

9-2-2022

Information Aggregation with Heterogeneous Traders

Cary Deck

Chapman University, deck@chapman.edu

Tae In Jun

University of Alabama

Laura Razzolini

University of Alabama

Tavoy Reid

University of Alabama

Follow this and additional works at: https://digitalcommons.chapman.edu/esi_working_papers



Part of the [Econometrics Commons](#), [Economic Theory Commons](#), and the [Other Economics Commons](#)

Recommended Citation

Deck, C., Jun, T. I., Razzolini, L., & Reid, T. (2022). Information Aggregation with Heterogeneous Traders. *ESI Working Paper 22-13*. https://digitalcommons.chapman.edu/esi_working_papers/374/

This Article is brought to you for free and open access by the Economic Science Institute at Chapman University Digital Commons. It has been accepted for inclusion in ESI Working Papers by an authorized administrator of Chapman University Digital Commons. For more information, please contact laughtin@chapman.edu.

Information Aggregation with Heterogeneous Traders

Comments

ESI Working Paper 22-13

Information Aggregation with Heterogeneous Traders

Cary Deck^{1,2}, Tae In Jun¹, Laura Razzolini¹, and Tavoy Reid¹

¹Department of Economics, Finance and Legal Studies, University of Alabama

²Economic Science Institute, Chapman University

September 2, 2022

Abstract

The efficient market hypothesis predicts that asset prices reflect all available information. A seminal experiment reported that contingent claim markets could yield market outcomes consistent with information aggregation when traders hold heterogeneous state-contingent values. However, a recent experiment found the rational expectation model outperformed the prior information and maxi-min models in contingent claim markets when traders hold homogeneous values despite the no trade equilibrium in that setting. But that same study failed to replicate the original result calling into question when, if ever, prices reliably reflect the aggregate information of traders with heterogeneous values. In this paper, we show contingent claim markets can robustly yield prices consistent with the efficient market hypothesis when traders hold heterogeneous values in certain circumstances. The key distinction between our environment and that of the previous studies is that we consider trader values that are correlated and not too dissimilar.

Keywords: Information Aggregation, Rational Expectations, Laboratory Experiments

We thank Brice Corgnet, Mark DeSantis, and Erik Kimbrough for helpful comments.

1 Introduction

Identifying conditions under which the efficient market hypothesis holds and prices reflect the aggregated information of the traders remains a topic of debate in economics and finance.¹ Relying on what became a canonical experimental design, [Plott and Sunder \(1988\)](#) demonstrate the potential for information aggregation to be reflected in prices in certain circumstances including markets in which traders hold *heterogeneous* values for state-contingent assets. However, [Corgnet et al. \(2022\)](#) fail to replicate the [Plott and Sunder \(1988\)](#) results, but [Corgnet et al. \(2022\)](#) do report clear evidence supporting the efficient market hypothesis in state-contingent markets with *homogeneous* traders. Our paper examines information aggregation with state-contingent assets and *heterogeneous* traders in an environment that falls between the two extremes examined previously. Specifically, in our setting the heterogeneous trader values are positively correlated and similarly sized. Ultimately, we find strong evidence in favor of the efficient market hypothesis.

While there have been numerous studies about information aggregation (e.g. [Asparouhova, Bossaerts, and Yang \(2017\)](#), [Page and Siemroth \(2017\)](#), [Choo, Kaplan, and Zultan \(2019\)](#)), as highlighted in the extensive literature review by [Corgnet et al. \(2022\)](#), very few studies after [Plott and Sunder \(1988\)](#) actually offered direct evidence for or against the rational expectations model, which requires alternative hypotheses based on competing models of behavior. The two alternative models considered by both [Plott and Sunder \(1988\)](#) and [Corgnet et al. \(2022\)](#) are the Prior Information (PI) model and the Maxi-min (MM) model. The PI model assumes traders use Bayes' rule to update their beliefs about an asset's expected value based on publicly available prior information and any private information, but do not infer information from market prices. The MM model also assumes that traders rely only on public and private information and do not infer information from market activity, but differs from the PI model in that it assumes traders are only willing to purchase an asset if its price is equal to or below the minimum possible value of the asset given the trader's information. By contrast, under the Rational Expectations (RE) model prices reflect the pooled information of all of the traders.

In their seminal paper [Plott and Sunder \(1988\)](#) consider three distinct combinations of market features (or "series" in their terminology). "Series A" markets are characterized by a single asset and heterogeneous state-contingent values among traders. The evidence presented in [Plott and Sunder \(1988\)](#) indicates the RE model is not a statically better predictor of market outcomes than the PI or MM models for Series A markets. [Plott and Sunder \(1988\)](#) also consider two other series that include features designed to facilitate information aggregation. "Series B" is similar to Series A, but

¹See for example [Fama \(1970, 1991, 2008\)](#), [Shleifer \(2000\)](#), [Thaler \(2005, 2015\)](#), and [Shiller \(2015\)](#).

involves state-contingent securities and “Series C” is similar to Series A, but involves homogeneous state-contingent values. It is these two series for which [Plott and Sunder \(1988\)](#) *do* and [Corgnet et al. \(2022\)](#) *do not* find evidence that market prices reflect aggregated trader information. That is, [Plott and Sunder \(1988\)](#) find statistical evidence that the predictive power of the RE model outperforms both the PI and MM models in Series B and Series C markets, while [Corgnet et al. \(2022\)](#) do not find statistical evidence in favor of the RE model for either of these two series. [Corgnet et al. \(2022\)](#) subsequently introduce “Series D” markets, which combine the state-contingent assets and homogeneous trader values features of Series B and C markets, and it is in these markets where they find compelling statistical evidence in favor of RE over the PI and MM models indicating successful information aggregation.

The demonstrated potential of markets with state-contingent assets and homogeneous trader values to aggregate information is impressive, not only in the laboratory, but in practice as well, as these features are standard in prediction markets.² However, this success is somewhat surprising in light of the no trade prediction highlighted by [Hirshleifer \(1971\)](#) that arises when traders have homogeneous values, at least in the zero *ex-post* risk setting of [Plott and Sunder \(1988\)](#) and [Corgnet et al. \(2022\)](#). In both of their experiments, there are three possible states of the world in which half of the traders are informed of one state that will not be realized and the other half of the traders are informed of the other state that will not be realized. Thus, no individual trader knows which state will be realized, but the aggregate information of the traders perfectly identifies the realized state. In such a setting, if the market aggregates information so that prices exactly reveal state-contingent asset values, there is no incentive for anyone to trade.³ By contrast, the heterogeneous values of Series B markets provide an incentive to trade even when the realized state is identified as there remain direct gains from exchange. Thus, one might expect heterogeneous asset values to encourage more trading than homogeneous asset values and, thus, yield prices that reflect information aggregation.

We conjecture that the apparent superiority of Series D markets over Series B markets in terms of facilitating information aggregation lies with the specific structure of the heterogeneous asset values used in prior studies. In particular, the Series B experiments of [Plott and Sunder \(1988\)](#) and [Corgnet et al. \(2022\)](#) are such that some traders have negatively correlated asset values across states, which increases the difficulty of interpreting the market activity of other traders. With homogeneous asset values, as in Series D, the interpretation of market activity is more straightforward. In this paper, we

²For discussions of the success of prediction markets in aggregating information see [Pennock et al. \(2001\)](#), [Chen and Plott \(2002\)](#), [Tetlock \(2004\)](#), [Gürkaynak and Wolfers \(2006\)](#), [Berg et al. \(2008\)](#), [Cowgill et al. \(2009\)](#), [Cowgill and Zitzewitz \(2015\)](#), and [Atanasov et al. \(2017\)](#).

³In naturally occurring prediction markets, aggregate information is unlikely to identify an outcome with certainty and in such cases differential risk attitudes may provide a motive to trade.

consider the case of heterogeneous, but positively correlated, state-contingent asset values in what we refer to as “Series B’.” These markets combine the incentive to trade of previous Series B markets with the ease of inference of Series D markets. Ultimately, we find strong evidence that Series B’ markets yield prices consistent with the efficient market hypothesis.

2 Experimental Design

2.1 Market Environment and Predictions

The market structure in our Series B’ markets closely follows the Series B markets discussed in [Plott and Sunder \(1988\)](#). As in their setting, our markets involve three types of traders: I, II, and III. Each trader is endowed with money, referred to as francs, and state-contingent assets, termed certificates. During the market, traders can buy and sell assets via a double auction. There are three possible states: X, Y, and Z. s -Certificates pay the holder a dividend of 0 if the state is not s and pay d_i^s to a trader of type i if the state is s for $s \in \{X, Y, Z\}$ and $i \in \{I, II, III\}$. [Table 1](#) gives the parameter values for the state-contingent dividend values for each trader type and is comparable to [Table 1](#) in [Plott and Sunder \(1988\)](#).

The distinction between Series B of [Plott and Sunder \(1988\)](#) and Series B’ lies with the relationship of dividend values across trader types. In Series B of [Plott and Sunder \(1988\)](#), $d_I^X < d_I^Y < d_I^Z$ and $d_{III}^X < d_{III}^Y < d_{III}^Z$, but $d_{II}^X > d_{II}^Y > d_{II}^Z$. That is, while the dividend values of Types I and III traders are perfectly positively correlated, the dividend values of Types I and II and Types II and III are perfectly negatively correlated. By contrast, in our Series B’, $d_i^X < d_i^Y < d_i^Z$ for all i . Thus, in our market environment there is agreement across types as to which states are more valuable. Of course, this agreement exists trivially in the homogeneous values of Series D in [Corgnet et al. \(2022\)](#), where there was a single trader type with $d_I^X < d_I^Y < d_I^Z$.

Other aspects of our experimental markets closely follow [Plott and Sunder \(1988\)](#) and [Corgnet et al. \(2022\)](#). Each state was equally likely to occur in each period and information was distributed so that half of the traders of a given type received each of the possible pieces of information.⁴ Further, there are multiple traders of each type in the market and each trader’s endowment is sufficient to purchase all outstanding certificates at the maximum price any trader would be willing to pay for those certificates. This level of liquidity and the competition between traders of the same type is

⁴If, for example, the state was X, then half of the traders of each type were given a signal that the state was Not Y and the other half were given a signal that the state was Not Z. We note that both [Plott and Sunder \(1988\)](#) and [Corgnet et al. \(2022\)](#) consider markets with equiprobable and non-equiprobable states, but the results reported in those studies suggest this distinction does not impact information aggregation.

theoretically sufficient for market prices to reflect the maximum value belief held by any trader. The endowment was treated as a loan, which had to be repaid at the end of the market period.

The RE, PI, and MM models differ with regards to the beliefs that traders hold and, thus, the prices that are predicted to prevail in the market. Under RE, information is pooled so that a trader of Type i believes s -Certificates are worth d_i^s when the state is s and are worth 0 otherwise. Thus, under RE, the price of s -Certificates equals $\max_i(d_i^s)$ when the state is s and is 0 otherwise. By contrast, under PI, traders use their private information to update their beliefs from the common uniform prior, but do not use market information to update their beliefs. As such, there are always traders of Type i who believe that s -Certificates are worth $d_i^s/2$ and no traders of Type i who believe that s -Certificates are worth more than $d_i^s/2$. Thus, under the prior information model, the price of s -Certificates equals $\max_i(d_i^s/2)$ regardless of the actual state. Finally, under MM, traders value the asset at the minimum amount that it could be worth to them given their private information. Since no trader's private information is sufficient to determine that a certificate has a positive value, under this model the price of s -Certificates equals 0 regardless of the actual state. Table 2 summarizes the price predictions under each model and the resulting asset allocations are given in Table 3. It is worth noting that the rational expectations and maxi-min models only make different predictions for s -Certificates when the state is in fact s and, thus, outcomes in the other markets cannot be used to distinguish between the two models. By comparison, the prior information model makes predictions that differ from the other models for all certificates.

2.2 Laboratory Procedures

Each of our eight sessions involved 12 subjects, who were randomly assigned to trader types such that there were always four traders of each type. The 96 subjects were all undergraduate students at The University of Alabama who had volunteered to participate in research studies at The Interactive Decision Experiment (TIDE) Lab. While some subjects had participated in previous unrelated studies, none had participated in any related studies. After providing informed consent, subjects were taken to a classroom in the lab where they were provided with paper instructions. The instructions, available in the Appendix and similar to those of [Corngnet et al. \(2022\)](#), were read aloud. Before each of the 10 market periods began, each subject was provided with an "Information and Record Sheet." These sheets contained relevant subject-specific information and allowed the subject to track their holdings of certificates and francs. After each period this sheet was checked by a researcher for the correctness and collected. At the end of the session, subjects were paid in private based upon their total experimental earnings and dismissed from the study.

While our experimental procedures closely follow previous studies in general, Table 4 highlights the various differences among the studies. One notable difference is our use of software for presenting market information during the double auction market. Our markets are hand-run with traders raising paddles to be individually recognized by the auctioneer and announcing bids, offers, and acceptances orally, just as in the previous studies. What differed is that rather than bids, asks, and acceptances being written by hand on the board in the front of the room, this information was entered in real time into a computer program by a researcher and concurrently displayed electronically on a screen at the front of the room that was visible to all traders. Following [Corgnet et al. \(2022\)](#), the auctioneer who called upon traders was unaware of the market predictions in order to avoid inadvertent influence on behavior. Because a researcher did not need to physically move and write information on the board, the duration of a decision period was reduced from 7 minutes to 5 minutes. The fact that trade volumes are ultimately greater in our Series B' markets than in previous Series B markets, the computerized display and the reduced duration do not appear to be limiting the effectiveness of the market.

Other differences in our procedures from previous studies include dropping the use of a separate clue sheet for providing signals. Instead, we included the signal directly on the “Information and Record Sheet” for the period. Further, we did not conduct an extensive training on how to interpret the clue sheet or the randomization process for drawing the state. Thus, to the degree that omitting this training increased trader confusion then our results underestimate the ability of Series B' markets to aggregate information. The exchange rate from francs to dollars differed nominally from previous studies, but was set so that the average real earnings would be comparable to previous studies. On average, our subjects earned \$25.64 in addition to a fixed participation payment of \$10 for the two hour session. Finally, the number of periods was reduced to 10. To the degree that experience helps subject traders understand how to interpret the market activity of others, this design feature can be considered to bias our results away from support for the efficient market hypothesis.

2.3 Hypotheses and Performance Criteria

Our main goal is to test whether markets with contingent claim assets in which traders have heterogeneous values are consistent with information aggregation as predicted by the efficient market hypothesis. Formally, we follow [Plott and Sunder \(1988\)](#) and [Corgnet et al. \(2022\)](#) in testing the null hypothesis that the market data are no better described by the RE model than by another model against the alternative hypothesis that the RE model better fits the observed data than the other model. Separate tests are conducted to compare the RE model with both the PI and MM models. As in the previous studies, the statistical analysis is based on Wilcoxon signed rank sum tests using data from

the periods in which a state occurred for the last time. For our markets this is period 8 (for State Z), period 9 (for State Y), and period 10 (for State X). For each model comparison we rely upon all of the performance criteria used by [Plott and Sunder \(1988\)](#) or [Corgnet et al. \(2022\)](#). Here we briefly describe each criterion. The first five pertain to observed prices.

Criterion 1. Mean Absolute Deviation (MAD) is the average difference between the observed prices and the price predicted to hold under a given model. This is perhaps the most common measure of model performance:

$$MAD = \text{average}_{j,t,s} |p_{j,t,s} - m_{t,s}|, \quad (1)$$

where j represents the j^{th} transaction, t denotes a period, $p_{j,t,s}$ corresponds to the price of the j^{th} transaction in period t for security s , and $m_{t,s}$ is the predicted price for security s under model m in period t , with $m \in [RE, PI, MM]$. The lower the MAD, the better a model's prediction fits the data.

Criterion 2. Total Absolute Deviation (TAD) is the total deviation between observed prices and the price predicted to hold under a given model. Thus, TAD complements MAD, which could be impacted by a few large deviations. The lower the value of TAD is, the better a model fits the data:

$$TAD = \sum_{j,t,s} |p_{j,t,s} - m_{t,s}|. \quad (2)$$

Criterion 3. Closing Absolute Deviation (CAD) is the deviation of the final transaction price in a market from the the predicted price. To the extent that prices evolve over the trading horizon, the closing price represents the final beliefs of the traders. The lower the value of the CAD is, the better the model fits the data:

$$CAD = \text{average}_{t,s} |p_{t,s}^* - m_{t,s}|, \quad (3)$$

where $p_{t,s}^*$ is the final transaction price (or closing price) of asset s in period t .

Criterion 4. Log Odds (LO) is computed by linearly regressing observed prices in a session on the predicted prices for the given model and then recovering the log likelihood value under the assumption of normally distributed error terms. LO measures a model's goodness of fit and, thus, the higher the value the better a model fits the data. Like MAD and TAD, LO only depends on the set of prices that are observed in each market and each period. The final price based criterion is more akin to CAD in that it depends on the price sequence.

Criterion 5. Percentage of Convergent Price Changes (PCPC) is computed as the ratio of convergent price changes to the total number of price changes. The $(j + 1)^{\text{th}}$ transaction in a market in a period is considered a convergent price change if $p_{j+1,t,s} \neq p_{j,t,s}$ and $|p_{j+1,t,s} - m_{t,s}| \leq |p_{j,t,s} - m_{t,s}|$. The denominator of PCPC can be less than the number of transactions minus one since $p_{j+1,t,s} = p_{j,t,s}$

is not considered a price change. The greater the percentage of convergent price changes, the better a model is considered to fit the data.

Beyond price, each model makes predictions about which types of traders will purchase different assets and how much each type of trader will earn given their private information. The final two criteria concern the holdings of each trader. However, all three models under consideration are silent as to how shares should be allocated among traders of the same type and with the same private information. As in the Series B markets of [Plott and Sunder \(1988\)](#) and [Corgnet et al. \(2022\)](#), in our markets exactly two of the twelve subjects have each of the six combinations of types and possible pieces of private information.

Criterion 6. Allocation Flow Percentage (Flow) is the percentage of the available assets, net of their own endowments, that the set of traders predicted to purchase the assets actually acquire. In Series B or B' markets, the MM model makes no prediction about which traders will acquire which assets as all traders hold the identical belief that each asset's value is zero. Thus, this criteria can only be used to compare the RE and PI models. For those two models, the larger the flow, the better the model's fit with the observed data.

Criterion 7. Sum of Squared Profit Deviation (SSPD) is based on the squared difference between the predicted average profit of a given trader type with a given piece of private information and the actual average profit of such trader. This criteria is measured by the following equation:

$$SSPD = \underset{t}{average} \left(\frac{\sum_i \sum_{\rho} 2 (\bar{\Pi}_{i,\rho,t,m,s} - \bar{\pi}_{i,\rho,t})^2}{12} \right), \quad (4)$$

where $\rho \in P$ represents the set of possible pieces of private information and Π and π denote the predicted and realized payoff for a given combination of i and ρ . While allocation flow only depends on who purchases the assets and not the prices, the sum of squared profit deviations depends both on who traded and the prices. The smaller the value of the sum of squared deviations, the better the model's fit.

3 Behavioral Results

This section closely follows both [Plott and Sunder \(1988\)](#) and [Corgnet et al. \(2022\)](#). We first present Figures 1, 2, and 3. These figures show the chronological sequence of transaction prices for all eight sessions over the ten market periods for Assets X, Y, and Z, respectively. The figures also provide the predicted asset price for each model for each period given the realized state. While it is difficult to follow the sequence of transaction prices in any specific session, what is quite clear

from the figures is that prices are concentrated near the rational expectations predictions each period, especially in the later portions of each session.

To formally evaluate the degree to which the markets aggregate the disparate information held by the traders, we rely upon the statistical analysis presented in Tables 5 and 6. For all seven criteria, the data are better described by the RE model than the PI model and the difference is highly statistically significant in each case. For the comparison between the RE and MM models, the analysis can only be conducted using six of the criteria, as the MM model makes no prediction regarding who should hold the assets. In five of the six possible comparisons between RE and MM, the data supports the RE model and the statistical results are highly significant. This leads to our main finding:

Main Finding: When traders have heterogeneous, but positively correlated values, then state-specific contingent claim securities facilitate information aggregation.

The only criterion over which RE does not outperform MM is the sum of squared profit deviations. However, this calculation for the MM model assumes that on average each person will hold an average number of shares. As can be seen in Table 6, the traders predicted to purchase shares under RE are acquiring approximately 40% of the outstanding assets. That these traders are not able to acquire all of the assets may explain why the RE model does not outperform the MM model in terms of profit deviation. Figure 4 shows that trades occurred throughout the trading period and, thus, it is possible that with a longer trading period even more of the assets would have been acquired by the traders with the highest value as predicted by RE. Still, we observed on average almost 28 trades per period and this is more than double the average trade volume observed by [Corgnet et al. \(2022\)](#) in Series B markets and about four times the volume they observed in Series D markets despite their periods having lasted two minutes longer.

While the evidence presented in Tables 5 and 6 is compelling, we calculate the Post-Study Probability (hereafter PSP), as suggested by [Maniadis, Tufano, and List \(2014\)](#), to help the reader assess the probability of our declared research finding being true. The PSP is defined as follows:

$$PSP = \frac{(1 - \beta)\pi}{(1 - \beta)\pi + \alpha(1 - \pi)}, \quad (5)$$

where π represents one's prior belief of the probability of alternative hypothesis being true, α reflects the Type I error rate, and $1 - \beta$ reflects the statistical power of the test. Consistent with [Corgnet et al. \(2022\)](#), we focus on MAD, as this is arguably the most common measure of performance among the criteria we consider.⁵ According to Table 5, the p-values for the Wilcoxon Signed Ranked Sum

⁵Each of the other criteria for which statistically significant evidence in favor of RE was found yield similar PSP values to MAD.

Tests for both RE vs PI and RE vs MM are 0.01172. We first convert this p-value of 0.01172 into the corresponding z-score of 2.27 and then use G*Power software (Faul et al., 2009), to derive $\beta = 0.0000523$ under the condition that $\alpha = 0.05$ and a sample size = 8. Table 7 provides the resulting PSPs under various priors one might hold. The PSP that the RE hypothesis outperforms the other models is quite high unless one holds a very low prior. We also provide a more conservative approach to calculating the PSP by assuming the actual effect size to be only $\frac{2}{3}$ of the observed effect size. These conservative values are also reported in Table 7 and again suggest that the evidence for information aggregation in our setting is quite strong.

4 Conclusion

Seminal experiments by Plott and Sunder (1988) demonstrated that contingent claim markets are capable of yielding outcomes consistent with the efficient market hypothesis. However, recent research by Corgnet et al. (2022) has called into question the reliability of information aggregation in such markets when the traders do not hold homogeneous values. While the results of Corgnet et al. (2022) convincingly show that information aggregation is not robust for the specific value environments studied by Plott and Sunder (1988), the results of our experiments clearly reveal that reliable information aggregation can arise when traders have heterogeneous values under certain circumstances.

The critical difference between our heterogeneous environment and that of Plott and Sunder (1988) is that in our setting that trader values are positively correlated and not too dissimilar, but not in their setting. This is a subtle, but critical distinction. First, it indicates that the domain of markets over which the efficient market hypothesis holds is larger than what is identified by Corgnet et al. (2022), albeit still smaller than the domain implied by Plott and Sunder (1988). Second, it suggests that market designers should consider introducing correlated heterogeneity into the market to facilitate information aggregation, as this encourages trading and eliminates the no trade equilibrium that arises with homogeneous values. Of course, further research is necessary to more precisely identify the boundary between heterogeneous value structures that are and are not likely to lead to outcomes consistent with the efficient market hypothesis.

References

- Asparouhova, E., Bossaerts, P., & Yang, W. (2017). “Costly information acquisition in decentralized markets: An experiment”. *Available at SSRN 3079240*.
- Atanasov, P., Rescober, P., Stone, E., Swift, S. A., Servan-Schreiber, E., Tetlock, P., . . . Mellers, B. (2017). “Distilling the Wisdom of Crowds: Prediction Markets vs. Prediction Polls”. *Management Science*, 63(3), 691–706.
- Berg, J., Forsythe, R., Nelson, F., & Rietz, T. (2008). “Results From a Dozen Years of Election Futures Markets Research”. *Handbook of Experimental Economics Results*, 1, 742–751.
- Chen, K., & Plott, C. (2002). “Information Aggregation Mechanisms: Concept, Design and Implementation for a Sales Forecasting Problem”. *Social Science Working Paper 1131*, California Institute of Technology.
- Choo, L., Kaplan, T., & Zultan, R. (2019). “Information Aggregation in Arrow-Debreu Markets: An Experiment”. *Experimental Economics*, 22(3), 625–652.
- Corgnet, B., Deck, C., DeSantis, M., Hampton, K., & Kimbrough, E. (2022). “When Do Security Markets Aggregate Dispersed Information?” *Management Science*, Forthcoming.
- Cowgill, B., Wolfers, J., & Zitzewitz, E. (2009). “Using Prediction Markets to Track Information Flows: Evidence from Google”. *AMMA*, pp. 3.
- Cowgill, B., & Zitzewitz, E. (2015). “Corporate Prediction Markets: Evidence from Google, Ford, and Firm X”. *Review of Economic Studies*, 82(4), 1309–1341.
- Fama, E. (1970). “Efficient Capital Markets: A Review of Theory and Empirical Work”. *Journal of Finance*, 25(2), 383–417.
- Fama, E. (1991). “Efficient Capital Markets: II”. *Journal of Finance*, 46(5), 1575–1617.
- Fama, E. (2008). “A Brief History of the Efficient Market Hypothesis”. American Finance Association speech.
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A. G. (2009). “Statistical power analyses Using G*Power 3.1: Tests for correlation and regression analyses”. *Behavior Research Methods*, 41, 1149–1160.
- Gürkaynak, R., & Wolfers, J. (2006). “Macroeconomic Derivatives: An Initial Analysis of Market-based Macro Forecasts, Uncertainty, and Risk”. *NBER Working Paper 11929*.
- Hirshleifer, J. (1971). “The Private and Social Value of Information and the Reward to Inventive Activity”. *American Economic Review*, 61(4), 561–574.
- Maniadis, Z., Tufano, F., & List, J. (2014). “One Swallow Doesn’t Make a Summer: New Evidence on Anchoring Effects”. *American Economic Review*, 104(1), 277–290.

- Page, L., & Siemroth, C. (2017). “An Experimental Analysis of Information Acquisition in Prediction Markets”. *Games and Economic Behavior*, 101, 354–378.
- Pennock, D., Lawrence, S., Lee, C., & Nielsen, F. (2001). “The Real Power of Artificial Markets”. *Science*, 291, 987–988.
- Plott, C. R., & Sunder, S. (1988). “Rational Expectation and the Aggregation of Dispersed Information in Laboratory Security Markets”. *Econometrica*, 56(5), 1085–1118.
- Shiller, R. (2015). *“Irrational Exuberance”* (3rd ed.). Princeton University Press.
- Shleifer, A. (2000). *“Clarendon Lectures: Inefficient Markets”*. Oxford University Press.
- Tetlock, P. (2004). “How Efficient are Information Markets? Evidence from an Online Exchange”. *Harvard University*, Working Paper.
- Thaler, R. (2005). *“Advances in Behavioral Finance”*. Princeton: Princeton University Press.
- Thaler, R. (2015). *“Misbehaving: The Making of Behavioral Economics”*. New York: W. W. Norton and Company.

Table 1: Market Parameters

Market (Series)	Trader Type	Number of Traders	Initial Endowment		Fixed Cost	Exchange Rate	Dividends			Probabilities		
			Certificates	Francs			X	Y	Z	X	Y	Z
Series B'	I	4	2	10,000	10,000	0.008	70	140	320	1/3	1/3	1/3
	II	4	2	10,000	10,000	0.008	80	150	300	1/3	1/3	1/3
	III	4	2	10,000	10,000	0.008	90	170	330	1/3	1/3	1/3

Table 2: Price Predictions

Market	Model	True State		
		X	Y	Z
X	RE	90	0	0
	PI	45	45	45
	MM	0	0	0
Y	RE	0	170	0
	PI	85	85	85
	MM	0	0	0
Z	RE	0	0	330
	PI	165	165	165
	MM	0	0	0

Table 3: Allocation Predictions

Market	Model	True State		
		X	Y	Z
X	RE	III	-	-
	PI	III	III (Not Z)	III (Not Y)
	MM	No predictions		
Y	RE	-	III	-
	PI	III (Not Z)	III	III (Not X)
	MM	No predictions		
Z	RE	-	-	III
	PI	III (Not Y)	III (Not X)	III
	MM	No predictions		

Table entry identifies Trader Type (signal) of traders who should hold the certificate according to a given model.

Table 4: Differences Among Experiments

	Plott and Sunder (1988)	Corgnet et al. (2022)	Our study
Wording for Market Period	Year	Year	Period
Periods per Session	13-16 periods	13 periods	10 periods
Time per Period	7 minutes	7 minutes	5 minutes
Display of Bids and Asks at Front of Room	Board	Board	Computer Projection
Participant pool	Caltech & IIM	Chapman University	University of Alabama
Simple calculator provided?	None	Yes	Yes
Exchange rate	0.0025-0.003	0.005-0.006	0.008
Auctioneer Uninformed of Model Predictions?	Not Reported	Yes	Yes
Use of Separate Clue Sheet and Clue Sheet Training	Yes	Yes	No

Note: Details from prior studies are only based on Series B and D markets in those studies.

Table 5: Comparison of Price Criteria Between Models

Market Series	Session	MAD			TAD			CAD			LO			PCPC		
		PI	RE	MM	PI	RE	MM	PI	RE	MM	PI	RE	MM	PI	RE	MM
B'	1	82.33	18.05	129.18	3211	704	5038	89.43	7.00	79.57	-232.21	-164.38	-255.03	31.58	68.42	42.11
	2	64.09	55.23	144.32	2820	2430	6350	77.86	18.57	92.86	-250.25	-255.60	-288.27	38.46	65.38	61.54
	3	83.41	15.65	159.32	3086	579	5895	90.33	14.67	107.33	-210.42	-148.00	-246.10	25.00	75.00	25.00
	4	86.65	18.68	128.87	2686	579	3995	79.00	26.00	90.50	-186.25	-149.11	-202.76	50.00	56.25	62.50
	5	85.49	12.68	163.90	3505	520	6720	75.00	22.86	104.29	-225.03	-177.32	-274.54	52.38	52.38	52.38
	6	87.41	11.95	159.13	3409	466	6206	102.00	3.00	99.00	-221.80	-165.40	-260.31	28.57	71.43	35.71
	7	81.61	45.00	153.57	4570	2520	8600	82.14	14.29	90.00	-325.91	-314.04	-370.38	37.93	79.31	62.07
	8	60.69	35.94	74.69	2913	1725	3585	85.89	20.22	70.22	-268.64	-254.51	-286.57	20.00	76.00	68.00
Wilcoxon Signed Ranked Sum Tests																
RE vs PI		0.01172(RE)			0.01162(RE)			0.01172(RE)			0.01729(RE)			0.01796(RE)		
RE vs MM		0.01172(RE)			0.01172(RE)			0.01172(RE)			0.01172(RE)			0.04252(RE)		

Table 6: Comparison of Non-Price Criteria Between Models

Market		Flow			SSPD		
Series	Session	PI	RE	MM	PI	RE	MM
B'	1	3.57	25.00	No Prediction	7508.94	21.90	9.68
	2	7.14	33.33	No Prediction	7488.75	362.05	359.18
	3	7.74	66.67	No Prediction	7406.33	15.70	14.25
	4	3.57	22.92	No Prediction	7511.75	25.17	16.28
	5	4.76	41.67	No Prediction	7399.14	13.74	10.81
	6	1.79	50.00	No Prediction	7424.24	18.33	7.21
	7	12.50	43.75	No Prediction	7448.11	202.98	209.71
	8	10.71	29.17	No Prediction	7045.55	61.73	65.98
Wilcoxon Signed Ranked Sum Tests							
RE vs PI		0.01172(RE)			0.01172(RE)		
RE vs MM		Undefined			0.2076(MM)		

Table 7: Post-Study Probability MAD is Lower Under RE than PI or MM

Probability RE outperforms PI			Probability RE outperforms MM		
Prior	PSP	Conservative PSP	Prior	PSP	Conservative PSP
0.99	0.9995	0.9995	0.99	0.9995	0.9995
0.95	0.9974	0.9973	0.95	0.9974	0.9973
0.9	0.9945	0.9944	0.9	0.9945	0.9944
0.75	0.9836	0.9833	0.75	0.9836	0.9833
0.5	0.9524	0.9515	0.5	0.9524	0.9515
0.25	0.8696	0.8673	0.25	0.8696	0.8673
0.1	0.6896	0.6855	0.1	0.6896	0.6855
0.05	0.5128	0.5080	0.05	0.5128	0.5080
0.01	0.1681	0.1654	0.01	0.1681	0.1654

PSP is calculated assuming the true effect size equals the observed effect size while the conservative PSP is calculated assuming the true effect size is $\frac{2}{3}$ of the observed effect size.

Figure 1: Sequences of Transaction Prices - Asset X




Figure 2: Sequences of Transaction Prices - Asset Y




Figure 3: Sequences of Transaction Prices - Asset Z





Figure 4: Cumulative Percentage of Trades vs Percentage of Time Elapsed



Appendix

The appendix consists of three parts. The first part is the instructions as read aloud to the traders. The second part is an example Information and Record Sheet that is consistent with the example given in the instructions. A subject was given an Information and Record Sheet at the start of each period. The rows of those sheets were blank except for rows 0 and 21. The sheet also showed the subject's state contingent values, which did not change period to period, and the subject's private information, which did change period to period. The third part is the Profit Sheet that each subject received.

Instructions

General. – This is an experiment in the economics of market decision making. The instructions are simple, and if you follow them carefully and make good decisions, you might earn a considerable amount of money which will be paid to you in cash.

In this experiment we are going to simulate a market in which you and the other investors will buy and sell certificates in a sequence of market periods. Attached to the instructions you will find a sheet, labeled information and record sheet, which helps determine the value to you of any decisions you might make. You are not to reveal this information to anyone. It is your own private information.

The type of currency used in this market is francs. All trading and earnings will be in terms of francs. Each franc is worth _____ to you. Do not reveal this number to anyone. At the end of the experiment your francs will be converted to dollars at this rate, and you will be paid in dollars. Notice that the more francs you earn the more dollars you earn.

Specific Instructions

Your profits come from two sources—from collecting certificate earnings on all certificates you hold at the end of the period and from buying and selling certificates. During each market period you are free to purchase or sell as many certificates as you wish, provided you follow the rules below. There are three types of certificates: x-certificate, y-certificate, and z-certificate. For each certificate of a given type that you hold at the end of the period you will be given either the number of francs listed on row 19 of your information and record sheet for that certificate type or 0. Note that earnings may be different for different investors. The method by which one of the two numbers for each certificate type is selected each period is explained later in these instructions. Compute your total certificate earnings for a period by multiplying the earnings per certificate of a given type by the number of certificates of that type held and then summing these three amounts. That is, $(\text{number of x-certificates held}) \times (\text{earnings per x-certificate}) + (\text{number of y-certificates held}) \times (\text{earnings per y-certificate}) + (\text{number of z-certificates held}) \times (\text{earnings per z-certificate}) = \text{total certificate earnings}$. Suppose, for example, that you hold five x-certificates at the end of period 1. If for that period your earnings are _____ francs per x-certificate (i.e., the number for x-certificates from row 19 is _____) then your total x-certificate earnings in the period would be $5 \times _____ = _____ \text{ francs}$. Suppose you hold 4 y-certificates and 3 z-certificates at the end of period 1. If for that period your earnings are 0 francs per certificate for y-certificates and z-certificates then your total y-certificate earnings in the period would be $4 \times 0 = 0 \text{ francs}$ and your total z-certificate earnings in the period would be $3 \times 0 = 0 \text{ francs}$. Your total certificate earnings would be $_____ + 0 + 0 = _____ \text{ francs}$. This number should be recorded on row 19 at the end of the period.

Sales from your certificate holdings increase your francs on hand by the amount of the sale price. Similarly, purchases reduce your francs on hand by the amount of the purchase price. Thus you can gain or lose money on the purchase and resale of certificates. At the end of each period all your holdings are automatically sold to the experimenter at a price of 0. At the beginning of each period you are provided with an initial holding of certificates. This is recorded on row 0 of the period's information and record sheet. You may sell these if you wish or you may hold them. If you hold a certificate, then you receive "earnings per certificate" at the end of the period.

In addition, at the beginning of each period you are provided with an initial amount of francs on hand. This is also recorded on row 0 of each period's information and record sheet. You may keep this if you wish or you may use it to purchase certificates.

Thus at the beginning of each period you are endowed with holdings of certificates and francs on hand. You are free to buy and sell certificates as you wish according to the rules below. Your francs on hand at the end of a period are determined by your initial amount of francs on hand, earnings on certificate holdings at the end of the period, and by gains and losses from purchases and sales of certificates. All francs on hand at the end of a period in excess of the initial amount of francs are yours to keep. These are your profits for the period.

Information about Dividends

Whether the dividend you receive from the certificates of a given type you hold is the amount shown on row 19 or 0 depends on the draw of an outcome from a bingo cage. Before the study began, for each period, we drew a ball from a bingo cage containing thirty balls numbered one through thirty. If the ball drawn was numbered one through ten, the outcome of the draw is called x; if a ball numbered eleven through twenty was drawn, the outcome is called y; if a ball numbered twenty-one through thirty was drawn, the outcome is called z.

If the outcome is x then x-certificates receive the x-dividend shown on row 19 while y-certificates and z-certificates each earn 0. If the outcome is y then y-certificates receive the y-dividend shown on row 19 while x-certificates and z-certificates each earn 0. If the outcome is z then z-certificates receive the z-dividend shown on row 19 while x-certificates and y-certificates each earn 0. Each period's outcome will be announced at the end of each period.

As a reminder, each period there is a $1/3$ chance that the outcome is x; a $1/3$ chance that the outcome is y; and a $1/3$ chance that the outcome is z.

At the beginning of each period, before trading starts, each investor will receive a clue at the top of your information and record sheet. Do not reveal this clue to anyone.

If the outcome is x then there are two possible clues: not y and not z

If the outcome is y then there are two possible clues: not x and not z

If the outcome is z then there are two possible clues: not x and not y

Which of the two possible clues that you receive was determined randomly. The clues that other investors receive are also determined randomly. But the clues are distributed in such a way that exactly half of all investors receive one of the two possible clues and the other half receive the other possible clue.

Trading and Recording Rules

- 1) All transactions are for one certificate of a given type at a time. After each of your sales or purchases you must record the TRANSACTION PRICE in the appropriate column depending on the nature of the transaction. The first transaction is recorded on row 1, and succeeding transactions are recorded on subsequent rows.
- 2) After each transaction you must calculate and record your new holdings of certificates and your new francs on hand. Your holdings of certificates of any type may go below zero. Your francs on hand may never go below zero.
- 3) At the end of the period, record your total certificate earnings in the last column of row 19. If you have negative certificate holdings of a type, you receive no dividends for that type, and you must pay a penalty of 300 francs plus the highest transaction price for that type during the period, for each certificate you are short. Compute your end of period totals on row 20 by listing certificate holdings and adding total certificate earnings to your francs on hand.
- 4) At the end of the period, subtract from your francs on hand the amount listed in row 21 and enter this new amount on row 22. This is your profit for the market period and is yours to keep. At the end of each market period, record this number on your profit sheet.
- 5) At the end of the experiment add up your total profit on your profit sheet and enter this sum on row 22 of your profit sheet. To convert this number into dollars, multiply by the number on row 23 and record the product on row 24. The experimenter will pay you this amount of money.

Market organization. – The market for these certificates is organized as follows. The market will be conducted in a series of periods. Each period lasts for five minutes. Anyone wishing to purchase a certificate of a given type is free to raise his or her hand and make a verbal bid to buy one certificate of a given type at a specified price, and anyone with certificates of the given type to sell is free to accept or not accept the bid. Likewise, anyone wishing to sell a certificate of a given type is free to raise his or her hand and make a verbal offer to sell one certificate of a given type at a specified price. If a bid or offer is accepted, a binding contract has been closed for a single certificate, and the contracting parties will record the transaction on their information and record sheets. Any ties in bids or acceptance will be resolved by random choice. Except for the bids and their acceptance, you are not to speak to any other subject. There are likely to be many bids that are not accepted, but you are free to keep trying. You are free to make as much profit as you can.

Demonstration of Market

We will now go through a demonstration of how the market will work.

The auctioneer will announce the market is open.

Suppose that Trader 13 wants to offer to buy an X certificate at a price of 25. Trader 13 would raise her paddle and the auctioneer would call on her saying something like “Trader 13.” Trader 13 would then respond with “Bid on X of 25.” The auctioneer will repeat the order: “Trader 13 Bid on X of 25.” This will be reflected on the screen. [The order is displayed on the screen.]

Now suppose that Trader 14 wants to offer to sell an X certificate at a price of 500. Trader 14 should raise his paddle. When the auctioneer calls on Trader 14, Trader 14 should announce “Ask for X of 500.” The auctioneer will repeat “Trader 14 Ask for X of 500.” This will then be reflected on the screen. [The order is displayed on the screen.]

This process will continue as long as someone wants to improve on the currently standing bid or ask or accept the standing bid or ask. For example, Trader 13 may raise her paddle and after being called state “Bid on X of 100.” The auctioneer will repeat “Trader 13 Bid on X of 100” and this will be reflected on the screen by displacing the previous smaller bid of 25. [The order is displayed on the screen.] Notice that you can improve upon your own bid or ask.

If Trader 15 is willing to sell X at 100, the current standing bid, Trader 15 can raise her paddle. When called upon Trader 15 will state “Accept Bid on X of 100.” The auctioneer will repeat that “Trader 15 Accepts Bid on X of 100.” At this point Trader 13 and 15 have traded a single X certificate and this will be reflected on the screen which will display the corresponding buyer, seller, and transaction price. [The transaction is displayed on the screen.]

Trader 13 will record on her Information and Record sheet that she purchased an X certificate by writing the price in the Purchase column in the row corresponding to Transaction 1. She will then update her Certificates on Hand and Francs on Hand. Similarly, Trader 15 will record on her Information and Record sheet that she sold an X certificate by writing the price in the Sale column in the row corresponding to Transaction 1. She will then update her Certificates on Hand and Francs on Hand.

All the previous bids and asks in the market for X certificate will be removed, indicating that all of the previous bids and asks for X certificates are no longer available. The auctioneer will then invite new bids or asks.

Trader 13 may state an Ask for X of 600.

Trader 14 may state an Ask for Y of 542.

Trader 17 may state a Bid on Y of 35

Trader 16 may state a Bid on X of 33

Trader 16 may state a Bid on X of 40.

Trader 17 may state a Bid on Y of 60.

Trader 14 may state an Ask on X for 500.

Trader 15 may Accept the Bid on Y of 60.

At this point, the transaction of Y at 60 by Trader 15 and Trader 17 will be reflected on the screen. Trader 17 would record the price of 60 under Purchase in the row for Transaction 1. Trader 17 would then update all of his Certificates on Hand and Francs on Hand. Notice that for Trader 15 this would be her second transaction, so she would record the Sale price on the row for Transaction 2 and then update the remainder of the row.

All the previous bids and asks in the market for Y certificates will be removed, indicating that all of the previous bids and asks for Y are no longer available and the auctioneer will again invite new bids or asks. Notice that the bids and asks for X are still available. This process continues until 5 minutes have passed at which point no more bids, asks, or acceptance can be made.

Suppose there are no more transactions involving Trader 15 in the market period. Also, suppose that at the beginning of the period, Trader 15 was initially provided with 10 X certificates, 10 Y certificates, 10 Z certificates, and 20,000 francs. Assume that Trader 15 earns 50 francs per X certificate if the state is X, 100 francs per Y certificate if the state is Y, and 150 francs per Z certificate if the state is Z. If the state is Y, given the fictitious values in this example, Trader 15 would complete their Information and Record Sheet as follows:

Total Certificate Earnings in Row 19 would be: $9 * 0 + 9 * 100 + 10 * 0 = 900$

Total Francs on Hand in Row 20 would be: $20,160 + 900 = 21,060$

End of Period Net Profit in Row 22 would be: $21,060 - 20,000 = 1,060$

INFORMATION AND RECORD SHEET

Period: 1

Clue: Not z

Trader No. 15

Beginning of the
Year Holdings

Transaction Number	Transaction Price		Certificates on Hand			Francs On Hand
	Sale	Purchase	x	y	z	
0	//////////////////////////////////// ////////////////////////////////////		10	10	10	20,000
1	100		9	10	10	20,100
2	60		9	9	10	20,160
3						
4						
5						
6						
7						
8						
9						
10						
11						
12						
13						
14						
15						
16						
17						
18						
19	Total Certificate Earnings = Dividend Rate × Certificates on Hand at the End of the Period					900
20	Total Francs on Hand at the End of the Period					21,060
21	Less: Fixed Costs					20,000
22	End of Period Net Profit					1,060

x-Dividend: 50
y-Dividend: 100
z-Dividend: 150

Transfer this amount to your Profit Sheet



Trader No. ___

PROFIT SHEET

Row	Market Period	Profit
1	1	
2	2	
3	3	
4	4	
5	5	
6	6	
7	7	
8	8	
9	9	
10	10	
11	11	
12	12	
13	13	
14	14	
15	15	
16	16	
17	17	
18	18	
19	19	
20	20	
21		
22	Total profit (in francs)	
23	Dollars per franc	0.008
24	Total dollars profit	