetric information, experience in a previous session promotes equilibrium pricing (e.g. Smith, et al (1988), and the large literature it spawned).

[7] The use of liquidity traders with trading targets has frequently been used in experimental asset market research (e.g. Schnitzlein (1996), Cason (2000), Bloomfield, O'Hara, and Saar (2005)).

[8] In the 19 CDA experimental sessions, one liquidity trader in one session with first-time participants was removed from trading due to a negative cash balance.

remaining in the current trading period. Screens also identify each subject's own transactions, asset positions, cash balance, and either information signals (informed traders) or end-of-period required positions (liquidity traders).

During the trading period, each trader is free to submit a limit buy (bid), a limit sell (ask), to remove an outstanding bid or ask, or to transact against the highest bid or the lowest ask with a market order. Each quote is for a single unit. We impose no constraints on the timing, number, or direction of trades. Limit orders are ranked by price and then time priority. When a trade occurs, all subjects observe the transaction and the price, but do not learn trader identities.

Subjects begin the first trading period with a monetary endowment. Informed and liquidity traders have starting cash balances of L$200 and L$300 respectively in the first nine sessions at University A. In the six sessions at University B and the last four sessions at University A, starting cash balances are L$280 and L$340.[9] Differences in initial balances are intended to minimize differences in profits by trader type. Subjects are not endowed with shares of the risky asset but are allowed to buy or sell-short an unlimited number of shares over the course of a trading period.

Each trading period consists of three minutes of "economic time." Since our interest is in studying information aggregation, and since market activity can be fast, the market is paused after each transaction or quote revision until all participants acknowledge a market activity message. This allows subjects to carefully update their beliefs after each action. Each trading period therefore took between six and twelve minutes.[10]

At the beginning of each session, subjects read the experimental instructions (which are available in the online appendix). An experimenter then reviewed the trading rules, parameter values, distributions, and information structure. Before trading commences, subjects took a written quiz that tested their understanding of the relationship between the information signals and the intrinsic asset value (see the appendix). Subjects had no difficulty with the quiz, but the correct answers were always discussed with the subjects prior to the first trading period.

We employ nine sets of random draws across the 19 sessions in order to fully represent the possible combinations of exogenous variables, but never repeat a set of draws with subjects that participate more than once. We used three of the same sets of draws with both inexperienced and experienced subjects to ensure differences by level of experience are not driven by differences in the draws. After 11 sessions (six initial sessions and five with subjects participating for the second time) we were surprised by the consistent failure of price to converge to intrinsic value. We therefore modified the experimental procedures in eight additional sessions as follows. First, we added a second quiz question pertaining to the relationship between the private information signals and intrinsic value. Second, from each of four new initial sessions (two at each location) we recruited the subjects with above median profits to participate in a second and third session to be completed within six days.[11]

At the end of a trading period, each subject is informed of the intrinsic value of the asset and any penalties incurred (in case of the liquidity traders). Trading profits are calculated by closing out all remaining long or short positions at the intrinsic value. Each subject is informed of their end-of-period cash balance after liquidation of positions at the intrinsic asset value, and these are carried forward to the start of the next trading period.[12] After the completion of the final trading period, subjects were given a brief questionnaire that assessed their understanding of the experiment, ending cash balances are multiplied by a pre-announced exchange rate to convert L$ to the local currency and each subject is paid their earnings in private. Payoffs average $34 per subject, including a payment of $5 for arriving to the experimental lab on time.

1.2. Prediction market (PM)

In this type of market, a market scoring rule (MSR) determines the payoff of a prediction as a function of its accuracy relative to the accuracy of the preceding prediction. In our simple implementation, each period begins with a computer making a public prediction of the asset value equal to its unconditional expectation. The first player then makes a prediction (P_1), which all participants observe. This is followed by the second player's prediction (P_2), and so on. After each player has made a single prediction, the asset value is revealed and each prediction is assigned a "score" equal to the squared prediction error relative to the true asset value. Each player's payoff is then determined by the score of her prediction relative to that of the previous prediction. Thus, the payoff of subject i depends on the score of the prediction relative to the score of the previous prediction,

$$\text{Score}_i = (\text{prediction}_i - V)^2 \tag{3}$$

$$\text{Payoff}_i = -1 \times (\text{Score}_i - \text{Score}_j) \tag{4}$$

where prediction j (and its associated score) precedes prediction i.

In this setting, the first predictor maximizes expected profit by predicting $P_1 = E[V|PI] + S_1$, and each subsequent prediction P_i equal to $P_{i-1} + S_i$ also maximizes expected profit. Under this sequence of predictions, the last prediction results in complete information aggregation, and is consistent with the result in Hanson (2003) that when risk-neutral traders predict only once, truthful predictions maximize expected profit and form the unique risk-neutral Bayesian-Nash equilibrium (RNBNE).

We conduct 14 PM sessions (seven at each university) on networked computers using software created with Z-tree (Fischbacher (2007)). There are eight initial sessions and six sessions in which subjects participate for the second time. In each session, four subjects participate in 12 independent prediction periods. We drew the subjects from

[9] We adjusted the difference between informed and liquidity trader starting balances due to smaller informed trader profits than we expected in the initial nine sessions at University A. Adjusting these sorts of "nuisance variables" has been shown to have little effect in asset market experiments (see for example Copeland and Friedman (1987)).

[10] Since we were interested in studying information aggregation and markets can move very fast for novice traders, our design allowed traders to consider the information content of activity, before taking actions. Trying to replicate one feature of extant markets (continuous trading) without other salient features (computer driven algorithmic trading, professional experience with the trading mechanism, for example) would move us away from meaningful results. Kocher et al (2019) show that subjects with exhausted self-control can make poor decisions, specifically a tendency for over-pricing. While our market pause extends the trading period, so potentially adding to fatigue, the pauses themselves give moments for reflection and resetting of emotions that can counteract this.

[11] Subjects in these four initial sessions were not informed that eligibility in future sessions would be performance dependent. Before beginning the second session, subjects knew they would be eligible to participate in a third session. Further detail on the distribution of subjects across the 19 CDA sessions is given in Table A-2a in the appendix.

[12] Since traders' ability to buy shares is not constrained by their cash balance and since their ability to sell shares is not constrained by their share balance, this has no effect on buying or selling power. It does make final payments depend on performance in every period and not a randomly selected period (which is the case in some experimental designs). It does have the possible disadvantage of introducing wealth effects but we believe this is outweighed by the increased salience due to every decision having payoff implications.

comparable subject pools as in the CDA sessions, although no subject had participated in those sessions. Of the 32 subjects in the initial session, 24 returned to participate in a second session.[13] Information signals and the asset value are determined exactly as in the CDA markets.

After reviewing the instructions, subjects took a quiz that tested their understanding of the relationship between the signals and the intrinsic value and how the score and profit (or loss) of each prediction is determined. Subjects had no difficulty with the quiz, but as a further review, the correct answers were discussed with the subjects.

At the beginning of each period, the three independent information signals are randomly distributed across the subjects: one subject does not receive information. Subjects begin the first period with a monetary endowment of L$250. The computer predicts L$100 to start each period, which is the common information, unconditional expected value of the asset. One of the four subjects is then randomly selected to make a prediction of the asset value and this process is repeated until all four predictions have been made. The full prediction history is reported after each prediction. Both the allocation of signals to subjects and the prediction order are independent and random across periods. There was no time limit once a subject's turn to predict commenced.

After each subject has made a single prediction, the intrinsic asset value (V) is disclosed and subjects privately learn their profits or losses for the period. The payoff for a subject is determined by the Market Scoring Rule.

At the completion of the final prediction period, subjects were given a brief questionnaire that assessed their understanding of the experiment. They were then paid their earnings in local currency in private, at a pre-announced exchange rate times their ending cash balances, and then dismissed. Sessions averaged approximately one hour and per-subject payments averaged $19.

2. Hypotheses and methods of analysis

2.1. Continuous double auctions (CDA) hypotheses

Foster and Viswanathan (1996) (FV) show in a multi-period framework with multiple informed traders, liquidity traders, and a trading protocol based on the Kyle (1985) model, the lower the conditional correlation between information signals, the slower and less complete the convergence of price to intrinsic value. A key feature of FV is that as trading proceeds, the correlation between the remaining informational advantage of insiders becomes negative, inducing traders to trade less aggressively to avoid revealing information to the market. Thus, the lower the assumed correlation between the signals, the more rapidly the conditional correlation between the signals becomes negative and the less informative the price process. Back, Cao, and Willard (2000) generalize FV to continuous time and show that with less than perfect correlation between the informed traders' signals, by some date and all dates thereafter, the market would have learned more from a monopolistic insider than competing informed traders.[14]

In these models, the correlation between signals is a parameter choice necessary to solve the model. Our setting is more complex with the information signals unconditionally independent, no restrictions on the timing of trades, and the realized relationship between signals (conflicting versus reinforcing) arising randomly. Since our setting is beyond the reach of tractable modeling, we use the intuition gleaned from these theoretical results that less information is revealed through trading when the signals have lower correlation, to establish our first hypothesis.

H1 (CDA): Information aggregation is incomplete

FV also show that toward the end of the trading interval (when the conditional correlation between signals becomes negative), market liquidity, as measured by the implicit bid-ask spread declines. Our second hypothesis is therefore:

H2 (CDA): Market liquidity declines at the end of the trading period.

Our third hypothesis pertains to the role of experience in a previous session on information aggregation. We note in our discussion of the experimental design (Footnote 5 above) that a large body of experimental asset market evidence shows that experience in a previous session promotes equilibrium outcomes. Indeed, Forsythe and Lundholm (1990) show that in a market with diverse information, experience in a previous session is a necessary (although not sufficient) condition for successful information aggregation. Therefore:

H3 (CDA): Information aggregation improves with experience in a previous experimental session.

2.2. Prediction markets (PM) hypothesis

Hanson (2003) established a myopic honesty result: a risk-neutral predictor with a single prediction opportunity will maximize expected profit by reporting his true belief conditional on his information. Since we limit each trader to a single prediction, the unique risk-neutral Bayesian-Nash equilibrium (RNBNE) is that the first trader will predict his conditional expectation given his private signal and the unconditional expectation, the second trader will predict her conditional expectation given the first prediction, and so on.[15] Given the additive information structure, this implies complete aggregation and leads to the following hypothesis:

H4 (PM): Information aggregation is complete.

2.3. Measuring aggregation

We measure information aggregation within a trading period or a prediction period in three ways.

2.3.1. End-of-period price error (PE)

Our first measure is the price error relative to intrinsic value on the final transaction (CDA) or final prediction (PM) of the period. Since the three information signals collectively determine intrinsic value, this is also a measure of convergence to the rational expectations equilibrium, and is a measure of "strong form" efficiency that has been widely used in the finance literature.[16]

$$\text{Convergence}: \mathrm{PE} = |\mathrm{V} - \mathrm{P_T}| \tag{5}$$

where V is the intrinsic asset value and P_T is the final price or prediction of the period. Full efficiency would imply a value of zero, while the maximum price error within the bounds of the distribution of the asset value is L$36 for periods with reinforcing signals and L$24 for periods with conflicting signals. To reflect this difference, we also report the price errors as a proportion of the corresponding maximum price error.

2.3.2. Aggregation relative to the average signal (AGG_{AS})

Even if the market were quite informationally efficient, exact convergence to intrinsic value might be infrequent in a market with a

[13] Further detail on the distribution of subjects across the 14 PM sessions is given in Table A-2b in the appendix.

[14] Ostrovsky (2012) shows that for "additive" securities, at the end of a bounded trading interval, the market price converges to the security's expected value conditional on the traders' pooled information but does not provide results on how the relationship between signals affects the speed of convergence.

[15] Although deviations from risk-neutral behavior in the low-stakes environment of the experimental design have been observed in many previous experiments, we use this prediction as a benchmark.

[16] As a robustness check, under the CDA we calculate this measure and subsequent results using the bid-ask spread midpoint at the time of the last transaction, and the average of the last two transactions. Results are almost identical with no changes in statistical significance and are not tabulated.

large state space and multiple unconditionally independent private signals. We therefore also measure the degree of aggregation relative to the average signal.[17] Defining the price error of the average signal, $PE_{AS} = |V - (E[V|PI] + (1/3)\sum_{i=1}^{3} S_i)|$, the absolute value of the difference between the intrinsic asset value and the average signal;

$$AGG_{AS} = PE_{AS} - PE \tag{6}$$

Since the average signal is calculated from the sum of three unconditionally independent signals, it represents one-third of the private information held by informed traders. We also calculate a scaled version of this measure, by dividing by PE_{AS}, that has the interpretation of the percentage improvement in efficiency relative to the average signal and is independent of the extremeness of the asset value. If the final price had been at the average signal, this scaled measure will take the value zero, and if the price error (PE) is zero, this measure takes the value unity.[18]

2.3.3. Aggregation relative to the unconditional expectation (SS_{EFF})

We also calculate the price error of the final price relative to what the price error would have been if the final price had been the unconditional expectations. This is a measure of semi-strong form efficiency, efficiency relative to the public information set:

$$SS_{EFF} = (|V - E[V|PI]| - PE) \tag{7}$$

A positive value indicates an improvement over semi-strong efficiency while a negative value indicates that attempts to profit from private information damage informational efficiency (the efficiency of the final price is worse than a price that is semi-strong efficient, i.e., is at E[V|PI]). We also scale this measure, by dividing by (|V – E[V|PI]|), so that it is independent of the extremeness of the asset value and to enhance interpretation of its magnitude. If the final price had been at the unconditional expectation, this scaled measure will take the value zero, and if the price error (PE) is zero, this measure takes the value unity.[19]

2.3.4. Exogenous variables and convergence to intrinsic value

We use panel regressions with trading period data (CDA) or prediction period data (PM) to examine how exogenous variables affect convergence to intrinsic value. We control for signal dispersion with the standard deviation of the signals within a period, INFO(σ), and we control for the absolute value of the distance of the asset value from its unconditional mean, which we refer to as asset value extremeness, abs (V-E[V|PI]).[20] We include indicators for experience in a previous session (EXP) and whether the signals are reinforcing (REINF) or in conflict (CONF), and interact the REINF and CONF indicators with EXP. In order to control for the possible autocorrelation in the residuals from each predictor (or set of traders) making multiple decisions in each session we report robust t-statistics that cluster on each predictor (PM) and each session (CDA), with the number of degrees of freedom set to the number of clusters minus one). In each regression, the measure of convergence is PE as defined above, and is related to the exogenous variables through (7) below:

$$\begin{aligned} PE_{i,t} = {} & b_0 + b_1 abs(V - E[V|PI])_{i,t} + b_2 INFO(\sigma)_{i,t} + b_3 REINF_{i,t} \\ & + b_4 REINF_{i,t} {*} EXP_{i,t} + b_5 CONFL_{i,t} {*} EXP_{i,t} + e_{i,t} \end{aligned} \tag{8}$$

where, the sessions are indexed by i=1,2,…n, the trading (or prediction) periods by t=1,2,…,T, and where the total number of observations is N=nT.

In our analysis of the CDA, we also assess whether the direction of the net liquidity demand affects information aggregation. In each period the independently and randomly determined liquidity trader demand and supply typically results in a net liquidity trader imbalance. LIQ DIR, a variable that takes on the value of one when the net liquidity trade is in the same direction as intrinsic value relative to the unconditional expectation, negative one when the opposite is true, and zero when net liquidity demand is zero allows us to examine how non-informational trading affects the aggregation process.

In the prediction markets, we also examine the effect of off-equilibrium predictions on convergence. We do this by including the variable NOT EQ %, the percentage of predictions in a period inconsistent with the equilibrium prediction under the assumption that the previous prediction was an equilibrium prediction.

2.4. Measuring and explaining order choice strategy in the CDA

The early theoretical market microstructure literature assumed fully-informed traders use only market orders, but later theoretical work established that informed traders will use both limit orders and market orders.[21] Bloomfield, O'Hara, and Saar (2005) use experimental asset markets to study liquidity provision in an electronic limit order market, and find that fully-informed insiders compete against each other aggressively with market orders until their information is impounded in price and then switch to limit orders and become the primary providers of liquidity. Their study provides evidence that a market-making role for insiders can arise endogenously. In contrast we explain information aggregation by examining the order submission strategies of partially informed traders with independent information.

We use logit regressions to examine the order choice between a limit and a market order, and then extend our analysis to consider the choice between putatively profitable and putatively unprofitable market orders. A market order is putatively profitable if it implies a profit conditional on the initiator's signal. A putatively unprofitable market order implies a loss conditional on the initiator's signal but may be profitable if the accumulated information in the order flow combined with the information in a trader's signal implies an updated conditional expectation. We use the following explanatory variables in the regressions. The expected profit from a market order conditional on a trader's signal (EPMO) is the difference between the inside bid and the informed trader's signal if the signal is below the midpoint of the bid-ask spread, or the difference between the signal and the inside ask if the signal is above the midpoint of the spread. As EPMO increases, the attractiveness of market orders may increase because of execution certainty. The inside spread at the time of the order submission captures the difference in expected profit (conditional on the signal) between submitting a limit order that is just exposed to the market, and a market order. The inclusion of this variable is motivated by the theoretical prediction from Foucault (1999) that the frequency of limit orders is increasing in the bid-ask spread. Bloomfield et al (2005) show that order choice changes over the trading period so we include the variable TIME, which is the number of seconds that have elapsed in the trading period. Since each trader makes multiple order choices, we estimate robust Z-statistics,

[17] The use of the average signal as a benchmark corresponds to the prior-information equilibria in Plott and Sunder (1988) and subsequent papers.

[18] We thank both anonymous referees for prompting us to reconsider the interpretation of these measures that has resulted in us reporting scaled versions of both AGG_{AS} and SS_{EFF} (that follows below), which have clearer interpretations. The scaled versions are both undefined when the asset value (V) equals the unconditional expectation (E[V|PI]). The ex-ante probability of this occurring is 5.7%. It occurs in our drawn signal sets in 7.9% [7.1%] of periods for the CDA [PM] and these periods are excluded from the calculation of these scaled measures.

[19] Numerical examples of each of the efficiency measures can be found in the appendix (Figure A-1).

[20] When signals are conflicting (reinforcing), the average dispersion of the signals is greater (less) than the unconditional standard deviation of the signals, while the distance of the asset value from the unconditional mean is less (greater) than the unconditional distance of the asset value from the unconditional mean. The ex-post correlation between these two variables is -0.37.

[21] See for example Chakravarty and Holden (1995), Kaniel and Liu (2006) and Goettler, Parlour and Rajan (2009).

clustering on each trader.

2.5. Measuring and explaining deviations from Equilibrium in the PM

Under the RNBNE, the first predictor's signal is the conditional expected asset value and a prediction equal to the signal maximizes expected profit. It does not however minimize the variance of profit because if signals are conflicting, an equilibrium prediction may imply a loss, and a payoff of zero is certain if the prediction equals the unconditional expectation. A deviation from the RNBNE by the first predictor that is a partial adjustment to the information in the signal can therefore be interpreted as due to risk aversion, although the interpretation on later predictions is more complicated because subjects are forced to make conjectures about the relationship between previous predictions and signals.

We measure deviations from the RNBNE with SDEQ: the signed deviation from the RNBNE under the assumption that at each prediction stage, each trader assumes the previous trader employed the RNBNE strategy. SDEQ is signed negative if the deviation is closer to the unconditional expectation (L$100) and positive if it is farther from the unconditional expectation than the RNBNE prediction. We also construct three indicator variables to help interpret deviations from equilibrium. "Closer-to-Previous-Prediction" is one if the prediction is a partial adjustment to the information in the signal from the previous prediction and zero otherwise. "Closer-to-L$100" is one if SDEQ is negative and zero otherwise. This variable helps interpret whether deviations from the RNBNE may be due to a risk-reduction motive. "Wrong Direction" is an indicator that takes on the value one if a prediction is in the opposite direction implied by the signal under the assumption that the previous prediction was consistent with the RNBNE prediction.

We first document the frequency and significance of SDEQ, and calculate foregone profits due to deviations from the RNBNE. We then use logit regressions to estimate the determinants of each type of out-of-equilibrium behavior (Closer-to-L$100, Closer-to-Previous Prediction, and "Wrong Direction"). We include the following explanatory variables. EXTeq is the absolute value of the extremeness of the RNBNE implied by the new signal minus the absolute value of the extremeness of the previous RNBNE. absPREV is the absolute value of the difference between the previous RNBNE and the unconditional expectation.[22] absSIG is the absolute value of the difference between the unconditional expectation and the signal held by the predictor. We also include the previously defined variables EXP and REINF, interact the CONF and REINF indicators with EXP, and include indicators for the second (2nd), third (3rd), and fourth (4th) prediction in a period.

3. Results

Our data are drawn from 19 CDA experimental sessions and 14 PM experimental sessions, each consisting of 12 periods. In each session the first two periods are compensated training periods, and are not analyzed.[23] This yields a CDA data set of 190 trading periods, 3473 transactions, and 17,726 order submissions or cancellations. In the analysis that follows, we refer to the ten initial sessions as inexperienced sessions, and the nine sessions in which subjects participated in a previous session as experienced sessions.[24] In the PM sessions, there are eight inexperienced sessions comprising 80 prediction periods (320 individual predictions), and six experienced sessions with 60 prediction periods (240 individual predictions). We divide our results into three parts.[25] In 3.A.1, we report information aggregation under the CDA, and in 3.A.2 report results for the PM sessions. In 3.B., we use period-level data to analyze the exogenous factors that affect convergence. In 3.C., we analyze the strategic behavior of the subjects in the CDA and PM with the goal of explaining observed aggregation patterns.

3.1. Information aggregation in the continuous double auction

The CDA markets are liquid in the sense that after the first five to ten seconds of the trading periods, there is almost always an outstanding bid and ask. At the end of the trading period, there is at least one bid and ask outstanding in 98% of the trading periods, and the average spread at the time of transactions averages (L$1.53), which is 26% of the average level of asset value extremeness abs(V-E[V|PI]). On average, there are 18.3 trades in each trading period.[26]

Our initial analysis of the CDA markets shows that the sessions with second- and third-time participants are statistically similar and that all significant changes due to experience occur between the first and the second sessions. We verify this by regressing convergence (PE) on the exogenous variables abs(V-E[V|PI]), INFO(σ), REINF, and indicator variables for second and third sessions. We find no evidence that the coefficient estimates on the experience indicators are different ($p=0.97$). We therefore group the trading periods from second and third sessions. This simplifies the exposition, and we have verified this grouping does not change the main results.[27]

In Table 1, we report the average levels of information aggregation and pricing errors across the sessions by level of experience and the relationship between information signals. We conservatively test hypotheses by averaging period results for each session and then treating the average from each session as a single observation. An additional measure of information aggregation is the frequency with which attempts to profit from private information result in end-of-period prices which are less efficient than the unconditional expectation, indicating that attempts to profit from private information resulted in a deterioration of efficiency. We call this "excess volatility" and define this to occur if the ending price error is greater than the absolute value of the difference between the intrinsic value and the unconditional expecta-

[22] We use the RNBNE in calculating each value of EXTeq and absPREV to avoid the endogeneity that would arise if we were to use previous predictions.

[23] Bossaerts, et al (2010) is an asset market experiment that employs two initial practice periods that are sometimes compensated. Their results suggest the compensation of practice periods does not affect outcomes.

[24] Under the CDA we thus group together sessions with second- and third-time participants. We show below that this is statistically justified.

[25] Summary statistics, by mechanism, level of experience and the ex post relationship between signals, for the number of sessions, periods and subjects and for the exogenous regressors: average asset value extremeness, average ex-post standard deviation of the signals and (for the CDA) average trading period volume, are reported in Table A-3 in the appendix.

[26] Consistent with the empirical literature (e.g. Carlin, et al (2014) and many previous papers), volume is positively related to disagreement. While this and other studies must use proxies for disagreement, we measure it directly with the standard deviation of the three pieces of information that determine intrinsic value. Regressing period-level volume ($N=190$) on the standard deviation of private information, the coefficient on the disagreement measure is significant ($p=0.01$). We also run this regression adding the percentage of informed orders that are market orders on the RHS and find the coefficient to be significant, both when we keep the "disagreement" variable, and when we exclude it ($p<0.01$). We show below that the aggressive use of market orders by informed traders with extreme signals explains our results pertaining to information aggregation.

[27] Recall that as part of the experimental design, two cohorts were formed with the best performers from four first sessions. Further breaking out the levels of experience to include a separate indicator for this group of subjects does not change the results.

Table 1
Information Aggregation and Excess Volatility.

Sessions	Periods	Experience	PE		AGG_{AS}		SS_{EFF}		Excess Volatility
Panel A. All Trading Periods: Continuous Double Auction (CDA)									
10	100	Inexperienced	4.56 (4.03-5.09)	{0.14}	-0.65** (-2.63)	[-33%]	1.31*** (4.70)	[11%]	10%
9	90	Experienced	4.03 (3.63-4.45)	{0.13}	-0.39** (-2.54)	[-34%]	1.44*** (8.40)	[11%]	12%
Panel B. Reinforcing Signals: CDA									
10	59	Inexperienced	6.25 (5.49-6.99)	{0.17}	-0.64* (-1.86)	[-14%]	2.17*** (5.86)	[24%]	10%
9	42	Experienced	5.81 (5.06-6.55)	{0.16}	0.49 (1.66)	[9%]	3.65*** (11.08)	[39%]	0%
Panel C. Conflicting Signals: CDA									
10	41	Inexperienced	2.19 (1.79-2.60)	{0.09}	-0.63*** (-3.99)	[-57%]	0.16 (0.84)	[-5%]	10%
9	48	Experienced	2.49 (2.00-2.98)	{0.10}	-1.15*** (-6.26)	[-90%]	-0.46** (-2.34)	[-26%]	23%
Panel D. All Prediction Periods: Prediction Markets (PM)									
8	80	Inexperienced	3.96 (3.16-4.76)	{0.12}	-0.20 (-0.65)	[-21%]	1.69*** (5.61)	[19%]	13%
6	60	Experienced	3.55 (2.29-5.87)	{0.12}	-0.08 (-0.17)	[-27%]	1.65** (3.37)	[11%]	10%
Panel E. Reinforcing Signals: PM									
8	52	Inexperienced	4.51 (3.76-5.26)	{0.13}	0.35 (1.18)	[7%]	2.78*** (9.49)	[38%]	6%
6	24	Experienced	(2.79-5.87)	{0.12}	2.00** (3.33)	[31%]	5.17*** (8.60)	[54%]	0%
Panel F. Conflicting Signals: PM									
8	28	Inexperienced	2.99 (1.42-4.56)	{0.12}	-1.20* (-1.90)	[-83%]	-0.32 (-0.51)	[-22%]	25%
6	36	Experienced	3.03 (1.46-4.59)	{0.13}	-1.47* (-2.42)	[-85%]	-0.69 (-1.14)	[-23%]	17%

We test hypotheses by treating the average (in L$) from each session as a single observation. Significance for a two-tailed test relative to zero at the 10%, 5%, and 1% level are indicated with *,**, and *** respectively with t-statistics in parentheses. PE is the absolute value of the average price error on the final transaction, with the 95th percentile confidence interval relative to zero in parentheses. In curly brackets we report the PE divided by the maximum possible ex-ante price error. Our primary measure of information aggregation is aggregation relative to the average signal (AGG_{AS}). We also report informational efficiency relative to the public information set (SS_{EFF}). In square brackets we report the average percentage improvement in efficiency relative to the average signal (AGG_{AS}) or the public information set (SS_{EFF}); the scaled measures. Excess volatility is the percentage of periods when either the ending price is less efficient than the unconditional expectation, in that it ends on the wrong side of unconditional expectations relative to the intrinsic value, or overshoots the intrinsic value by more than the difference between the intrinsic value and the unconditional expectation. In the case of conflicting signals, where undershooting and overshooting are more likely, we apply a hurdle equal to two-times the ending bid-ask spread (L$3) for the under- or over-shooting for the CDA (PM) markets.

tion.[28] This occurs in 21 (11.0%) of the periods. Most of these periods (81.0%) are due to price moving in the "wrong" direction, ending on the opposite side of the unconditional expectation to the intrinsic value, with the reminder due to price overshooting intrinsic value. In the aggregate, 71.4% of these periods are when signals are conflicting, including 100% of the instances with experienced subjects. In these periods, the asset value is 47% less extreme than average and the dispersion of signals is 34% higher than average.

Information aggregation relative to the average signal (AGG_{AS}) fails in the CDA markets (Table 1, Panel A). In inexperienced sessions, the last transaction price is 11% more efficient than the unconditional expectation ($p<0.01$), although 10% of market periods are less efficient than the unconditional expectation and consistent with excess volatility. Last prices average 33% less efficient than the average signal: AGG_{AS} is L$-0.65 (N=10, $p<0.03$). In experienced sessions, the ending price is 11% more efficient than the unconditional expectation ($p<0.01$), but still 34% less efficient than the average signal: AGG_{AS} is L$-0.39 (N=9, p=0.04). The improvement in information aggregation relative to the average signal with experience is not significant (p=0.38).

The failure of information aggregation relative to the average signal implies that price does not converge to intrinsic value.[29] In Table 1, we show 95th percentile confidence intervals around final price errors PE, which are large and significantly greater than zero. Despite small overall improvements in informational efficiency, with experience in a previous session the percentage of periods in which market activity culminates in prices consistent with excess volatility with respect to fundamentals increases from 10% to 12%.

In Table 1, Panels B, we report results for periods in which the signals are reinforcing. In sessions with inexperienced subjects, ending prices are 24% more efficient than the unconditional expectation by an average of L$2.17, ($p<0.01$), but 14% less efficient than the average signal: AGG_{AS} = L$-0.64 ($p<0.10$). In experienced sessions, ending prices are 39% more efficient than the unconditional expectation. AGG_{AS} increases to L$0.49 and is significantly greater than in inexperienced sessions (p=0.02), but not significantly greater than the average signal (p=0.13). The effect of experience on AGG_{AS} is depicted in Fig. 1, Panel A.

Despite the highest levels of information aggregation when signals are reinforcing, price errors on the final transaction (PE) are also highest in these periods because reinforcing signals are associated with more extreme asset values, but are still higher when adjusted for the greater range of possible price errors. In the inexperienced CDA sessions, only 24% of the information in the private signals is reflected in price by the last transaction, rising to 39% in the experienced sessions.[30]

In contrast to reinforcing signals, aggregation and convergence

[28] Since the maximum values of the ending price error, PE, and the absolute value of the difference between the intrinsic value and the unconditional expectation, $|V - E[V|PI]|$, imply that it is more likely for periods with conflicting signals to display excess volatility, we define a period with conflicting signals (in the CDA) to display excess volatility if PE exceeds $|V - E[V|PI]|$ by more than twice the ending bid-ask spread. Similarly, in the periods with conflicting signals where $|V - E[V|PI]|=0$, we define a period to display excess volatility if PE alone exceeds twice the ending bid-ask spread. Periods with $|V - E[V|PI]|=0$ comprise 19% [16%] of the excess volatility periods in the CDA [PM] markets. In the Prediction Market for periods with conflicting signals, we replace the hurdle of twice the bid-ask spread with a hurdle of L$3. This hurdle is consistent with the CDA market hurdle as the average ending spread in the CDA is not significantly different from 1.50 ($p>0.29$). Our definition of excess volatility concerns only the price error on the final transaction, relative to the distance between the intrinsic value and the unconditional expectation, it does not imply a particular behavior for price changes within the trading period.

[29] Recall that with the asset value equal to the sum of three independent pieces of information, the average signal only incorporates 1/3 of the information held by the informed traders in aggregate: convergence to intrinsic value requires inferring the information of other traders from market activity.

[30] Incomplete aggregation is not due to extreme results in a relatively small number of sessions. In the online appendix (Figure A-2, Panel A) we show that this result obtains in all of the 19 sessions. The percentage of information in the private signals reflected in price by the last transaction is calculated as $(P_T-E[V|PI])/(V-E[V|PI])$. It is undefined in periods where $V=E[V|PI]$. For values within the interval $(-\infty,1]$ it equals $SS_{EFF}/|(V-E[V|PI]|$.

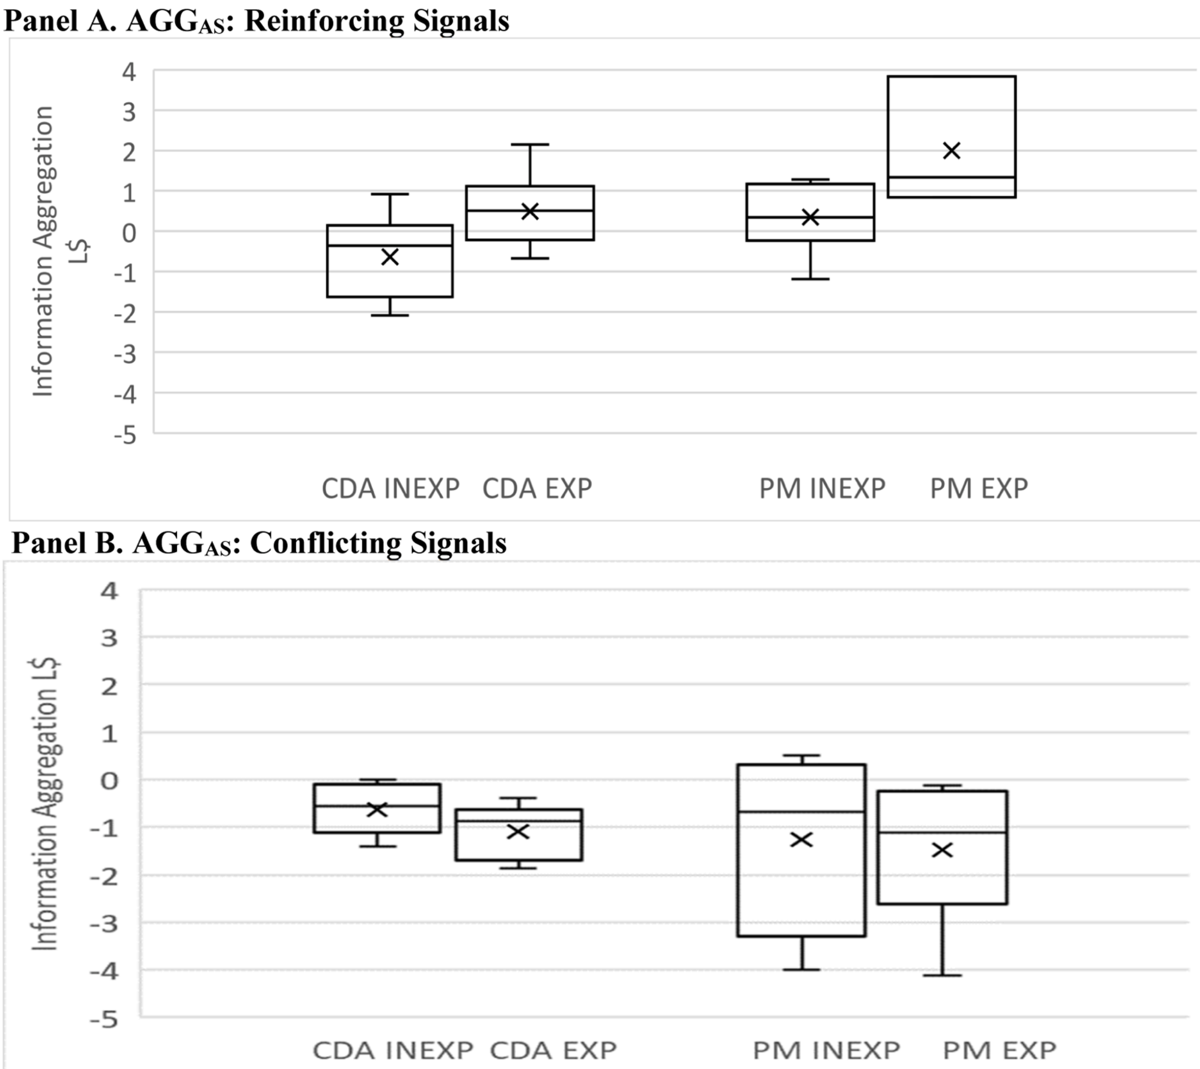

Fig. 1. Information Aggregation Relative to the Average Signal.
We group the results by mechanism, whether subjects have experience in a previous session, and by the relationship between signals (reinforcing versus conflicting). The interquartile range (IQR) is defined as quartile 3 (Q3) minus quartile 1 (Q1). The mean and median are indicated with an X and the proximate line respectively.

decline with experience when signals are conflicting. Informational efficiency relative to the unconditional expected value is L$0.16 ($p=0.44$) in the inexperienced sessions, and declines to L$-0.46 ($p=0.05$), which is 26% worse than the public information set, in experienced sessions: efficiency relative to the unconditional expectation (the public information set) is lower with experience ($p=0.04$). In Table 1, Panel C, we report the aggregation measure when the signals are in conflict. The values of AGG_{AS} for CDA markets are L$-0.63 ($p<0.01$) (57% less efficient than the average signal) in inexperienced sessions, and L$-1.15 ($p<0.01$) (90% less efficient than the average signal) in experienced sessions: Information aggregation is lower with experience, ($p=0.04$). Periods characterized by excess volatility only occur in sessions with experienced subjects when information signals are conflicting, increasing from 10% in the inexperienced sessions to 23% of the periods when subjects have experience in a previous session.

In sum, in the CDA sessions with inexperienced subjects there is no difference in AGG_{AS} due to the relationship between signals ($p=0.97$). However, with experience, 1) AGG_{AS} increases significantly when signals are reinforcing, and the percentage of periods with excess volatility declines from 10% to 0%; 2) when signals are conflicting, AGG_{AS} declines significantly and the percentage of periods with excess volatility increases from 10% to 23%; and 3) the difference in AGG_{AS} (reinforcing vs. conflicting) is significant ($p<0.01$). In net, there is a small but insignificant increase in AGG_{AS} in experienced sessions. In Section 3.B, we examine the determinants of our measures of aggregation and excess volatility, controlling for the differences in asset value extremeness and signal dispersion. In Section 3.C. we examine the behaviors that generate the aggregation patterns that we have described above.

3.2. *Information aggregation in the prediction markets*

Despite implementing a strategically simple version of a prediction market with a unique RNBNE that is consistent with complete information aggregation, information aggregation fails in the prediction markets, with similar patterns as in the CDA markets. In inexperienced sessions, the final prediction is 19% more efficient than the unconditional expectation ($p<0.01$), but less efficient (-21%) than the average signal, although not significantly: AGG_{AS} is L$-0.20 ($p=0.54$), Table 1, Panel D. In experienced sessions, final predictions are 11% more efficient than the unconditional expectation (p=0.55), and are less efficient (-27%) than the average signal, but again, not by a significant margin (AGG_{AS} = -0.08, p=0.87). The improvement in information aggregation with experience is not significant (p=0.85)

When information signals are reinforcing, ending prices are 38% more efficient than the unconditional expectation for inexperienced sessions, rising to 54% more efficient with experience (p<0.01). Table 1, Panel E shows that ending prices are 7% more efficient than the average signal for, but not significantly, in inexperienced sessions: AGG_{AS} is L$0.35, (p=0.28). In experienced sessions, AGG_{AS} averages L$2.00, and is significantly larger than in both inexperienced sessions (p=0.04) and (31% more efficient) the average signal (p=0.02). We depict the effect of experience on AGG_{AS} in Fig. 1, Panel A.

As in the CDA markets, despite the highest levels of information aggregation when signals are reinforcing, price errors on the final transaction are also highest in these periods because reinforcing signals are associated with more extreme asset values, though less so after adjusting for this extremeness. In sessions with inexperienced subjects and reinforcing signals, 38% of the private information is incorporated in the final prediction, rising to 54% in experienced sessions.

Similar to the CDA markets, aggregation and convergence decline with experience when signals are conflicting. In the prediction markets, final predictions are less efficient than the unconditional expected value in both inexperienced (-22%) and experienced sessions (-23%), although in neither case the difference is significant; L$-0.32 (p=0.62) and L$-0.69 (p=0.31) respectively. The values of AGG_{AS} are L$-1.20 for inexperienced subjects (p=0.10) and L$-1.47 for experienced sessions (p=0.06). The decline in aggregation (-83% to -85%) with experience is not significant (p=0.76). We show these differences in Fig. 1, Panel B. As in the CDA markets, in experienced sessions, periods characterized by excess volatility only occur when signals are conflicting, decreasing from 25% of the periods in inexperienced sessions to 17% in experienced sessions.

Overall, in the prediction markets, AGG_{AS} is significantly higher when signals are reinforcing than when signals are conflicting. As in the CDA markets, the significance of this difference increases with experience, rising from L$1.53 (p=0.05) in inexperienced sessions to L$3.47 (p<0.01) in the experienced sessions.

3.3. Information aggregation and convergence to intrinsic value: summary results

In sum, these results indicate that under both mechanisms, with both inexperienced and experienced subjects, ending prices are more efficient than the unconditional expectation. This result is consistent with aggregation relative to the public information set, as measured by SS_{EFF} (Eq. (7)), and hence ending prices that are better than semi-strong form efficient. Ending prices equal to the average signal imply the aggregation of one-third of the information in the private information signals, yet under both mechanisms, this threshold is not attained, even in experienced sessions. We therefore do not reject H1 (CDA information aggregation is incomplete) but reject H3 (CDA information aggregation improves with experience) and H4 (PM information aggregation is complete). However, under both mechanisms, and our most notable finding, is that experience in a previous session improves information aggregation when information signals are reinforcing, but lowers aggregation with experience when signals are in conflict. Thus, H3 is rejected for periods with conflicting signals but not rejected for periods with reinforcing signals. In addition, periods characterized by excess volatility increase in experienced sessions, and only occur when signals are in conflict.

3.4. Exogenous determinants of convergence of price to intrinsic value

In Table 2, Panel A, we examine the exogenous determinants of convergence to intrinsic value using the CDA trading periods (N=190). The high significance of the coefficient on the abs(V-E[V|PI]) variable (p<0.01) indicates prices do not converge to intrinsic value.[31] Consistent with the session-level results, price errors (PE) decline significantly with experience when the signals are reinforcing (p=0.03), but increase when the signals are conflicting although not significantly (p=0.12). The difference between these two coefficient estimates is significant (p<0.01). In the second regression, we also include LIQ DIR, a variable that captures the direction of the net liquidity demand relative to the direction of the intrinsic asset value relative to the unconditional expectation. The estimated coefficient is negative and significant (p=0.02), indicating the direction of an imbalance in non-informational trades supports aggregation when it is in the direction of the revision in price necessary for aggregation, but harms aggregation when the opposite is true.[32] In periods exhibiting excess volatility, the asset value is 47% less extreme than average and the dispersion of signals is 34% higher than average. In the third regression, we use these same exogenous variables as controls and examine how the relationship between signals and experience are related to excess volatility (EV). The signal dispersion coefficient is positive and significant (p=0.01). Consistent with the session-level results in Table 1, instances of EV decline with experience when the signals are reinforcing (p=0.12), but increase when the signals are conflicting (p=0.17). Although neither is significant, the difference between these two coefficient estimates is significant (p=0.02). This adds to the evidence that while the more aggressive attempts to exploit private information with experience can improve aggregation when signals are reinforcing it can do the opposite when signals are in conflict.

In Panel C, we report qualitatively similar results for the prediction market periods (N=140). PE declines with experience when signals are reinforcing (p=0.08) and increases with experience when signals are conflicting, although not significantly (p=0.74). The difference between the coefficient estimates for REINF*EXP and CONF*EXP is not significant (p=0.15). The effect of experience in the PM periods is attenuated because in contrast to the CDA markets, higher (although insignificant) aggregation with reinforcing signals occurs in the inexperienced periods.

3.4.1. Order submission behavior in the CDA

Under the CDA we have found: 1) the consistent failure of convergence to intrinsic value when signals are reinforcing but a modest (but significant) improvement in aggregation and convergence with experienced subjects; and 2) a deterioration in information aggregation and convergence with experienced subjects when signals are in conflict. In this section we examine the order submission strategies that generate these results. Bloomfield, O'Hara and Saar (2005) show informed traders use market orders when the value of their information is high. We can thus use the choice of order type, market order versus limit order, as a way to infer beliefs and whether, or not, they are being updated from market activity. For example, a choice of a putatively unprofitable order, conditional on a signal, could indicate an updating of information.

Informed traders comprise three-fourths of the traders in each market and are responsible for 72% of the market orders submissions. Before analyzing their behavior, we first characterize the behavior of liquidity

[31] In sections 3.B and 3.C, p-values are from significance tests of coefficients from regression equations, in contrast to the mean comparison tests that featured in section 3.A. The details are in Tables, 1, 2 and 3.

[32] We also interacted LIQ DIR with the conflicting and reinforcing indicators, but neither was significant (results not tabulated).

[33] Because the number of clusters is small, we conservatively test hypotheses using degrees of freedom equal to the number of clusters minus one. As a robustness test, we supplement this with a wild bootstrap procedure (Djogbenou, Mackinnon, and Nielsen (2019)), and obtain almost identical results.

Table 2
Price Errors (PE) and Excess Volatility (EV) with Period-level Data and Controls for Exogenous Variables.

	abs(V- E[V|PI])	INFO(σ)	REINF	LIQ DIR	REINF* EXP	CONF *EXP	cons.	R^2
Panel A. Price Errors under the CDA (N=190)								
(1) PE	0.64*** (13.79)	0.04 (0.44)	0.09 (0.16)		-1.04** (-2.33)	0.41 (1.62)	0.63 (1.64)	62.8%
(2) PE	0.64*** (13.45)	0.02 (0.23)	0.06 (0.11)	-0.29** (-2.46)	-1.03** (-2.29)	0.33 (1.27)	0.74* (2.08)	63.7%
(3) EV	-0.02* (-1.76)	0.03** (2.84)	0.17 (1.73)	-0.01 (-0.57)	-0.08 (-1.66)	0.17 (1.44)	-0.00 (-0.01)	11.0%
Panel B. Informed Trader Market Orders and Price Errors under the CDA (N=190)								
	abs(V- E[V|PI])	INFO(σ)	REINF	LIQ DIR	MO*REINF*EXP	MO*CONF*EXP	cons.	R^2
(4) PE	0.65*** (13.63)	0.01 (0.16)	0.09 (0.20)	-0.32*** (-2.88)	-0.08** (-2.53)	0.04*** (3.00)	0.60 (1.60)	65.7%
Panel C. Price Errors in the Prediction Markets (N=140)								
	abs(V- E[V|PI])	INFO(σ)	REINF	NOT EQ %	REINF* EXP	CONF *EXP	cons.	R^2
(5) PE	0.50*** (5.20)	0.09 (0.72)	-0.61 (-0.65)		-1.27* (-1.92)	0.33 (0.74)	1.24 (1.26)	25.6%
(6) PE	0.51*** (3.91)	0.04 (0.31)	-1.03 (-1.23)	3.21*** (3.91)	-1.02 (-1.72)	0.34 (0.42)	-0.74 (-0.88)	32.1%

Significance for a two-tailed test relative to zero at the 10%, 5%, and 1% levels are indicated with *,**, and *** respectively. t-statistics in parentheses are calculated with cluster-robust standard errors.[33] abs(V- E[V|PI]) is the extremeness of the asset value in absolute value relative to the unconditional expectation. INFO(σ) is the standard deviation of the signals. REINF is an indicator for periods when the information signals are reinforcing. CONF is an indicator for periods when the signals are in conflicts and along with REINF is interacted with EXP, an indicator for periods with experienced subjects. In the CDA markets, in Panel A we include LIQ DIR, a variable that takes on the value of one when the net liquidity trade is in the same direction as intrinsic value relative to the unconditional expectation, negative one when the opposite is true, and zero when net liquidity demand is zero. Excess Volatility (EV) is an indicator variable taking the value unity when the ending price meets our definition of excess volatility. In Panel B we multiply the REINF*EXP and CONF*EXP variables by the number of informed trader market orders in a period. In the Prediction markets (Panel C), NOT EQ % is the percentage of predictions in a period inconsistent with the equilibrium prediction under the assumption that the previous prediction was an equilibrium prediction.

traders. We showed in the previous section that the direction of the exogenous net liquidity demand relative to intrinsic value (LIQ DIR) is a significant determinant of information aggregation. Other measures of liquidity trader activity (the total number of trades in a trading period, or trades in excess of the required number) are insignificant. LIQ DIR also explains most of the variation in liquidity traders' profits: the simple regression of liquidity traders' profits per period on LIQ DIR yields an R^2 of 67% and a t-statistic of 24.25. Liquidity traders use limit orders to complete 47% of their trades and most frequently use them in order to lower the cost of attaining the exogenously determined share balance by placing a bid or ask inside the bid-ask spread.

We next analyze the determinants of the informed trader's choice between a limit order and market order submission. Conditional on there being both a bid and ask at the time of order submission, approximately three-fourths of order submissions are limit orders in both inexperienced and experienced sessions. This percentage is slightly higher when signals are in conflict, but not significantly.

In Table 3, we report the results of logit regressions designed to explain the choice between a market order and a limit order. Panel A shows that market order submissions are increasing in the expected profit (EPMO) ($p<0.01$), and decreasing in the size of the inside spread (IS) ($p<0.01$). The coefficients on Time and the experience in a previous session indicator are insignificant. Repeating this analysis by level of experience, results are qualitatively similar.

Further decomposing market orders into those that are either putatively profitable conditional on the trader's signal (MOP) or putatively unprofitable (MOU) lends insight into patterns of information aggregation. While MOU orders may be due to mistakes or speculation, they can also be helpful in revealing information if they are due to a trader updating beliefs based on market activity, and for example, buying at a price above a trader's signal. The percentage of MOU orders is roughly similar with inexperienced and experienced subjects at 8.2% and 6.6%

Table 3
Order Choice in the Continuous Double Auction.

	EPMO	Inside Spread	TIME	EXP	Pseudo R^2	N
Panel A: Determinants of a market order relative to a limit order						
All Sessions (1-19)	0.09***	-0.33***	-0.0003	0.07	0.07	9624
	(3.10)	(-7.32)	(0.29)	(0.78)		
Inexperienced Sessions (1-10)	0.10***	-0.26***	0.0004		0.07	5078
	(3.13)	(-5.46)	(0.36)			
Experienced Sessions (11-19)	0.09*	-0.44***	0.0005		0.09	4546
	(1.73)	(-5.42)	(0.25)			
Panel B: Determinants of market order type relative to a limit order: Inexperienced Sessions						
Putatively Unprofitable Market Orders	-0.19**	-0.21***	0.002		0.07	5078
	(-2.53)	(-2.58)	(1.22)			
Putatively Profitable Market Orders	0.22***	-0.35***	-0.002			
	(6.90)	(-5.54)	(-1.57)			
Panel C: Determinants of market order type relative to limit order: Experienced Sessions						
Putatively Unprofitable Market Orders	-0.31***	-0.60***	0.004*		0.10	4546
	(-2.84)	(-5.69)	(1.69)			
Putatively Profitable Market Orders	0.17***	-0.41***	-0.002			
	(2.69)	(-4.29)	(-1.04)			
Panel D: Determinants of each type of market order relative to limit order: Experienced sessions, Reinforcing Signals						
Putatively Unprofitable Market Orders	-0.68***	-0.75***	0.007***		0.14	2072
	(-4.32)	(-5.13)	(2.86)			
Putatively Profitable Market Orders	0.23***	-0.46***	-0.002			
	(2.90)	(-4.60)	(-0.70)			

Significance for a two-tailed test relative to zero at the 10%, 5%, and 1% level are indicated with *,**, and *** respectively with cluster robust z-statistics in parentheses. EPMO is the expected profit of a market order conditional on the trader's signal at the time of order submission. "Inside spread" is the size of the bid-ask spread. "TIME" is the number of seconds that have elapsed in the trading period at the time of order submission. EXP is an indicator for all sessions with subjects with experience in a previous session.

respectively. However, patterns as a function of the relationship between signals are significantly different. We verify this by repeating the regressions, but distinguishing between MOU and MOP orders. With inexperienced subjects (Panel B), MOP orders depend on EPMO and IS (as in Panel A). MOU orders are less likely when EPMO is high (p=0.01) and when IS is large (p=0.01). In the experienced sessions (Panel C), results are similar but MOU are more likely with the passage of time (p=0.09). The influence of time on putatively unprofitable market orders is entirely due to trading behavior in trading periods when the signals are reinforcing. We show this by breaking out these trading periods in the experienced sessions (Panel D). Results are similar except for the strong significance of Time on the propensity to submit putatively unprofitable market orders (p<0.01). This can also be seen in Fig. 2, Panels A and B. When the signals are reinforcing, MOU become more frequent in experienced sessions. In contrast, when the signals are in conflict, MOU do not increase with time. The increased use of putatively unprofitable markets orders when signals are reinforcing helps explain the higher degree of convergence with experience in these trading

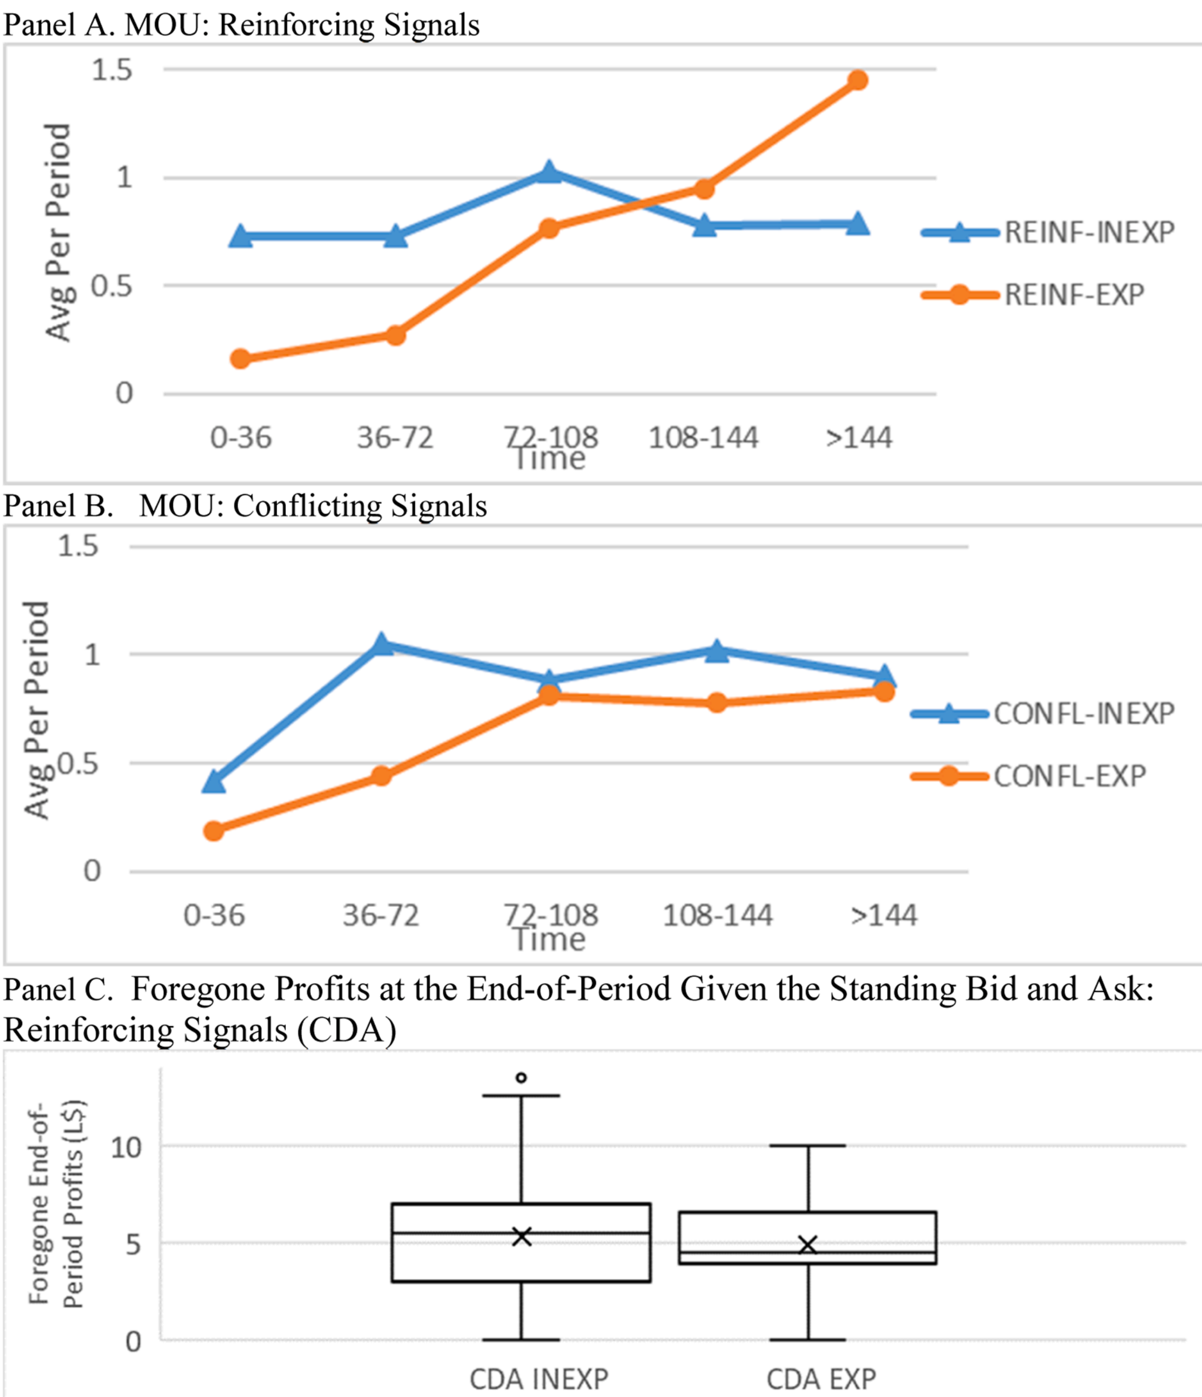

Fig. 2. Unprofitable Orders Conditional on a Trader's Signal.
The first two panels depict market orders that are unprofitable conditional on a trader's signal (MOU) for trading periods when information signals are reinforcing (Panel A) and conflicting (Panel B). Periods are broken out by inexperienced sessions (INEXP) and session in which subjects have experience in a previous session (EXP). Panel C depicts foregone profits at the end of the trading period due to not making a single transaction at outstanding quotes.

periods: traders are learning to use the information in market activity to trade beyond their information.

In contrast, the decline in efficiency with experience in a previous session when signals are conflicting is due to informed traders with extreme signals submitting more market orders. On average, informed traders with signals that imply a revision of the unconditional expected asset value of L$4 or more submit 39.2% more market orders when signals are in conflict, and 13.7% more market orders when signals are reinforcing than informed traders in inexperienced sessions, despite no difference in the magnitude of bid-ask spreads.[34] When signals are in conflict, the more aggressive use of market orders by experienced traders helps explain the deterioration in information aggregation: informed traders lose an average of L$5.57 per period on these market orders compared with losses of L$2.41 in inexperienced sessions.

In Table 2, Panel B we examine the impact of the number of market orders in a period on convergence to intrinsic value by multiplying the REINF*EXP and CONF*EXP variables by the number of informed trader market orders in a period. This modification increases the R^2 from 63.7% to 65.7% and increases the significance of the variables that interact conflicting or reinforcing signals with experience. This shows the role that the aggressive use of market orders plays in improving aggregation when signals are reinforcing and reducing aggregation when signals are conflicting.

What explains this apparently perverse adaptation? Although more aggressive informed trading with more extreme signals harms informational efficiency when signals are in conflict, it improves profits by 12% and increases convergence when signals are reinforcing. We have seen that expected profits are driving market orders and that this doesn't decline with time. When an informed trader holds an extreme signal, aggressive trading at prices between the signal and the unconditional expectation often results in a number of small losses and inefficient pricing if signals are in conflict, but large profits and more efficient pricing if ex post, the signals are revealed to be reinforcing. Together these findings mean that when signals are in conflict, the relatively small realized losses from market orders (expected to generate profits) from traders with an extreme signal are insufficient to make the traders update their information and so market orders that then reduce aggregation continue apace. By contrast, when signals are reinforcing, the relatively large profits realized from market orders for traders with an extreme signal encourage an updating of information and subsequent market orders improve aggregation. That there is no significant difference in the profits of informed traders holding extreme signals as a function of the level of experience, suggests that these traders are consistently responding to large profits but not to small losses, and are willing to trade-off the small losses against the large profits.

It is possible that the failure of aggregation could be further impeded if liquidity dried up before the end of the trading period. In trading periods in which the signals are reinforcing (which have the highest price errors), there is an outstanding quote that could have been transacted against profitably with a market order at the end of 97% of the trading periods. On average, market orders against these quotes would have provided a profit of L$5.14 (Fig. 2 Panel C). In the appendix (Fig. A-2, Panel B) we show this result obtains in all sessions. This compares with an average loss of L$0.23 on the 2,541 trades in which informed traders submit market orders. Also, the average bid-ask spread at the end of the trading period (L$1.54) is significantly less than average level over the trading period (L$2.35), and not significantly different from the average spread at the time of all transactions (L$1.53).[35] We therefore conclude the failure of aggregation is not due to a breakdown in liquidity at the end of the period and we reject H2 (liquidity declines at the end of the trading period).

3.4.2. *Prediction behavior in the prediction markets*

In Table 4, we report results pertaining to prediction behavior. These results show the frequency and way in which predictions deviate from the RNBNE predictions. In Panel A, we show that under the assumption that the previous prediction was an equilibrium prediction, 24% of the predictions with inexperienced subjects are consistent with the RNBNE, rising to 32% for experienced subjects. Predictions consistent with the RNBNE are most frequent on the first prediction in a period (48% with inexperienced and 63% with experienced subjects), and least frequent on the final prediction (13% and 10% respectively). Panel A also shows that both inexperienced and experienced subjects are 1.8 times more likely to follow the equilibrium strategy when they receive an informative signal (recall that in each period, one subject does not receive a signal, and matching the previous estimate ensures a payoff of zero). This difference is significant ($p<0.01$, $N=14$).

Deviations from the RNBNE significantly affect convergence to intrinsic value. We show this by adding a variable to the Table 2, Panel C regression equal to the percentage of predictions in a period inconsistent with the equilibrium prediction under the assumption that the previous prediction was an equilibrium prediction (NOT EQ %). This variable is significant ($p<0.01$) and increases the R^2 of the regression from 25.6% to 32.1%.

Given the preponderance of estimates that do not conform to the RNBNE, in Table 4, Panel B we report signed deviations from equilibrium predictions (SDEQ) by level of experience, prediction order, and whether a subject received a signal. In calculating these averages, predictions closer to the unconditional expectation than the RNBNE prediction receive negative values, while predictions more extreme than the RNBNE are signed positive. First predictions in each period tend to be more extreme than the equilibrium prediction with both inexperienced and experienced subjects, although not significantly in either case. This behavior however depends on whether a subject received a signal. While subjects that receive a signal tend to predict closer to the unconditional expectation than consistent with the signal, subjects that do not receive a signal make estimates that are more extreme than the unconditional expectation (which is the estimate that maximizes expected profit and minimizes the variance of profit).[36] On average, subsequent predictions deviate significantly toward the unconditional expectation with both experience levels. Overall, weighting each session average as a single observation, SDEQs are negative and significant for each level of experience ($p=0.02$ and $p=0.06$ for sessions with inexperienced and experienced subjects respectively). The high frequency with which subjects deviate from the equilibrium strategy leads to significant price errors at the end of each prediction period, as shown in Table 1, Panels D to F.

We explore the determinants of the systematic deviations from the RNBNE using logistic regressions with the indicator variables defined in section 2.E. as dependent variables: Closer-to-Previous-Prediction, Closer-to-L$100, and Wrong Direction. Summarizing the main results: 1) the RNBNE is less likely when signals are reinforcing ($p<0.01$); 2) predictions closer to the previous prediction than the RNBNE are more frequent when the RNBNE implied by the new signal is more extreme than the previous equilibrium ($p<0.01$) and when the signal is extreme relative to the unconditional expectation ($p<0.01$); 3) deviations from the RNBNE toward the unconditional expectation increase in frequency in the extremeness of the RNBNE ($p<0.01$); and 4) predictions in the "wrong" direction conditional on the signal are more likely when the new equilibrium prediction is more extreme than the previous

[34] The larger increase in market orders when signals are conflicting is due to the more frequent availability of putatively profitable market orders when information signals are dispersed around the unconditional expectation.

[35] On average, informed traders earn profits when using either market or limit orders in trades with liquidity traders, but they lose money when demanding liquidity from other informed traders. Overall, liquidity traders lose money.

[36] This result is quite general. This same pattern obtains in the subset of 14 subjects that predict first multiple times, sometimes receiving an informative signal and other times not.

Table 4
Deviations from Equilibrium in the Prediction Markets.

Number of Sessions	Experience Level	Number of Predictions	Prediction Order				
			1	2	3	4	All
Panel A. Proportion of predictions consistent with equilibrium							
8	Inexp. Sessions (All)	320	0.48	0.24	0.13	0.13	0.24
	No Signal		0.38	0.12	0.12	0.04	0.15
	Informative Signal		0.50	0.29	0.13	0.16	0.27
6	Exp. Sessions (All)	240	0.63	0.20	0.33	0.10	0.32
	No Signal		0.33	0.22	0.25	0.06	0.20
	Informative Signal		0.71	0.19	0.35	0.12	0.36
Panel B. Signed deviation from equilibrium prediction (SDEQ)							
8	Inexp. Sessions (All)	320	0.09	-0.80***	-0.61	-0.86**	-0.55**
			(0.31)	(-4.62)	(-1.57)	(-2.97)	(-3.07)
	No Signal	80	1.31***	-1.25*	-1.00	0.29	-0.23
			(3.59)	(-1.94)	(-1.05)	(0.50)	(-0.57)
	Informative Signal	240	-0.22	-0.61**	-0.52	-1.36***	-0.65***
			(-0.72)	(-2.84)	(-1.48)	(-4.69)	(-4.01)
6	Exp. Sessions (All)	240	0.20	-1.43***	-0.50**	0.18	-0.39*
			(0.59)	(-7.51)	(-2.31)	(0.54)	(-2.45)
	No Signal	60	1.83*	-1.22**	0.00	0.72	0.22
			(2.48)	(-3.20)	(0.00)	(0.83)	(0.56)
	Informative Signal	180	-0.21	-1.52***	-0.63**	-0.05	-0.59***
			(-0.73)	(-5.54)	(-3.95)	(-0.11)	(-5.81)

When conducting statistical tests averages from each session are treated as a single observation. Significance for a two-tailed test relative to zero at the 10%, 5%, and 1% level are indicated with *,**, and *** respectively with t-statistics in parentheses.

equilibrium ($p<0.01$), but less likely when the signal itself is extreme ($p<0.01$). We report complete results in the appendix (Table A-4). Fig. A-2, Panel C shows the preponderance of negative SDEQs when signals are reinforcing.

In the CDA markets, we saw evidence that traders were changing their order choice as trading progressed in response to their assimilation of information from transaction prices. This meant that putatively unprofitable trades conditional on their signal become more frequent later in the trading period. We now consider whether similar learning activity could play a role in prediction strategies in the PM markets. On average traders earned significant profits from their predictions, averaging L $5.22 on each prediction in sessions with inexperienced subjects, and L $5.31 in experienced sessions. Deviations from equilibrium play however are costly, with per prediction profits averaging L$5.77 less with inexperienced subjects ($p<0.01$) and L$4.72 less with experienced subjects ($p<0.01$) than they would have been if all predictions had been consistent with the RNBNE. In each case, the standard deviation of profits along the equilibrium path would also have been significantly lower: by 29% with inexperienced subjects and 16% with experienced subjects, with both differences significant at the 1% level.

Since out-of-equilibrium play predominates, we also test whether subject behavior is adaptive by comparing actual profits with what profits would have been if each subject had made the prediction consistent with the RNBNE, under the assumption that prior predictions were equilibrium predictions. Again, foregone profits due to not playing the equilibrium strategy are significant, averaging L$7.79 ($p<0.01$) and L$3.85 ($p=0.03$) per prediction with inexperienced and experienced subjects respectively. The standard deviation of profits is also higher: 19%, ($p<0.01$), and 8%, ($p=0.15$) respectively. We report complete results in the appendix (Table A-5).

In sum, PM predictions do not follow the equilibrium path. Experience in a previous session only implies a marginally significant increase in equilibrium play. Although aggregation relative to the average signal is higher when signals are reinforcing, the rate of equilibrium play is lower. When either a predictor's signal or the previous prediction is extreme, predictions are closer to the previous prediction than consistent with equilibrium. An extreme previous prediction also leads to a subsequent prediction closer to the unconditional expectation than consistent with the RNBNE. Overall, these patterns lead to incomplete aggregation and we reject H4 (complete aggregation in the PM).

In the CDA we observe the aggressive use of market orders by traders with more extreme signals that explained the incomplete aggregation. By contrast, subjects in the PM that received more extreme signals appear to act more conservatively, predicting closer to the previous prediction than the equilibrium prediction. This apparent contrast may be explained by how the mechanisms shape behavior. [37] In the CDA, subjects with extreme signals trade aggressively when the price is close to the unconditional expectation. With an extreme signal and experience, informed traders typically make multiple trades, some at more aggressive prices than others, but on average, not moving price up to their signals. Subjects are willing to trade aggressively near the E[V|PI] because they have a margin of safety built in. In the PM in contrast, there is only one way to create a margin of safety: by predicting closer to the previous signal than suggested by the RNBNE.

4. Concluding comments

Our primary result is that the ex post relationship between independent signals determines the level of information aggregation, and more generally the time series properties of asset prices. When private information is reinforcing, aggregation improves significantly with experience in previous sessions but information aggregation is incomplete. At the end of the trading (or prediction) period, less than half of private information impounded in price. When information is in conflict, aggregation becomes worse with experience in a previous session. In about one-fifth of these periods, ending prices are less efficient than the unconditional expectation, implying that attempts to use private information in the pursuit of profit reduces informational efficiency and increases volatility with respect to fundamentals.

We show that in the continuous double auction sessions, the aggressive use of market orders by experienced traders with signals that imply large revisions in their conditional expectations facilitates aggregation when information signals are reinforcing, but impedes aggregation when signals are conflicting. In the prediction market sessions,

[37] We thank an anonymous reviewer for inviting us to consider this contrast between the CDA and PM.

convergence is hindered by the propensity of subjects to predict closer to the unconditional mean or to the previous signal than is consistent with a risk neutral Bayesian-Nash equilibrium, a strategy which is consistent with an attempt to accept lower profits for reduced variance of profits.[38]

Our subjects are sophisticated (including MS finance students), in many cases receive experience in multiple sessions, and interact under a stationary information structure: a feature absent from field markets that likely impedes aggregation. Under the assumption that each piece of private information eventually becomes public, given the empirically robust result that the disclosure of bad news is typically delayed relative to the disclosure of good news, we find patterns consistent with price momentum when signals are reinforcing; and excess volatility, weak negative serial correlation, and inefficient aggregation when signals are in conflict. We use this private information setting to complement the literature that shows that proxies for public disagreement are associated with inefficient information aggregation, but without requiring illiquidity to explain mispricing.

CRediT authorship contribution statement

Charles Schnitzlein: Writing – review & editing, Writing – original draft, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Patricia Chelley-Steeley:** Software, Investigation, Formal analysis. **James M Steeley:** Writing – review & editing, Writing – original draft, Supervision, Software, Project administration, Methodology, Investigation, Formal analysis.

Data availability

Data will be made available on request.

Supplementary materials

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jbankfin.2024.107300.

References

Avery, C., Zemsky, P., 1998. Multidimensional uncertainty and herd behavior in financial markets. Am Econ Rev 88, 724–748.
Back, K., Cao, H., Willard, G., 2000. Imperfect competition among informed traders. J Finance 55, 2117–2155.
Baginski, Stephen P., Campbell, John L., Hinson, Lisa A., Koo, David S., 2018. Do career concerns affect the delay of bad news disclosures. Account. Rev. 93, 61–95.
Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. J financ econ 49, 307–343.
Biais, Bruno, Hilton, Denis, Mazurier, Karine, Pouget, Sebastien, 2005. Judgmental overconfidence, self-monitoring, and trading performance in an experimental financial market. Rev. Econ. Studies 72, 287–312.
Bloomfield, R., O'Hara, M., Saar, G., 2005. The "make or take" decision in an electronic market: evidence on the evolution of liquidity. J financ econ 75, 165–199.
Boehme, Rodney D., Danielsen, Bartley R., Sorescu, Sorin M., 2006. Short-sale constraints, differences of opinion and overvaluation. J. Financ. Quant. Anal. 41, 455–487.
Bossaerts, P., Ghirardato, P., Guarnaschelli, S., Zame, W.R., 2010. Ambiguity in asset markets: theory and experiment. Rev. Financ. Studies 23, 1325–1359.
Brown, Lawrence D., Call, Andrew C., Clement, Michael B., Sharp, Nathan Y., 2015. Inside the "black box" of sell side financial analysts. J. Account. Res. 53, 1–47.
Carlin, Bruce I, Longstaff, Francis A., Matoba, Kyle, 2014. Disagreement and asset prices. J financ econ 114, 226–238.
Cason, TN, 2000. The Opportunity for conspiracy in asset markets organized with dealer intermediaries. Rev Financ Stud 13, 385–416.
Chakravarty, S., Holden, C., 1995. An integrated model of market and limit orders. J. Financ. Intermediation 4, 213–241.
Charness, Gary, Gneezy, Uri, Imas, Alex, 2013. Experimental methods: eliciting risk preferences. J. Econ. Behavior Organization 87, 43–51.
Chen, Y., Dimitrov, S., Sami, R., Reeves, D.M., Pennock, D.M., Hanson, R.D., Fortnow, L., Gonen, R., 2010. Gaming prediction markets: equilibrium strategies with a market maker. Algorithmica 58, 930–969.
Copeland, T.E., Friedman, D., 1987. The effect of sequential information arrival on asset prices: an experimental study. J Finance 42, 763–797.
Corgnet, Brice, Mark DeSantis, and David Porter, 2015, Revisiting Information Aggregation in Asset Markets: Reflective Learning & Market Efficiency, ESI Working Paper 15-15.
Daniel, K., Hirshleifer, D., Subrahmanyam, A, 1998. Investor psychology and security market under and over-reaction. J Finance 53, 1839–1885.
De Jong, Cyriel, Koedijk, Kees G., Schnitzlein, Charles R., 2006. Stock market quality in the presence of a traded option. J. Business 79, 2243–2274.
Diether, Karl B., Malloy, Christopher J., Scherbina, Anna, 2002. Differences of opinion and the cross section of stock returns. J Finance 57, 2113–2141.
Djogbenou, Antoine A., Mackinnon, James G., Nielsen, Mortern Orrgaard, 2019. Asymptotic theory and wild bootstrap inference with clustered errors. J Econom 212, 393–412.
Eyster, Erik, Rabin, Matthew, Vayanos, Dimitri, 2019. Financial markets where traders neglect the informational content of prices. J Finance 74, 371–399.
Fischbacher, Urs, 2007. z-tree: a toolbox for readymade economic experiments. Exp Econ 10 (2), 171–178.
Forsythe, R., Lundholm, R., 1990. Information aggregation in an experimental market. Econometrica 58, 309–347.
Foster, F.D., Viswanathan, S., 1996. Strategic trading when agents forecast the forecasts of others. J Finance 51, 1437–1478.
Foucault, T., 1999. Order Flow composition and trading costs in a dynamic limit order market. J. Financ. Markets 2, 99–134.
Goettler, R., Parlour, C., Rajan, U., 2009. Informed traders and limit order markets. J financ econ 93, 67–87.
Hanson, R., 2003. Combinatorial information market design. Inform. Syst. Front. 5 (1), 107–119.
Healy, Paul J., Linardi, Sera, Richard Lowery, J., Ledyard, John O., 2010. Prediction markets: alternative mechanisms for complex environments with few traders. Manage. Sci. 56, 1977–1996.
Holden, C., Subrahmanyam, A., 1992. Long-lived private information and imperfect competition. J Finance 47, 247–270.
Jian, Lian, Sami, Rahul, 2012. Aggregation and manipulation in prediction markets: effects of trading mechanism and information distribution. Manage. Sci. 58, 123–140.
Kaniel, R., Liu, H., 2006. What orders do informed traders Used. J. Business 79, 1867–1913.
Kocher, Kocher Martin G, Lucks, Konstantin E, Schindler, David, 2019. Unleashing animal spirits: self-control and overpricing in experimental asset markets. Rev Financ Stud 32, 2149–2178.
Kothari, S.P., Shu, Susan, Wysocki, Peter D., 2009. Do managers withhold bad news? J. Account. Res. 47, 241–276.
Kyle, A., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335.
Linardi, Sera, 2017. Accounting for noise in the microfoundation of information aggregation. Games Econ Behav 101, 334–353.
Ostrovsky, Michael, 2012. Information aggregation in dynamic markets with strategic traders. Econometrica 80, 2595–2647.
Page, Lionel, Siemroth, Christoph, 2020. How much information is incorporated into financial asset process? experimental evidence. Rev. Financ. Stud., forthcoming.
Plott, C.R., Sunder, S., 1988. Rational expectations and the aggregation of diverse information in laboratory security markets. Econometrica 56, 1085–1118.
Sadka, Ronnie, Scherbina, Anna, 2007. Analyst disagreement, mispricing, and liquidity. J Finance 62, 2367–2403.
Schnitzlein, CR, 1996. Call and continuous trading mechanisms under asymmetric information: an experimental investigation. J Finance 51, 613–636.
Smith, VL, Suchanek, GL, Williams, AW, 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. Econometrica 56, 1119–1151.

[38] While such behavior contradicts the assumption of risk neutrality, which underpins the theoretical predictions of prediction markets, it is common for subjects to display behavior consistent with risk aversion in an experiment with relatively low stakes, though the way to measure risk across settings and mechanisms is not settled, see Charness, Gneezy, and Imas (2013).