Author/s:
Anderson, Ryan Joshua Acea

Title:
Essays on Perfect Foresight in Asset Pricing

Date:
2020

Persistent Link:
https://hdl.handle.net/11343/249528

Essays on Perfect Foresight in Asset Pricing

Ryan Anderson
Orcid ID: 0000-0002-7023-9864

A thesis presented in partial fulfilment for the degree of
Doctor of Philosophy

Finance Department
University of Melbourne
Australia

August 2020

Abstract

Through experimental and theoretical analysis, this thesis addresses the question of how the context of information in capital markets can affect the occurrence of perfect foresight equilibria. It contains three essays that build on theoretical and experimental applications of the perfect foresight assumption.

The first essay contrasts theoretical and experimental strains of information aggregation – the ability of prices to aggregate disparate pieces. It contains a methodological guide for designing robust information aggregation experiments with a detailed description of the pilot studies that were used to develop the experimental study in the second essay.

The second essay introduces two new theoretical concepts to the analysis of rational expectations equilibrium models. These concepts stem from "stability" characteristics inherent to the perfect foresight state-price mapping. Differences in stability characteristics are shown to arise from differences in the initial information structure underlying the aggregation problems. In the experiment, we test information aggregation with two fundamental information structures in continuous time double auction asset markets:

The first information structure is motivated by the canonical information aggregation model in theoretical asset pricing. In this setting, the asset traded pays according to the average privately held information signal in the market. This setting has a stable state-price mapping and is shown to aggregate information well.

The second information structure is motivated by prediction markets and studies in experimental finance. Both feature winner-take-all contracts where binary pay depends on the signal type held by the majority of agents. This setting is unstable under the theoretical stability concepts and is shown to aggregate information less efficiently.

The third essay examines the use of perfect foresight when modelling disagreement in financial markets. In particular, we examine the conditions under which the perfect foresight approach can be used in a rational expectations equilibrium model. We show that an agent's perfect foresight may be inconsistent with their own beliefs (based on subjective probabilities) unless their higher-order beliefs (about other participants' beliefs) are correct.

Declaration

I declare:

i This thesis is comprised of only my original work towards the PhD except where indicated in the preface.

ii Due acknowledgement has been made in the text to all other material used.

iii This thesis is less than 100 000 words in length, exclusive of tables, maps, bibliographies and appendices.

Ryan Anderson

Preface

This thesis contains original research in Chapters 2 through 4.

Chapter 3 is based on the following working paper:

> Anderson, R., Bossaerts, P. and F. Fattinger, (2019). Efficiency and Stability of Information Aggregation in Markets. Working Paper.

Chapter 4 is based on the following working paper:

> Anderson, R., Bossaerts, P. (2019). Modelling Asset Prices Under Heterogeneous Beliefs. Working Paper.

Chapter 4 also contains excerpts from:

> Bossaerts, P. (1997). Rational Expectations when Priors are Inconsistent. Caltech Social Science Working paper.

The conclusion of this thesis also contains wording found in these papers.

All co-authorship has taken place in accordance with the Graduate Research Training Policy of the University of Melbourne.

Acknowledgements

I want to thank the University of Melbourne, the Finance Department, and the taxpayers of Australia for providing my scholarship and financial support during my studies. I am incredibly grateful for the opportunity I was awarded.

I have received excellent training and coaching along the way. I am indebted in perpetuity to my principal supervisor, Peter Bossaerts, for shaping me into the researcher that I am today. Peter's guidance and encouragement continuously made me want to work harder in my search for answers. His criticisms were well-founded and always came with a chance to improve. I am also lucky to have been supervised by Felix Fattinger, whose attention to detail has saved me from a lot of unnecessary pain over the past few years. I also couldn't have had a better teammate for running experiments.

I again thank the Department of Finance at the University of Melbourne for providing me with the generous recourses that allowed me to complete my studies and attend external conferences. I want to thank all the professors who have taught me and have been involved in the PhD program. My committee members, Qi Zeng and Spencer Martin deserve a special thanks. They both have been involved in my progression since day one. Spencer passed away during my final year, and I wish I could tell him how much of an impact he had on my time in the program. I also give a special thanks to Michael Gallmeyer and Sean Crockett for teaching me and for providing guidance with the job market. I thank Bruce Grundy, Garry Twite, and Neil Galpin for their involvement with the PhD program. The professional staff has also been outstanding. Thank you to Suzie Collins, Silvia Barberoglou, and Rachel Soderstrom. A special thanks to Jennifer Decolongon for her assistance, especially during the COVID pandemic. She went above the call of duty to ensure that all the PhD students were given the support they required.

I would like to acknowledge everyone involved in the FIRN organization. The FIRN PhD courses and conferences were very insightful and enjoyable. I made a lot of friends across the country who are now spread around the world. FIRN is an excellent organization, and I hope it continues to grow.

I thank everyone at the BMM laboratory: Carsten Murawski, Shinsuke Suzuki, Nitin Yadav, Elizabeth Bowman, Pablo Franco, Xiaping Lu, Jeremy Metha, Karlo Doroc, Anthony Hsu, Michelle Lee, Abhijeet Anand, and Rob Woods. They are an eclectic mix of researchers who foster creativity and free-thinking. They truly helped me and made things fun, even when I was stressed out.

Thanks to my fellow PhD students Yifan Zhou, Bo Zhang, Vivian Xu, Angie Wu, Bochen Wu, Mengchuan Wang, Yolanda Wang, Jiamu Sun, Grace Meng, Linxiang Ma, Leslie Haoyi, Shupei Jia, Sheng Huang, Muhan Hu, Rob Manolache, Nguyet Nguyen, Yichao Zhu, Bowei Li, Roc Lv, Emma Li, Heath Blake, and most importantly my classmate Phuong Nguyen.

From my personal life, I would like to thank my parents and Kristen Anderson, Tom Wally, Stewart Thomas, and Ivanna Dankova for providing constant encouragement for the last five years. Thank you to the entire Mailes family for teaching me how to be a true blue Australian and for sharing the beauty of Phillip Island. Thanks to everyone back in Canada, especially Matt Coady, Quinn Perrick, Kyle Byblow, Allie Kilmartin, Anna Shoemaker, and Elaina Collis. Finally, I would like to acknowledge Kaikoa Anderson and the musicians of Hawaii, mahalo nui loa, and Maui nō ka ‘oi.

Dedication

Dedicated to Emily Bell

for helping me across the finish line and making me smile during a pandemic.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Perfect Foresight

Throughout this thesis I refer to perfect foresight (in the spirit of Radner (1974)) as the ability to correctly forecast all state-contingent equilibrium prices, which is equivalent to knowing the equilibrium state-price mapping. Agents are assumed to obtain perfect foresight by using private economic models that produce expected prices for all goods in all possible states of nature. In equilibrium, the expected prices are correct and common to those who have perfect foresight. However, these agents cannot predict the future in a time series sense. They know the prices that will obtain in a given state; but, they do not know if, or when, a given state will occur. Based on the information they have, each agent forms a subjective probability distribution over the set of possible states. In a perfect foresight rational expectations equilibrium, all agents possess perfect foresight and are never surprised by the observed prices in a given state of nature.

Through experimental and theoretical analysis, this thesis addresses the question of how the context of information in capital markets can affect the occurrence of perfect foresight equilibria.

1.2 The Evolution of Perfect Foresight

1.2.1 Equilibrium and Time

In standard static equilibrium models, agents and/or firms optimize choice variables with respect to budget constraints, and the solution is obtained by equating marginal rates. To advance a static model into the intertemporal class, time must be incorporated. With time naturally comes uncertainty, which leads to the importance of information and forecasting assumptions.

If markets are assumed to be complete and in equilibrium, prices and decisions are obtained as if all information is public knowledge. This happens because the state contingent forward prices reflect the true supply and demand for each state. The complete set of forward contracts thus allows agents to optimize with known equilibrium prices. Complete markets remove price uncertainty.

When markets are incomplete there are no forward markets for some states and these state contingent prices are uncertain. This means that expected equilibrium prices rather than true equilibrium prices must be used to formulate decisions for those states. If the expected equilibrium prices for those states are correct and common, agents have perfect foresight and the equilibrium is pareto-optimal.

With complete markets or perfect foresight, solving an intertemporal model is essentially the same as a solving a static model. After the reduction in complexity, the only issue remaining boils down to the palatability of the complete markets or perfect foresight assumptions.

1.2.2 Hayek and Thinking in Equilibrium

The importance of foresight in neo-classical economic thinking began in the first half of the twentieth century, when this notion of intertemporal equilibrium was taking shape. Hayek believed that, by definition, an equilibrium must include time. In a famous 1937 paper he stated, "since equilibrium is a relationship between actions, and since the actions of one person must necessarily take place successively in time, it is obvious that the passage of time is essential to give the concept of equilibrium meaning." The necessity of time meant that future plans must be a part of today's immediate equilibrium, which meant an agent's ability to form plans depended on how well they could foresee the future.

Critics of the theory of equilibrium were quick to focus on foresight assumptions, which appeared to be the weak point in the theory (Morgenstern, 1935). Hayek (1937) acknowledged this attack on equilibrium theory: "In general it seems that we have come to a point where we all realize that the concept of equilibrium itself can be made definite and clear only in terms of assumptions concerning foresight, although we may not yet all agree what exactly these essential assumptions are."

After admitting the importance of foresight assumptions, Hayek's key point in the same paper was that foresight is a defining characteristic of a system in equilibrium rather than a precondition for arriving at equilibrium. Furthermore, he suggested that the foresight of agents only needed to be correct for as long as an equilibrium existed and that it did not need to be correct for all states of nature. Foresight only needed to be correct for the states that agents took into consideration when forming their optimal plans. This notion of foresight as a characteristic of an equilibrium was the precursor to the "perfect foresight assumption" and the "rational expectations equilibrium."

1.2.3 Intertemporal Equilibrium with Complete Markets: Contingent Claims Arrow-Debreu Model

One solution to the issue of foresight in an intertemporal model is to simply have no need for it. This was famously done in Arrow-Debreu (AD) theory (Arrow (1953), Arrow and Debreu (1954), Debreu (1959)). [1]

In the AD general equilibrium model with uncertainty, time, place and state of nature are included and separated in the commodity space. Goods of the same type are considered different commodities at different moments in time and space, and in different conditions. For example, vanilla ice cream on 20 April 2021 in Melbourne, Australia when the sun shines is a different commodity than vanilla ice cream on 20 April 2021 in Melbourne when rain falls, as is vanilla ice cream on 19 April 2021 in Melbourne when the sun shines or on 20 April 20 2021 in Vancouver, Canada when the sun shines, etc.

By having a "complete market" with separate forward markets for the delivery of each state-contingent commodity, Arrow and Debreu removed the need for foresight. Market clearing prices for every future good in every possible state of the world are observable. Agents do not need to infer anything from prices. The entire state-price mapping is common knowledge and the Walrasian equilibrium is pareto-optimal.

In finance, the AD model is easily extended to a model of contingent claims which is a core model in asset pricing. The contingent claims model also has a pareto-optimal equilibrium under the same set of standard assumptions as the AD model.

The elegance of the AD model comes at a cost. The assumption of a complete set of forward contracts results in all production and consumption decisions being made before the passage of time. This leads to all trade decisions for the entire future occurring at once, before a state is even observed. Everything except the sequence of states in an AD

[1] The early models by Arrow and Debreu were for the cases of certainty, but the addition of uncertainty was straightforward and only required an expansion of the set of tradable commodities.

economy is predetermined and there is no need for institutions, asset markets or money after the initial round of trading; only delivery occurs after the initial round of trade. [2]

Even though the AD model is unrealistic, its importance to asset pricing cannot be overstated. Arrow-Debreu securities, and the state-price mapping they yield when in a complete set, are the most fundamental tools used in time-state modelling. Any model utilising a state price mapping, or stochastic discount factor, is essentially using a variation of what was originally modelled as state-contingent commodities in the AD model. This thesis is no exception; the essays contained within all make use of perfect foresight and equilibrium state-price mappings.

1.2.4 Intertemporal Equilibrium with Incomplete Markets: Radner's Plans Prices and Rational Expectations Equilibrium

To make the AD model more realistic, Radner (1972) relaxed the complete markets assumption and modelled a general equilibrium without a market for every state-contingent commodity. The idea was that when markets are incomplete agents can no longer do all their trading before the start of time. Agents must trade in spot markets when a state without a forward market occurs. Before time starts, agents form an optimal plan for what type of trading will occur in these state-dependent spot markets. They formulate these plans based on their expectations. In equilibrium, all plans, prices and price expectations in the economy are mutually consistent. [3]

In Radner's setting, agents are forced to optimize on state-contingent expected prices. Suddenly, the issue of foresight is reintroduced as states without forward contracts have no observable market price. To solve this issue, Radner imposed that agents have common expected prices for these states. If these common expectations are also the correct equilibrium prices, the agents have "perfect foresight." The equilibrium which is

[2] Radner (1968) discusses several other issues in the AD economy, including the computational requirements needed by agents to optimize when the state space is large and new forward markets open.

[3] Mutual consistency here means that individuals' plans can coexist without violation of equilibrium conditions.

obtained as a result was described by Radner as a "rational expectations equilibrium." In this type of equilibrium, observed prices are consistent with the prices predicted by an agent's model of the economy. This definition has the same flavour as the rational expectations definition introduced by Muth's (1961) partial equilibrium model, where producers' forecasts aligned with equilibrium analysis.

With a slight addition to Muth (1961) and Radner (1972), Radner (1979) defined rational expectations equilibrium in the setting of differential information. [4] The expanded version is in the following excerpt:

> "The possibility for one trader to make inferences from market prices about the information possessed by other traders rests upon his having a 'model' or 'expectations' of how equilibrium prices are determined, i.e., how equilibrium prices are related to the information initially possessed by the various traders. But this relationship is endogenous to the market system, and if traders have any opportunity to compare the results of the operation of the market with their own models, then a suitable equilibrium concept would require that their models not be obviously controverted by their observations of the market. This motivates the term 'rational expectations equilibrium.' "

The subtle difference in this definition emphasizes the duality of the rational expectations equilibrium. In one direction (like Muth's definition), agents can use their models to forecast future prices which in equilibrium are consistent with the observed equilibrium prices. In the other direction (Radner, 1979), agents can use prices to infer the underlying state of nature, and in equilibrium verify their models of the economy. One way to ensure both directions hold is to have perfect foresight. Perfect foresight makes price forecasts consistent and validates the information implied by prices. Both of these directional interpretations of a rational expectations equilibrium will be used throughout this thesis.

[4] Differential information here means that agents have different information about which state may obtain. Agents still have common expectations for state contingent prices.

1.2.5 Radner's Equilibrium in Asset Pricing

The Arrow-Debreu equilibrium has a contingent-claim equilibrium in finance. Similarly, the Radner equilibrium of plans, prices and price expectations also has an equivalent in finance, which is referred to as a security-market equilibrium and can be thought of as a version of the contingent-claim equilibrium with incomplete markets.

An important feature of the security-market setting is the possibility of having dynamically complete markets, where a few long-lived securities can be traded in such a way that they yield the same consumption patterns and pareto-optimal solution as in the Arrow-Debreu contingent-claim setting with complete markets. This was proved by Duffie and Huang (1985) in a continuous time Radner equilibrium setting, and was predicted by Merton (1982) who stated,

> "In the intertemporal version of the Arrow-Debreu model with complete markets, it is well known that the market need only be open 'once' because individuals will have no need for further trade. The continuous trading model is, of course, the opposite extreme. However, the continuous time model appears to have many of the properties of the Arrow-Debreu model without nearly so many securities. Hence, it may be that a good substitute for having so many markets and securities is to have fewer markets and securities, but the existing markets open more frequently. The study of this possibility will be left as a topic for future research."

In a dynamically complete market, a set of tradable securities can be combined in such a way as to create Arrow-Debreu contingent claims for the states with non-existent forward markets. If we think of a basic payoff tree, a necessary condition for a dynamically complete market is that the number of future states in the next period (possible future nodes connected to the current node) can be spanned by the set of tradable securities. If the securities are linearly independent, this will occur if there are at least the same number of tradable securities as there are states in the next level of the tree. In this case, the set of securities will span the state space and can be combined in different ways to create state contingent claims for each possible state. If this property holds for each level in the tree, the market is dynamically complete. If we take this idea to the limit, and shrink the time between nodes in the tree, we get the concept of continuous trading behind Duffie and Huang, which Merton was referring to.

Note that in equilibrium, in order for agents in a dynamically complete setting to make optimal trading plans that arrive at the same pareto-optimal solution as the complete markets setting, agents must have perfect foresight (as in the standard Radner equilibrium setting). To linearly combine a set of securities to make a portfolio that pays one dollar in a given state, an agent must know the equilibrium values of the underlying securities in the given state. Perfect foresight (or an equivalent) is fundamental to the dynamically complete solution.

In finance, perfect foresight can also be seen in many other time-state models, such as the Black-Scholes option pricing model and even the corporate finance model of Modigliani and Miller (for examples see Bossaerts, Fattinger, van den Bogaerde, Yang, 2019). [5] [6]

1.2.6 A Final Note on Perfect Foresight

Perfect foresight began as a general equilibrium concept before appearing in a variety of theoretical applications. It is a powerful modelling technique; however, it was not born out of convenience. It evolved with the theory of a rational expectations equilibrium and should be used accordingly.

Perfect foresight was used by Lucas in his famous 1978 model, where it appeared within the framework of a rational expectations equilibrium and was used to "close" the model. Lucas's agent's models affected prices and prices affected the agent's models.

[5] Outside of finance, perfect foresight also shows up in game theory as subgame perfection, where an agent at any node is assumed to know the future equilibrium values of the possible future strategies.

[6] Perfect foresight also appears in macroeconomics; see Glasner (2018) for a survey and another interpretation of the concepts presented here.

Perfect foresight made these effects coherent while satisfying the equilibrium conditions. In Lucas's model, perfect foresight and rational expectations together were the keystone that allowed the equilibrium to be internally consistent. Echoing Hayek's argument, these features were a part of the equilibrium rather than an assumption leading to one. This is how perfect foresight should be used – with equilibrium in mind.

Throughout this thesis, I reference perfect foresight with respect to fully revealing equilibrium prices in a competitive market. This is how perfect foresight was originally used.

If agents without perfect foresight exist, many different equilibria may occur depending on the extent of their price impact as a group. Deviations from rationality allow the behavioural zoo to increase the equilibrium state-space, dramatically stretching the mental capacity of those assumed to maintain perfect foresight. For this reason, I keep things simple and restrict the extent of perfect foresight to fully revealing prices. In the second essay, I discuss the relationship between fully revealing and private value equilibrium prices. However, I do not examine perfect foresight with respect to any other equilibria.

1.3 Thesis Structure

This thesis presents three essays that build on theoretical and experimental applications of the perfect foresight assumption:

The first essay contrasts theoretical and experimental strains of information aggregation – the ability of prices to aggregate disparate pieces. It contains a methodological guide for designing robust information aggregation experiments with a detailed description of the pilot studies that were used to develop the experimental study in the second essay.

The second essay introduces two new theoretical concepts to the analysis of rational expectations equilibrium models. These concepts stem from “stability” characteristics inherent to the perfect foresight state-price mapping. Differences in stability characteristics are shown to arise from differences in the initial information structure underlying the aggregation problems. In the experiment, we test information aggregation with two fundamental information structures in continuous time double auction asset markets:

The first information structure is motivated by the canonical information aggregation model in theoretical asset pricing. In this setting, the asset traded pays according to the average privately held information signal in the market. This setting has a stable state-price mapping and is shown to aggregate information well.

The second information structure is motivated by prediction markets and studies in experimental finance. Both feature winner-take-all contracts where binary pay depends on the signal type held by the majority of agents. This setting is unstable under the theoretical stability concepts and is shown to aggregate information less efficiently.

The third essay examines the use of perfect foresight when modelling disagreement in financial markets. In particular, we examine the conditions under which the perfect foresight approach can be used in a rational expectations equilibrium model. We show that an agent’s perfect foresight may be inconsistent with their own beliefs (based on subjective probabilities) unless their higher-order beliefs (about other participants’ beliefs) are correct.

Chapter 2

Information Aggregation Theory, Experimental Designs, Literature Review and Pilot Studies

2.1 Introduction

This essay covers the relationship between information aggregation theory and experimental design. It aims to provide the theoretical and methodological knowledge necessary to evaluate and design experiments on information aggregation. I will do this by reviewing relevant literature, a classic theoretical model of information aggregation, and outlining my piloting process. Throughout this essay, I will highlight the importance of perfect foresight.

2.1.1 What Does Information Aggregation Mean?

Information aggregation in capital markets occurs when prices correctly incorporate disparate pieces of information privately held by those who trade in the market. If aggregation occurs such that the equilibrium price reflects all relevant information (the aggregate information set), the equilibrium is said to be fully revealing. If the equilibrium price only reflects some information, but not the entire set, it is said to be a partially revealing equilibrium.

2.1.2 The Historical Significance

The importance of the informational content of prices to the efficient allocation of economic resources was made apparent during the middle of the twentieth century. During this time tensions between political ideologies were high. The debate over the relative efficiency of central planning versus the free market was in full swing.

In another famous paper by Hayek (1945), he emphasized the efficiencies of "man on the spot" decisions made by individuals who posses local knowledge of "time and place." He argued that if "man on the spot" decisions are made in conjunction with the information in prices, and occur spread out over the entire economy, a decentralized system can produce a more efficient outcome than a central plan made without market prices or local knowledge.

The key to this advantage was that price changes naturally cause decentralized economic agents to shift their allocations toward an efficient solution. For this to be true, market prices must accurately reflect the true information in the economy; prices must aggregate information. This economic idea of price guided allocational efficiency has an equivalent in finance and relies on what is known as the "efficient market hypothesis."

2.1.3 The Efficient Markets Hypothesis

In Fama's 1970 paper, he reviewed theoretical and empirical work related to the now famous efficient market's hypothesis (Bachlier (1900); Mandlebrot (1963); Samuelson (1965)). He began the review with a very impactful statement.

> "The primary role of the capital market is allocation of ownership of the economy's capital stock. In general terms, the ideal is a market in which firms can make production-investment decisions, and investors can choose among the securities that represent ownership of firms' activities under the assumption that security prices at any time 'fully reflect' all available information. A market in which prices always 'fully reflect' available information is called 'efficient.' "

Fama's opening statement declares the fundamental importance of information aggregation. He implies that without information aggregation, capital markets can not perform their primary function of ownership allocation.

Before reviewing the empirical studies done on market efficiency, Fama broke the notion of efficiency into three nested levels based on the incorporation of different information sets. His definitions of efficiency are now common place in academic finance.

Weak-form efficient prices reflect historical information (i.e., historical prices, past accounting statements, etc.). In addition, semi-strong form efficient prices also immediately reflect released public information (i.e., earnings announcements, stock splits, etc.). Finally, strong-form efficient prices include historical, public, and all available relevant private information to the formation of prices.

Fama's strong-form efficiency is directly related to information aggregation. If equilibrium prices aggregate all private information then prices are also strong-form efficient. Understanding information aggregation provides knowledge and insight about price efficiency which ultimately helps us understand the health of our financial markets and our social welfare.

2.2 Information Aggregation Theory

Theories of information acquisition developed concurrently with those of information aggregation. The connecting idea is that it may be profitable to acquire costly information if markets do not fully aggregate private information. In 1971, Hirshleifer used a production-based model to show private incentives to collect information can result in an over-investment in information generation beyond the social value of the information generated. In response, Grossman released a series of papers (1975, 76, 78) where informative prices and costly information led to the opposite conclusion. He argued when prices are fully revealing, and information is costly, the private incentives to acquire information can be reduced below the social value of the information. To overcome the cost of information acquisition, he concluded that there must be noise in the price system so those who collect information can obfuscate it from others while trading. [1]

The case of costless information generation in Grossman's 1976 paper was one in which the equilibrium price aggregates all dispersed information such that each trader achieves the same allocations as if they knew the entire set of private information. In the

[1] Grossman's papers here were concerning the case of differential informational aggregation, rather than the asymmetric information case like his famous 1980 paper with Stigilitz.

model, information aggregation is independent of agents' risk aversion. Still, it depends on their knowledge of the covariance between noise in their private signals and noise in the price. By understanding the covariance structure, individuals can drop their private information in favour of the information in the fully revealing equilibrium price.

Hellwig (1980) objected to Grossman's logic. In Grossman's model, the covariance between noise in prices and noise in individual signals results from a finite number of agents. The finite nature of the set up means that each agent's signal impacts the average signal and the price. The conflict in logic arises because the agents are aware of the covariance structure between signals and prices, yet are naive of their potential to influence price. The non-negligible agents are equipped with private information and the knowledge to act strategically; but, they do not.

To mitigate the impact of individual traders and remove the possibility of strategic behaviour, Hellwig increased the number of traders. However, a large N (number of agents) and a fully informative equilibrium price still subjected Hellwig's model to a paradox that plagues information aggregation theory. The paradox is as follows: if all information is correctly aggregated into the price, agents have no reason to utilize their information; however, if agents do not use their information, then the price cannot aggregate all information.

To account for this paradox, Hellwig introduced a small amount of noise to the traded asset's supply. This noise provided the individual traders with an incentive to maintain a positive weight on their private signals as they could no longer infer the state of nature from the price with certainty. In the limit, prices again aggregated information. Hence, models of this type would became known as the noisy rational expectation equilibrium models (Hellwig (1980); Diamond and Verrechia (1981)).[2].

I have elaborated on an example of Grossman's model presented below. I emphasize the problem's Bayesian nature by including a few new equations and steps in the derivation. Please note that I am in no way claiming origination of the model below,

[2]For asymmetric information models with noise in supply, see (Grossman and Stiglitz (1980); (Kyle (1985))

it is a well-known model and taught in many graduate-level classes on asset pricing. It is provided here for reference and to provide an intimate understanding of the information structure used in theoretical modelling, which will be vital to my discussion of experimental designs and pilot studies.

2.2.1 Grossman's Information Aggregation Equilibrium

The classic Grossman (1976) model of information aggregation has two periods, time zero and time one. It has n different trader types, each with a unique piece of information. There is a risk-free and a risky asset. The risky asset pays $\widetilde{P}_1$ in period one. In period zero, each trader must decide how much of the risky asset they will hold. A chosen allocation depends on an agent's expected value of $\widetilde{P}_1$ conditional on their information set I and private model of $\widetilde{P}_1$. The conditional expectations are inherently Bayesian.

We can see this with the Bayesian posterior distribution $P(\widetilde{P}_1|I)$ composed of a prior $P(\widetilde{P}_1)$, likelihood function $P(I|\widetilde{P}_1)$, and a marginal $P(I)$ distribution:

$$P(\widetilde{P}_1|I) = \frac{P(I|\widetilde{P}_1) \cdot P(\widetilde{P}_1)}{P(I)}. \tag{2.1}$$

The structure of the aggregation problem and ease of solving the model depends heavily on the assumed structure of the likelihood function.

Agents' wealth

Assume each agent i is endowed with wealth W_{0i}, and allocated x_{fi} into the risk free asset, and x_i into the risky asset. The period zero budget constraint is:

$$\widetilde{W}_{0i} = x_{fi} + \widetilde{P}_0 x_i. \tag{2.2}$$

The period one wealth can be expresses as:

$$\widetilde{W}_{1i} = (1+r)x_{fi} + \widetilde{P}_1 x_i. \tag{2.3}$$

Substituting 2.2 into 2.3 gives:

$$\widetilde{W}_{1i} = (1+r)W_{0i} + [\widetilde{P}_1 - (1+r)P_0]x_i. \tag{2.4}$$

Private Signals

At time zero, each agent receives a noisy signal about the payoff of the risky security. Assume agent i receives signal s_i, where P_1 is the realization of the random variable $\tilde{P}_1$,

$$s_i = P_1 + \epsilon_i.$$

Note, the ϵ_i term prevents any agent from knowing the true payoff with certainty.

Trader i's signal is normally distributed about the state of nature drawn, P_1. The signal is normal with mean P_1 and variance 1. The set of signals are i.i.d, so the covariance of ϵ_i and ϵ_j are zero for $i \neq j$. The set of signals $S = (s_1, s_2, ..., s_n)$ follows a joint density, $f(S|P_1)$ which is a multivariate normal with $\mu_S = (P_1, P_1, ..., P_1)$ and covariance matrix equal to the identity matrix.

The marginal distribution of $\tilde{P}_1$ is assumed to be Normal(P_1, σ^2) and independent of $\epsilon = (\epsilon_1, \epsilon_2, ..., \epsilon_n)$.

Agents are aware of this information structure and use it to formulate their Bayesian posterior means.

Time Zero Equilibrium Price of the Risky Asset

At time zero, the equilibrium price of the risky asset is a function of all the unique signals; $P_0(s_1, s_2, ..., s_n)$. If this time zero price is a sufficient statistic for S, agents will be able to infer P_1 if they know the pricing function (have perfect foresight). This means, by conditioning on P_0 and inferring P_1 they will have a higher utility than if they conditioned on their private information alone.

The ease of inferring P_1 from P_0 depends on the underlying relationship between the joint distribution of information signals and $\tilde{P}_1$. We will see the normal structure of the signals and $\tilde{P}_1$ prices will allow P_0 to be expressed as a linear function of the average signal $\bar{S}$. This simplifies the processes of solving the equilibrium coefficients. It also makes inferring $\bar{S}$ and P_1 from P_0 straight forward.

Utility Maximization

All agent's have CARA utility functions, with agent i having a_i coefficient of absolute risk aversion. She maximizes her utility conditional on her private information set I_i,

$$U_i(\widetilde{W_{1i}}) = -e^{a_i \widetilde{W_{1i}}}, \qquad a_i > 0. \tag{2.5}$$

If $\widetilde{W_{1i}}$ is normal conditional on I_i, then the expected utility function is:

$$E[U_i(\widetilde{W_{1i}})|I_i] = -e^{-a_i[E[\widetilde{W}_{1i}|I_i] - \frac{a_i}{2} Var[\widetilde{W}_{1i}|I_i]]}. \tag{2.6}$$

This is equivalent to maximizing:

$$E[\widetilde{W}_{1i}|I_i] - \frac{a_i}{2} Var[\widetilde{W}_{1i}|I_i]. \tag{2.7}$$

Where,

$$E[\widetilde{W}_{1i}|I_i] = (1 + r)W_{0i} + [E[\widetilde{P}_1|I_i] - (1 + r)P_0]x_i, \tag{2.8}$$

and,

$$Var[\widetilde{W}_{1i}|I_i] = x_i^2 Var[\widetilde{P}_1|I_i]. \tag{2.9}$$

The maximization problem can be written as:

$$\max_{x_i}[(1+r)W_{0i} + [E[\widetilde{P}_1|I_i] - (1+r)P_0]x_i - \frac{a_i}{2}x_i^2 Var[\tilde{P}_1|I_i]]. \tag{2.10}$$

The optimal demand for the risky asset is:

$$x_i^d = \frac{E[\tilde{P}_1|I_i] - (1+r)P_0}{a_i Var[\tilde{P}_1|I_i]}. \tag{2.11}$$

Equilibrium Conditions

Let X be the total supply of the risky asset. In equilibrium,

$$\sum_{i=1}^{n} x_i^d = X. \tag{2.12}$$

Rational Expectations Equilibrium and Perfect Foresight

From the above, we can see that each agent's demand is a function of their Bayesian posterior mean for P_1, which depends on their information set. Defining the information sets used affects the likelihood function's structure and will determine what type of equilibrium obtains.

At one extreme, in Walrasian equilibrium, each agent's knowledge is limited to the private signal they receive. The Bayesian agents simply update their posteriors according to their private signals and do not use any information from prices when forming their optimal demands.

At the other extreme, in a fully revealing rational expectations equilibrium, each agent's information set includes their private signal and their "model" of equilibrium. As in Muth (1961) and Radner (1972), Grossman's agents' models of equilibrium are also the correct model of equilibrium. The presence of these "models" in agents' information sets makes the Bayesian likelihood functions used to form expectations of $\tilde{P}_1$ in the rational expectations equilibrium different from those used in the Walrasian equilibrium. Knowledge of the structure of the relationship between aggregate information and price formation allows agents to use the information from time-zero equilibrium prices to estimate the average signal and subsequently estimate time-one prices.

In Grossman's rational expectations model, agents' models are the equilibrium price function equivalent to Radner's perfect foresight condition: agents know the aggregate set of signals corresponding to each equilibrium price and vice versa.

To define the equilibrium we express the aggregate set of private signals as $S \equiv (s_1, s_2, ..., s_n)$, and the rational expectations equilibrium price function as $P_0^*(S)$. We can now express the rational expectations (perfect foresight) equilibrium condition as:

$$\sum_{i=1}^{n} \frac{E[\tilde{P}_1|s_i, P_0^*(S)] - (1+r)P_0^*}{a_i Var[\tilde{P}_1|s_i, P_0^*(S)]} = X. \tag{2.13}$$

Deriving the Conditional Mean and Variance

In Grossman's model, he proves a lemma that the sample mean $\bar{S}$ of the s_i is a sufficient statistic for the family of densities $f(S|P_1)$. In general terms, once a trader knows the sample mean of the distributed signals, she has the same amount of information as if she knew the entire sample of signals. Flowing from the first lemma, Grossman then proves a second lemma that the Bayesian posterior mean and variance of P_1 conditional on $\bar{S}$ is the same as the estimates of P_1 conditional on $\bar{S}$ and s_i. [3]

[3] In Grossman's derivation of the lemmas, the structure which Hellwig objected to can be seen. It is not presented here for brevity.

The second lemma gives:

$$E[\tilde{P}_1|s_i, \bar{S}] = E[\tilde{P}_1|\bar{S}], \tag{2.14}$$

$$Var[\tilde{P}_1|s_i, \bar{S}] = Var[\tilde{P}_1|\bar{S}]. \tag{2.15}$$

The above equations imply that once an agent knows the average signal $\bar{S}$, she is no better off by using her private signal. [4]

The result of the second lemma combined with the distributional assumptions above, is then plugged into a theorem from DeGroot (1970) to obtain the posterior mean and variance estimates. The theorem is as follows,

Theorem 1, Sec, 9.5, DeGroot (1970) *Suppose that $X_1, ..., X_n$ is a random sample from a normal distribution with an unknown value of the mean W and a specified value of the precision r $(r > 0)$. Suppose also the prior distribution of W is a normal distribution with mean μ and precision τ such that $-\infty < \mu < \infty$ and $\tau > 0$. Then the posterior distribution of W when $X_i = x_i (i = 1, ..., n)$ is a normal distribution with mean μ' and precision $\tau + nr$ where:*

$$\mu' = \frac{\tau\mu + nr\bar{x}}{\tau + nr}. \tag{2.16}$$

The resulting expressions used by Grossman are [5]:

$$E[\tilde{P}_1|s_i, \bar{S}] = E[\tilde{P}_1|\bar{S}] = \frac{\bar{P}_1 + n\sigma^2\bar{S}}{1 + n\sigma^2}, \tag{2.17}$$

$$Var[\tilde{P}_1|s_i, \bar{S}] = Var[\tilde{P}_1|\bar{S}] = \frac{\sigma^2}{1 + n\sigma^2}. \tag{2.18}$$

Following DeGroot (1970), the posterior mean is more intuitively represented as a

[4] This idea is the basis for what later evolved into the the paradox discussed earlier and became a focal point of Hellwig (1980) which led to his model of the noisy rational expectations equilibrium.

[5] The assumed prior distribution in Grossman's model has unit variance, so the precision, τ, is also 1.

weighted average of the prior mean $\bar{P}_1$ and the signal sample mean $\bar{S}$,

$$E[\tilde{P}_1|\bar{S}] = \frac{1}{1+n\sigma^2} \cdot \bar{P}_1 + \frac{n\sigma^2}{1+n\sigma^2} \cdot \bar{S}, \tag{2.19}$$

$$\text{Let,} \qquad \theta = \frac{1}{1+n\sigma^2}, \tag{2.20}$$

$$E[\tilde{P}_1|\bar{S}] = \theta \cdot \bar{P}_1 + (1-\theta) \cdot \bar{S}. \tag{2.21}$$

Note, more weight is put on the average signal when the total number of signals is higher, or when the variance of the price is higher, which makes intuitive sense.

Using the former representation of the conditional moments, Grossman proves his main result,

Theorem 1, Grossman 1976 *Under the above assumption about the joint distribution of S and $\tilde{P}_1$, if $P_0^*(S)$ is given by*

$$P_0^*(S) = \alpha_0 + \alpha_1 \bar{S}, \qquad \textit{where} \tag{2.22}$$

$$\bar{S} \equiv \sum_{i=1}^{n} \frac{s_i}{n}, \qquad \textit{and} \tag{2.23}$$

$$\alpha_0 \equiv \frac{\bar{P}_1 \sum_{i=1}^{n} \frac{1}{\alpha_i} - \sigma^2 \bar{X}}{(1+n\sigma^2)(1+r)\sum_{i=1}^{n} \frac{1}{\alpha_i}}, \tag{2.24}$$

$$\alpha_1 \equiv \frac{n\sigma^2}{(1+n\sigma^2)(1+r)}, \tag{2.25}$$

then $P_0^*(S)$ is an equilibrium. That is, it is a solution to 2.13

Grossman uses a rational expectations equilibrium framework, meaning the proposed price function is known by traders. Hence, by the above lemmas and theorem from DeGroot (1970), in equilibrium the observation of the initial price combined with the pricing function (perfect foresight) yields the posterior mean as a function of the average signal from the sample,

$$E[\tilde{P}_1|s_i, P_0^*(S)] = E[\tilde{P}_1|s_i, \alpha_0 + \alpha_1\bar{S}] = E[\tilde{P}_1|s_i, \bar{S}] = E[\tilde{P}_1|\bar{S}] = \frac{\bar{P}_1 + n\sigma^2\bar{S}}{1 + n\sigma^2}, \quad (2.26)$$

$$Var[\tilde{P}_1|s_i, \alpha_0 + \alpha_1\bar{S}] = Var[\tilde{P}_1|s_i, \bar{S}] = Var[\tilde{P}_1|\bar{S}] = \frac{\sigma^2}{1 + n\sigma^2}. \quad (2.27)$$

The above moments and the proposed price function can now be plugged into the aggregate demand equilibrium condition:

$$\sum_{i=1}^{n}\left(\frac{E[\tilde{P}_1|s_i, P_0^*(S)] - (1+r)P_0^*}{\alpha_i Var[\tilde{P}_1|s_i, P_0^*(S)]}\right) = \bar{X}. \quad (2.28)$$

This yields the equilibrium condition:

$$\sum_{i=1}^{n}\left[\frac{\frac{(\bar{P}_1 + n\sigma^2\bar{S})}{(1+n\sigma^2)} - (1+r)(\alpha_0 + \alpha_1\bar{S})}{\frac{\alpha_i\sigma^2}{(1+n\sigma^2)}}\right] - \bar{X} = 0. \quad (2.29)$$

To solve for the α_0 and α_1 coefficients, we do some algebra to separate the terms that do not contain $\bar{s}$ from those that do:

$$\sum_{i=1}^{n}\left[\frac{(\bar{P}_1 + n\sigma^2\bar{S})}{(1+n\sigma^2)\alpha_i} - \frac{(1+r)\alpha_0}{\alpha_i}\right] - \sum_{i=1}^{n}\left[\frac{(1+r)\alpha_1\bar{S}}{\alpha_i}\right] - \frac{\bar{X}\sigma^2}{(1+n\sigma^2)} = 0. \quad (2.30)$$

Now the terms can be grouped and the above equation split such that [6]:

$$\sum_{i=1}^{n}\left[\frac{\bar{P}_1\frac{1}{\alpha_i}}{(1+n\sigma^2)} - \frac{(1+r)\alpha_0\frac{1}{\alpha_i}(1+n\sigma^2)}{(1+n\sigma^2)} - \frac{\bar{X}\sigma^2}{(1+n\sigma^2)}\right] = 0, \quad (2.31)$$

$$\sum_{i=1}^{n}\left[\frac{(n\sigma^2 - (1+n)\sigma^2(1+r)\alpha_1)\bar{S}}{(1+n\sigma^2)\alpha_i}\right] = 0. \quad (2.32)$$

[6] The proof is one of equilibrium existence. Splitting the above equation, which contains two unknowns, turns a necessary condition into two sufficient conditions. The two equations then allow the coefficients to be expressed independently of one another, and makes solving the equilibrium easier (i.e., $a + b = 0$, holds if $a = 0$ and $b = 0$).

Solving the first equation yields the equilibrium α_0,

$$\sum_{i=1}^{n}[\frac{\bar{P}_1\frac{1}{\alpha_i}}{(1+n\sigma^2)} - \frac{(1+r)\alpha_0\frac{1}{\alpha_i}(1+n\sigma^2)}{(1+n\sigma^2)} - \frac{\bar{X}\sigma^2}{(1+n\sigma^2)}] = 0, \tag{2.33}$$

$$\frac{1}{(1+n\sigma^2)} \cdot \sum_{i=1}^{n}[\bar{P}_1\frac{1}{\alpha_i} - (1+r)\alpha_0\frac{1}{\alpha_i}(1+n\sigma^2) - \bar{X}\sigma^2] = 0, \tag{2.34}$$

$$\sum_{i=1}^{n}[\bar{P}_1\frac{1}{\alpha_i} - (1+r)\alpha_0\frac{1}{\alpha_i}(1+n\sigma^2) - \bar{X}\sigma^2] = 0, \tag{2.35}$$

$$\bar{P}_1\frac{1}{\sum_{i=1}^{n}\alpha_i} - (1+r)\alpha_0\frac{1}{\sum_{i=1}^{n}\alpha_i}(1+n\sigma^2) - \bar{X}\sigma^2] = 0, \tag{2.36}$$

$$\alpha_0 = \frac{\bar{P}_1\sum_{i=1}^{n}\frac{1}{\alpha_i} - \bar{X}\sigma^2}{(1+r)(1+n\sigma^2)\sum_{i=1}^{n}\frac{1}{\alpha_i}}. \tag{2.37}$$

Solving the second equation yields the equilibrium α_1,

$$\sum_{i=1}^{n}[\frac{(n\sigma^2 - (1+n)\sigma^2(1+r)\alpha_1)\bar{s}}{(1+n\sigma^2)\alpha_i}] = 0, \tag{2.38}$$

$$\sum_{i=1}^{n}\frac{1}{(1+n\sigma^2)\alpha_i} \cdot [(n\sigma^2 - (1+n)\sigma^2\alpha_1)\bar{s}] = 0, \tag{2.39}$$

$$\alpha_1 = \frac{n\sigma^2}{(1+n\sigma^2)(1+r)}. \tag{2.40}$$

This result proves the existence of a fully revealing equilibrium with a price function that is a function of the average signal.

Individual Demand in Equilibrium

Now that we have solved for the coefficient in the linear pricing function. Assuming agents have the same risk aversion coefficients, we can use the coefficients to express the optimal demand of each agent as:

$$x_i^d = [\frac{\frac{(\bar{P}_1 + n\sigma^2 \bar{S})}{(1+n\sigma^2)} - (1+r)(\alpha_0 + \alpha_1 \bar{S})}{\frac{\alpha_i \sigma^2}{(1+n\sigma^2)}}] \tag{2.41}$$

$$= [\frac{(\bar{P}_1 + n\sigma^2 \bar{S})}{(1+n\sigma^2)} - (1+r)(\alpha_0 + \alpha_1 \bar{S})] \cdot \frac{(1+n\sigma^2)}{\alpha_i \sigma^2} \tag{2.42}$$

$$= \frac{(\bar{P}_1 + n\sigma^2 \bar{S})}{(1+n\sigma^2)} \cdot \frac{(1+n\sigma^2)}{\alpha_i \sigma^2} - (1+r)(\alpha_0) \cdot \frac{(1+n\sigma^2)}{\alpha_i \sigma^2} - (1+r)(\alpha_1 \bar{S}) \cdot \frac{(1+n\sigma^2)}{\alpha_i \sigma^2} \tag{2.43}$$

$$= \frac{(\bar{P}_1 + n\sigma^2 \bar{S})}{(1+n\sigma^2)} \cdot \frac{(1+n\sigma^2)}{\alpha_i \sigma^2} - (1+r)(\frac{\bar{P}_1 \sum_{i=1}^n \frac{1}{\alpha_i} - \bar{X}\sigma^2}{(1+r)(1+n\sigma^2)\sum_{i=1}^n \frac{1}{\alpha_i}}) \cdot \frac{(1+n\sigma^2)}{\alpha_i \sigma^2} - (1+r)(\frac{n\sigma^2}{(1+n\sigma^2)(1+r)})(\bar{S}) \cdot \frac{(1+n\sigma^2)}{\alpha_i \sigma^2} \tag{2.44}$$

$$= \frac{1}{\alpha_i \sigma^2} \cdot [(\bar{P}_1 + n\sigma^2 \bar{S}) - \bar{P}_1 + \frac{\bar{X}\sigma^2}{\sum_{i=1}^n \frac{1}{\alpha_i}} - n\sigma^2 \bar{S}] \tag{2.45}$$

$$= \frac{1}{\alpha_i \sigma^2} \cdot [\frac{\bar{X}\sigma^2}{\sum_{i=1}^n \frac{1}{\alpha_i}}] \tag{2.46}$$

$$= \frac{1}{\alpha_i} \cdot \frac{\bar{X}}{\frac{N}{\alpha_i}} = \frac{\bar{X}}{N}. \tag{2.47}$$

The fact individual demand can reduce to such a pure expression is of no coincidence. Grossman's proof used the lemma result that agents do not need to condition their private information when prices are fully revealing in equilibrium. The lemma result is present in the individual demand for the risky asset. It is simply the individual's share of the asset according to the socially optimal sharing rule. In the case above, I assumed the risk aversion coefficients to be the same for each agent, so the aggregate supply is divided equally in equilibrium.

Note, the demand equation is independent of price, and agents' demands are inelastic. Agent's are price takers. Again, this is because the equilibrium price is fully revealing; for any price, all information is aggregated, and agents have no reason to take a price dependent position.

As eluded to earlier, this price-taking behaviour combined with self-awareness of price influence is what led Hellwig to refer to Grossman's agents as schizophrenic.

2.2.2 Information Structure Recap

In Grossman's model, agents' signals are i.i.d Normal and centered around the P_1 value drawn by nature in period zero and experienced in period one. The agents all have CARA utility allowing their optimization problems to be expressed in terms of their Bayesian posterior mean and variance estimates. The Normal structure of the signals combined with rational expectations (perfect foresight), and the period zero equilibrium price, allows for the posterior moments to be expressed as a function of the average signal sample. It also results in a linear pricing function, where the price is a function of the average signal. The beauty and elegance of the solution all stem from the information structure. Unfortunately, as we shall see in the next section, there has not been an experimental study that has utilized the same structure as in the theory.

2.3 Information Aggregation in Experimental Markets

2.3.1 Why are Experiments Necessary?

The classic model from above proved the existence of a fully revealing equilibrium in the differential information setting with Normal signals and CARA utility. The proof is specific to the initial conditions and does not guarantee that the fully revealing equilibrium will obtain in a system that starts off equilibrium. For this reason, it is crucial to test information aggregation in the field and the laboratory. However, field testing comes with serious challenges.

To test for the informational efficiency of prices, one must have a model of efficient prices as a benchmark. Fama (1970) pointed out the joint-hypothesis problem: Testing prices for informational efficiency simultaneously tests the proposed model of equilibrium. Thus, it is impossible to tell whether prices are inefficient if there is no way to tell if the proposed equilibrium model is correct. Even if the proposed equilibrium model is right, there is no way to verify the setting for an aggregation problem in the field. Access to the private information of market participants is simply unattainable.

These problems have led to the empirical testing of informational efficiency through price changes rather than levels. If prices are efficient, they should follow a random walk (Sameulson (1965)). There has been much research done in this area (see Fama (1970) (1991); Lim and Brooks (2009)). However, using price changes to test for efficiency is also problematic as there is no reason why efficient price changes guarantee efficiency in price levels.

Experiments provide a partial solution to the limitations of empirical research on informational efficiency. They allow direct control of initial endowments and information signals, as well as aggregate level characteristics and institutional design. Even though laboratory experiments will never match the institutional detail, or scale, of a market in the field, they are the ideal tool for testing theories with initial conditions which are too complex to observe in the field.

In the next two sections, I will discuss two of the most commonly used experimental designs for testing informational price efficiency. Both configurations have tested efficiency in settings of differential information (information aggregation), information dissemination and amplification (asymmetric information), and information acquisition. I will present the basic designs with a focus on information aggregation.

2.3.2 The Diverse Divedend Schedule Design

In 1988, Plott and Sunder conducted their seminal experiment on information aggregation. Their design was an extension of earlier experiments by Forsythe, Palfrey, and Plott (1982), and Plott and Sunder (1982).

The classic design is as follows: there are three distinct states of nature (x, y, z), and each occurs with known probability. Once a state is chosen, each trader receives a signal about the underlying state. In isolation, the individual signals do not reveal the selected state but provide information about which state was not chosen. For example, if the underlying state were z, then half the traders would receive "not x", and the other half would receive "not y". Together, both signals reveal the correct state to be z.

A continuous oral double auction is open for a single trading period, then the traded assets pay a dividend. The dividends paid not only depend on the state of nature but also on the trader type that held the asset at the end of the period. Hence, the diverse dividend schedule.

Plott and Sunder (1988), tested three "series" of markets. Series A had incomplete markets with one traded security and two or three trader types. Series B had complete markets with three contingent claims and three trader types. Series C also had incomplete markets, but only one trader type (uniform dividends).

The design of the dividend schedule was chosen to allow for separation of the private value (Walrasian) equilibrium from the fully revealing rational expectations equilibrium. The private value equilibrium price reflected the highest expected value of the dividend according to a Bayesian update done on private information only. The fully revealing price was simply the highest dividend value across the trader types in the selected state of nature.

To make thing clear, I have provided an example of each series from Plott and Sunder (1988) in the Table 2.1.

Table 2.1: Plott and Sunder Design

Market	Series	Trader Type	Number of Traders	Initial Endowment Certificates	Initial Endowment Francs	Dividends x	Dividends y	Dividends z	Probabilities x	Probabilities y	Probabilities z
5	A	I	4	4	16,000	50	240	590	1/3	1/3	1/3
		II	4	4	16,000	170	450	110			
		III	4	4	16,000	310	190	390			
4	B	I	4	2	10,000	70	130	300	0.35	0.20	0.45
		II	4	2	10,000	230	90	60			
		III	4	2	10,000	100	160	200			
7	C	I	12	4	25,000	50	240	490	1/3	1/3	1/3

Plott and Sunder found that information aggregated in their series B and series C markets. Series A featured diverse dividend schedules with a single security and did not aggregate information very well.

In series B, the contingent claims paid zero if their underlying state was not chosen. Therefore, traders receiving a signal about a non-occurring state received information that the value of the respective contingent claim was zero with certainty. Traders had complete knowledge of the true value of one of the three securities.

The presence of the state-contingent claims mapped the information aggregation problem into a multi-asset setting where two out of three securities had traders with complete information. Competition amongst these fully informed traders likely drove the prices of the non-occurring contingent claims to zero, which in turn revealed the correct state. Even though this setting is one of information aggregation, the complete set of contingent claims and information structure turned two of the asset markets into cases of information dissemination/amplification. It simplified the aggregation problem in the remaining market.

There was only one security traded in series C, and all traders were of the same type (common dividend schedule). Having common knowledge of the same dividend schedule gave the agents perfect foresight. They knew what the price of the security should be conditional on any state.

The successful aggregation of information in series B and C was likely due to these unique information structures.

Recently, the information aggregation performance of Plott and Sunder (1988) has been called into question after a thorough replication of the series B and C markets done by Corgnet et al. (2019). They followed the replication recommendations of Camerer et al. (2016) and increased the sample size by collecting approximately four times the amount of data. After using the same measures of efficiency, their results did not support information aggregation in series B or series C. This replication attempt has shown how fragile information aggregation and strong form efficient prices are. In the second essay of this thesis, I will examine two possible reasons for this fragility.

Replication and performance issues aside, variations of the dividend schedule experimental design have been a popular choice to test a wide variety of hypotheses. A selection of some of the early examples is described below.

Ang and Schwarz (1985), tested different levels of information provided to insiders and altered the presence of common knowledge in a dividend schedule experiment with three states and two trader types. Their goal was to test for differences in behaviour between markets populated with speculative traders and markets with conservative traders. They found markets with the less risk-averse speculators had higher volatility, but prices converged faster to either the prior information equilibrium or rational expectations equilibrium.

Copeland and Friedman (1987) tested how information arrival and information content could affect market prices. They used a base setting with three trader types and two states. Each trader received information about the state at some point over each trading round. In their homogenous content setting, the state was the same for every trader type. In the heterogeneous content setting, each trader type could experience a different state. The heterogeneous environment increased the size of the aggregate level signal space. The release of information was also sent out simultaneously or sequentially. Traders started trading rounds with knowledge of their payoff structure but were unsure of others' structure. Overall, they found evidence in favour of rational expectations equilibrium over private information equilibrium. They also found prices were closer to the strong-form informationally efficient (fully revealing) rather than semi-strong.

Copeland and Friedman (1991) tested for partial revelation between strong-form and semi-strong price efficiency in a similar experiment with an information acquisition component. They again found support for fully revealing equilibrium prices. Still, the partial revelation equilibrium performed as good as any other benchmark in terms of asset allocation, the market price for information, and the allocation of information acquired.

In an extension of the two previous studies, Copeland and Friedman (1992) added an information auction to the 1987 experiment and additional modelling and tests to the 1991 study. They found support for the fully revealing rational expectations equilibrium in simple environments where traders could more easily infer the information of others. They found support for the non-revealing rational expectations in their more complex environment, along with higher prices for information.

In a similar experiment, Sunder (1992) used the Plott and Sunder (1982 1988) design and added a costly information component. In series A, a fixed quantity of information was auctioned. In this setting, asset prices reflected the fully revealing equilibrium prices, and the price of information eventually fell to zero. In series B, the price of information was fixed; in this setting, prices were less informative than serries A. In both configurations, the net profits of the information acquires was similar to the uniformed.

Forsythe and Lundholm (1990), extended Plott and Sunder (1988) by making all market parameters common knowledge, running experiments over two nights to give subjects experience, and increasing the number of traders participating in a market. They found that trading experience and common knowledge of dividends in tandem led to fully revealing equilibrium prices. However, neither was capable of this result in isolation. Their first night's results were indistinguishable from Plott and Sunder (1988).

More recently, discrete state-contingent payout schedule style assets have also been used in designs by Noeth et al. (1999), Hanson et al. (2006), and Corgnet et al. (2018). The popularity of such a design reflects the ease of use and flexibility.

2.3.3 The Two-Urn Design

The two-urn design is more complex than dividend-schedule experiments. An early example of the basic design appeared in the informational cascades experiment of Anderson and Holt (1997). More recently, it has become a very popular design for information aggregation experiments. This design is also featured in my experiment on information aggregation in Ch 2. In this section, I will review the details of the design and some of the experimental studies that featured the two-urn design.

An example of the basic design is as follows: There are two urns, "Urn Black," with, six black and four white balls, and "Urn White," with four black and six white balls. Figure 2.1 illustrates these urns. A coin is flipped to determine which urn signals are drawn from. Drawing is done with replacement. Participants do not know which urn has been chosen. Subjects privately see balls drawn with replacement from the chosen urn and then trade an asset that pays a fixed amount ($10) if the underlying urn is the black urn.

Variations of the above design have been used in many different settings. However, most experimental studies using this design have at least one treatment that resembles the basic version. In fact, Page and Siemroth (2019) were able to compile data on the basic treatment from several studies. They then ran a novel structural model on each to estimate the amount of information aggregated into prices. They found that public information was almost always reflected into prices but less than %50 of private

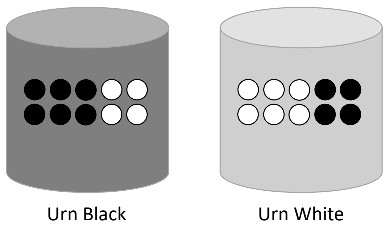

Figure 2.1: Basic Two-Urn Design

information made it into prices.

Included in Page and Siemroth (2019) was data from one of their previous studies, Page and Siemroth (2017). In this earlier study, they used the two-urn design along with an information acquisition component. They analysed the link between trader behaviour and trader characteristics. They found that individuals with more substantial endowments, inconclusive information, lower risk aversion, and less experience would acquire more private information. They proposed the success of prediction markets is linked to the propensity to over-acquire information as the presence of more information in the experimental markets led to more efficient prices even though traders could not recoup their costs.

Deck et al. (2010) were also included in the compilation. In their study, they used the two-urn design with the presence of manipulators and forecasters. The manipulators were paid if they could influence the forecasters. They found that well-funded manipulators could negatively affect a prediction market's ability to aggregate information, misleading those who made forecasts based on the information in prices. This result is different from that of Hanson et al. (2006), who used a dividend schedule design and that of Oprea et al. (2007), who used a two-urn design. In both of these papers, there was a public goods problem (not present in Deck et al. (2010)) in which the manipulators would prefer to free-ride on the manipulation of others.

Also included in Page and Siemroth (2019) was data from Fellner and Theissen

(2014). Who used the two-urn design to test whether the effect of overvaluation implied by the existence of short sale constraints increased with the degree of divergence of opinion among traders. Their measure of divergence of opinion was the cross-sectional variance of the asset's conditional expected value, which is non-linear in the probability of receiving the correct signal given the urn. They showed that short-selling constraints indicated higher prices, but overvaluation did not increase with the degree of opinion divergence. This result was based on subjects receiving a single signal only. It may be interesting to revisit this experiment with changes to the signal and parameter structure.

Halim et al. (2019) were included in the compilation. They used a two-urn experiment to test different network structures and the effect of social communication on prices and incentives to acquire information. They found their markets to under-react to privately held information, and that this under-reaction was decreasing in network density. However, this increase in efficiency from the increase in network density was counteracted by a reduction in information acquisition.

In addition to these papers, there is a set by Andrea Morone and co-authors that feature the two-urn design in the context of information aggregation, information acquisition, and public information. Hey and Morone (2004), used a two-urn design with a fixed price for information and a dynamic acquisition component where subjects could buy a signal at any time during trading. They found gross profits correlated with signal acquisition and net profits negatively correlated with signal acquisition. They found price variance to be higher and price efficiency to be lower when the balls' shade ratio was six to four, than when it was eight to two. This makes intuitive sense as the signals are more likely to be correct in the eight to two case. In both settings, they found evidence of "socially undesirable herding" and bubbles. In my experiment, I will refer to similar price patterns, as informational mirages rather than socially undesirable herding. I will also offer an alternative explanation to the informational cascade model as information is not continuously released in my setting.

In a similar set up with the inclusion of public signals, Ferri and Morone (2014) showed that the accuracy of the public signal reduces socially undesirable herding. In a follow-up experiment that also varied the accuracy of the public signal, Alfrarano et al. (2015) tested for an over-reliance on the public signal. They found the presence of a public signal crowded out demand for private information but left the overall informativeness of prices unchanged. The public signal dominated when the informativeness of the private signals was low. Their results warn of the danger of incorrect public signals in the two-urn setting.

2.4 Information Aggregation Pilot Sessions

To understand the lack of information aggregation in the two-urn design experiments discussed above, my supervisors and I ran a series of pilot studies examining several possible causes.

One feature of interest in the two-urn design was the non-linear state price mapping arising from the information structure. When ordering the aggregate information states on the x-axis, and the fully revealing equilibria on the y-axis, it is easy to see the shape of the mapping is not a straight line but a sigmoidal pattern.

The curved state price mapping gives price dependent marginal impacts for additional units of information (draws from the chosen urn). The non-linear structure of the problem adds a layer of complexity to the aggregation process. [7]

2.4.1 Perfect Foresight

We initially hypothesized that the lack of basic knowledge of the state price mapping (see Figure 2.2) could be generating some of the results of the two-urn studies. Without perfect foresight, subjects must compute (or infer) the equilibrium price for a given aggregate signal. The reverse is also true when backing out the aggregate signal from

[7] Features of the two-urn state price mapping with be discussed in detail, in the experimental paper.

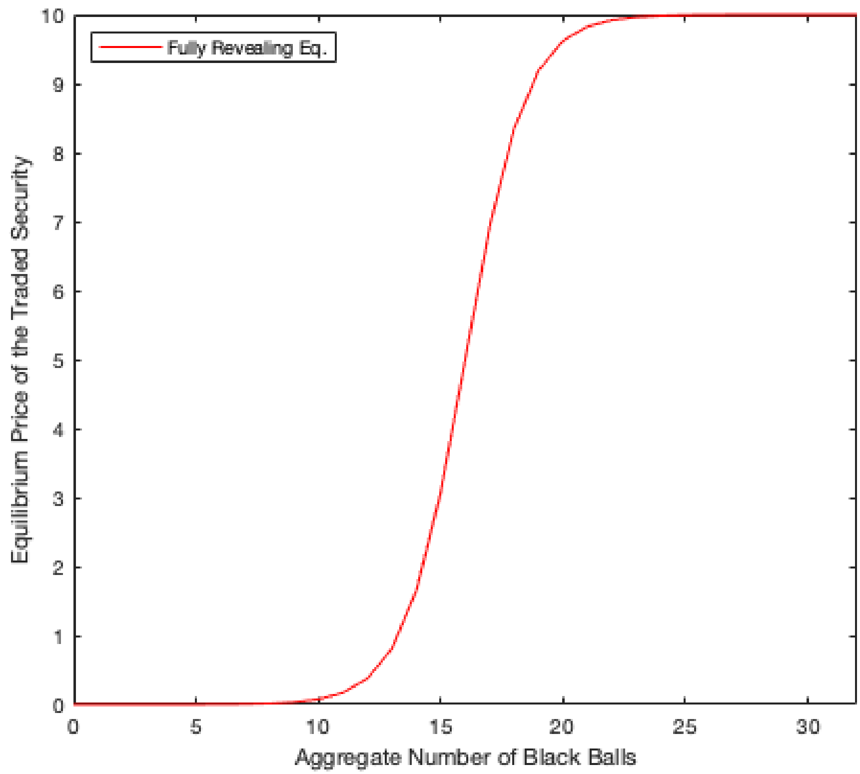

Figure 2.2: Two-Urn State Price Mapping

a given price. Similarly, they would need to infer the marginal impact of their private information. For these reasons, we believed information aggregation to be extremely challenging and that traders may have an intuition for the difficulty in this setting.

To test the effect of perfect foresight in the two-urn design, we ran the first half of the first pilot without giving subjects information about the state price mapping. In the second half, we provided them with an additional set of instructions and the perfect foresight state price mapping.

2.4.2 Signal Endowment Combinations

Besides testing perfect foresight, we explored whether portfolio balancing based on private information could stimulate or dampen the aggregation process for different combinations of signals and initial endowments. We wanted to know if information aggregation could be kick-started in a favourable direction under the right conditions. We were flirting with off equilibrium dynamics, but because the hypothesis was related to the problem's initial conditions, we didn't need to make many assumptions about the dynamics.

We thought aggregation might be stimulated in the correct direction if subjects with posteriors closer to the fully revealing posterior had substantial incentives to trade in a way that would pressure the order book to move mid-point prices toward the fully revealing equilibrium. This shift in the order book could precipitate further information aggregation.

For example, aggregation may be stimulated at the start of trading if subjects with high private values want to buy to balance their portfolios at a time when the underlying fully revealing equilibrium price is above the prior. A desire to buy, coupled with a higher valuation, should lead these subjects to remove liquidity from the book's ask-side. If subjects with lower private values start from fully hedged positions, they should demand compensation to be knocked off their balanced positions. Following this intuition, the liquidity provided on the ask-side should demand a premium. These effects should raise the mid-point price. Hence, kick-starting aggregation in the correct direction.

In contrast, aggregation may be dampened for the same high priced fully revealing equilibrium if those with high private values start from a hedged position while those with low private values wish to sell. In this case, those with high values would demand a discount to be knocked off their balanced positions. If this discount is substantial, information aggregation may be kicked off in the wrong direction when those with low private valuations wish to sell.

To test this hypothesis, we created a variety of signal endowment combinations for different equilibrium prices (see Appendix A.1). We maintained the same aggregate quantity of tradable assets across all treatments. Table 2.2 contains a basic description of the treatments.

Table 2.2: Pilots 1 and 2 Design

Period	FRE Price	Perfect Foresight Mapping	Predicted Signal Endowment Combination Effect
1	$1.65	No	Dampen Aggregation
2	$5.00	No	No Effect
3	$8.35	No	Dampen Aggregation
4	$1.65	No	No Effect
5	$8.35	No	Stimulate Aggregation
6	$8.35	Yes	No Effect
7	$1.65	Yes	Dampen Aggregation
8	$5.00	Yes	No Effect
9	$1.65	Yes	Stimulate Aggregation
10	$8.35	Yes	Stimulate Aggregation

2.4.3 Parameters for Pilot 1 and 2

Security Design

To generate the signal endowment combinations, we used two assets: stock-A and stock-B. Stock-A was tradable while the stock-B was non-tradable. Stock-A paid $10 if signals were drawn from the urn with a black ball majority, while stock-B paid $10 if draws were from the urn with a white ball majority. The two assets paid $10 with certainty when held in tandem.

Endowing some subjects with only stock-A and others with only stock-B allowed hedging incentives to create liquidity on both sides of the stock-A market. The non-traded securities, and complementary nature of the two assets, created a natural reason for trade and avoided the no-trade theorem of Milgrom and Stokey (1982). The two-asset security design was first used by Bossaerts, Frydman, and Ledyard (2014) in an asymmetric information (information amplification) experiment.

Individual and Aggregate Information Signals

Each participant was given two draws from the chosen urn, similar to the base case in other studies.

If individuals observed more draws, their posteriors would become more extreme and have less variation. The aggregation problem would become less impressive. The price implied by the average individual signal would become closer to the expected equilibrium price implied by the size of the aggregate signal. In simple terms, if individuals were shown more balls, they would almost know with certainty which urn was chosen. The price implied by the average private signal would be very close to the equilibrium price, and the variance of individual posterior means would decrease substantially.

In aggregate, the number of draws given to subjects also affects the conditional distribution of aggregate signals. More balls in aggregate increases the expected difference between the total number of black and white balls observed. The fully revealing equilibrium price is a function of this difference, and when the difference is large, the price is either close to zero or one. If we gave subjects more draws, the probability of observing an interior equilibrium price close to the midpoint in Figure 2.2 would decrease.

2.4.4 Number of Participants and Aggregate Information Signals

The expected difference between the total number of black and white balls would also increase if we increased the total number of traders in the market, keeping the private signal size constant. In this case, equilibrium prices would contain more information and become more extreme on average, while traders would become relatively less informed. The gap between the average private signal and the fully revealing equilibrium price would increase on average.

Increasing the number of traders would potentially increase the effect of the paradox discussed in Hellwig (1980). In a two-urn perfect foresight equilibrium, larger signals with extreme prices provide a stronger signal about the underlying urn. Prices near the boundaries also have smaller marginal price impacts for additional draws. In these equilibria, individuals can garner more information from prices and have less incentive to follow their signals; however, if individuals do not use their signals, it is unclear how prices become informative.

For these reasons, we chose to start with ten traders with two draws each in the first two pilots. We then increased this number to twenty in pilots three and four (discussed shortly) and settled on sixteen participants for the remainder of the experiments.

Aggregate Signals

We set most aggregate signals close to the expected aggregate signal conditional on the number of traders and chosen urn. We also included counterfactual aggregate signals that were half black and half white.

Aggregate Risk

If the total number of stock-A and stock-B outstanding are not equal, then aggregate consumption is not the same across states. There exists aggregate risk, and the level of risk aversion influences the market risk premia. Aggregate risk means general equilibrium forces are at play, and in theory, the individual risk premia should be correlated with aggregate consumption risk. If we momentarily ignore aggregate information and assume the total number of stock-A is greater than the total number of stock-B, then the risk premia on stock-A should be larger than that of stock-B. Stock-A should be cheaper than stock-B. The importance of aggregate risk in information aggregation experiments is discussed again following the first pilot results.

Aggregate Risk and Signal Endowment Combinations

If equilibrium prices are not equivalent to the fifty-fifty prior, then there must be more individuals with Black-Black signals than White-White signals or vice versa. In either case, imposing signal endowment correlations yields aggregate risk if the per-capita holdings of each group are equal. There simply cannot be an unequal number of Black-Black and White-White signals, zero aggregate risk, the desired signal endowment combinations, and uniform wealth across signal groups. One of the parameters must be free to compensate for the rigidity of the others.

This notion may appear complicated but is quite simple. To illustrate the point, I have provided a short logical example.

Assume every trader has either stock-A or stock-B. If the traders with Black-Black signals hold an equal number of stock-A and stock-B as a group, and the traders with Black-White signals as a group also hold an equal number of stock-A and stock-B; and, the traders with White-White signals want to balance their portfolios by being either net buyers or net sellers of stock-A (as in the signal endowment combinations) then as a group they cannot hold an equal number of stock-A and stock-B. Therefore, based on the traders' endowments, in aggregate, the quantity of stock-A does not equal stock-B, and there is aggregate risk.

In our first two pilots, we focused on maintaining a uniform wealth distribution across signal groups (according to the fifty-fifty prior). To do this, we let aggregate risk vary while generating our signal endowment combinations. [8]

Aggregate Liquidity

To control for potential liquidity differences when testing aggregation in a variety of settings, we fixed the aggregate quantity of the traded security to forty stocks in total.

[8] The other studies discussed above had only one traded security, so aggregate risk was always present.

This was similar to the two-urn studies mentioned above. By fixing the amount of stock-A and varying the signal endowment combinations, we unintentionally imposed a relationship between aggregate risk and the traded security. Again, this will be discussed in the presentation of the results and subsequent design changes.

Pseudo-Randomness Across Periods

In each period at fully revealing prices, some individuals were worse off than others in terms of the fully revealing equilibrium value of their initial endowments. According to the fifty-fifty prior, they all started with the same wealth, but their endowments were not of equal value at equilibrium prices. We decided that it would be fair if, on average, they received similar starting positions throughout the experiment.

Similarly, the signals individuals received in each period could be unbiased or biased towards or away from the fully revealing equilibrium. Again, it would be fair if subjects received signals across the entire experiment that were not biased in one way or another.

To assign endowments and signals fairly, we designed a simple algorithm. It first created a label for each subject in the experiment and then randomly assigned them to an endowment signal pair for each period. The algorithm then checked to see if their equilibrium priced endowments and signal correctness were in a range deemed fair. The algorithm was solving a combinatorial optimization problem, and the length of time required to solve such a problem was determined by the constraints of the fairness criteria (set manually). If the experiment were entirely fair, the constraints would be tight, and finding a solution would be difficult or impossible. If we didn't care about fairness, any solution would work. We chose a fairness criterion that was close to the maximum but allowed the algorithm to find an answer in a reasonable amount of time.

Endowments, Cash, Exchange Rate, and Payments

The endowments were chosen, such that the aggregate quantity of the tradable security was equal to forty. The cash distributed to subjects allowed them always to hedge their untraded risk at the maximum possible value of the traded asset. Subjects received pay according to their cash holdings and dividends from a randomly chosen period. The exchange rate ensured each individual would receive $45 Aud on average with a fixed $20 Aud component. The parameters for each pilot are in appendix for this chapter.

Instructions and Training

Instructions and training were a significant part of the piloting process. The experiments utilized a continuous-time open book market-based software system (flex-e-markets.com). The subjects needed to understand the task and how to trade in a financial market before commencing the experiment. Instructions described the task and trading system in the simplest terms possible. Each subject received a copy of the instructions that were read aloud and projected at the laboratory's front and rear. [9]

A hands-on practice session took place after presenting the first set of instructions. During which subjects received information signals and practiced making trades using the flex-e-markets system. During the practice sessions, members of the University of Melbourne BMM laboratory were present to answer questions about the trading system and other technical issues.

The first set of instructions and training process took approximately one and a half hours to complete.

The second set of instructions, released after the first half of the experiment, provided the state price mapping and examples of how to calculate the price of the traded

[9] Several diagrams in the instructions were borrowed or modified from Page and Siemroth (2019) and used to present the two-urn design.

asset from an aggregate signal. Copies of the instructions for each pilot are also in the appendix.

2.4.5 Pilot 1 and 2 Results

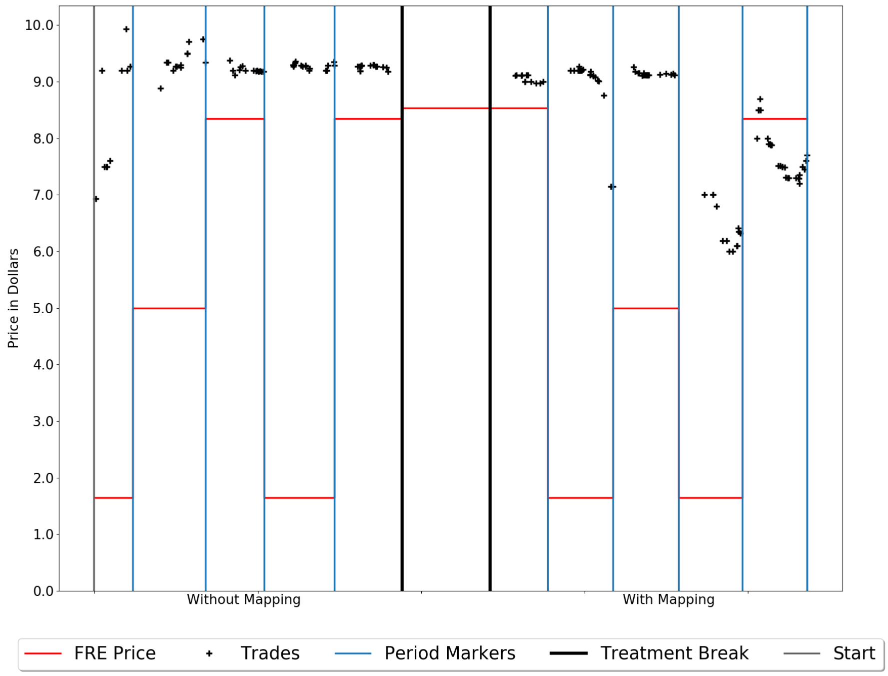

Figure 2.3: Pilot 1 Trade Prices

The first and second pilot experiments had identical parameterizations, and the results were nearly indistinguishable. The transaction prices are displayed in Figures 2.3 and 2.4 [10] Prices in all rounds were consistently high and well above the fully revealing equilibrium prices. Access to the state price mapping did not affect information aggregation, nor did the signal endowment correlations.

One possible cause of the high prices was the presence of aggregate risk. When the total quantity of tradeable securities did not equal that of the non-tradable securities, there was aggregate risk.

[10] We ran out of time do to technical difficulties in the second pilot.

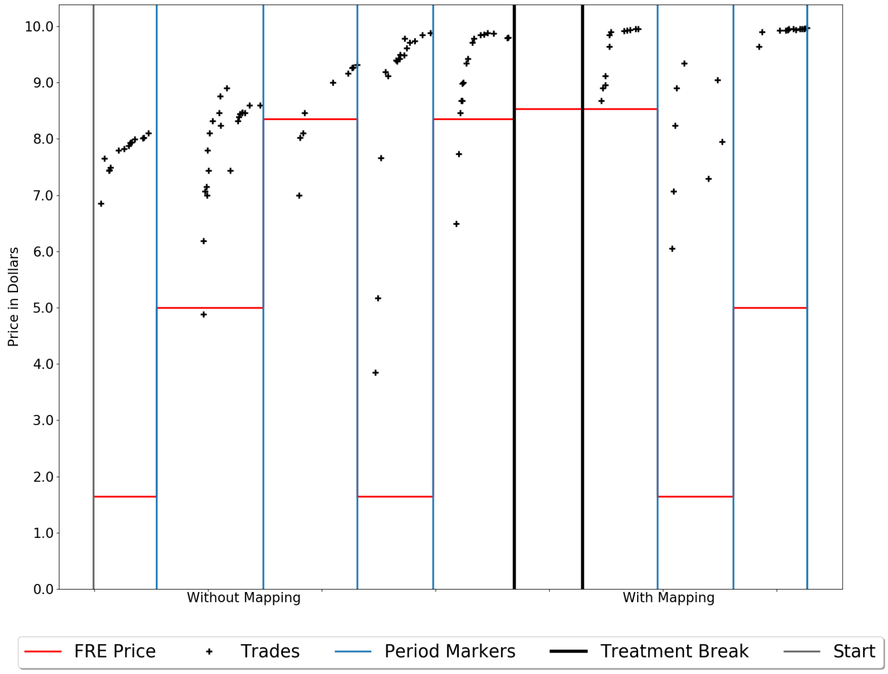

Figure 2.4: Pilot 2 Trade Prices

When reviewing the design, it became apparent that this was a major confounding factor. When ignoring the informational component, the aggregate risk levels implied stock-A was worth more than stock-B in eight out of ten periods. Only in period three and period nine did aggregate risk support lower stock-A prices. In period three, the information supported a high price of stock-A. Only in period nine, which had the lowest prices in both pilots, did the aggregate information and aggregate risk imply a low price of stock-A.

Allowing aggregate risk to vary freely without concern was a major oversight and a reminder that macro-level general equilibrium characteristics can be pervasive in even such a simple environment. These forces should not be underestimated when designing information aggregation experiments.

2.4.6 Pilot 3 Design

Number of Traders

We designed the third pilot for twenty subjects to see if information aggregation might improve with more traders. [11] The increase in traders increased the length of the aggregate signals. On average, fully revealing equilibrium prices were more informative about the underlying urn. We chose equilibrium prices and aggregate signals to be close to the expected aggregate values conditional on the true urn. We included the same counterfactuals as the first two pilots.

By increasing the number of participants to twenty, we increased the total quantity of the tradeable security to eighty from forty keeping the per capita quantity the same.

Perfect Foresight State Price Mapping

We provided subjects with the state-price mapping and improved the examples.

Language

We changed the term "urn" to "bucket" in the instructions. Talking about urns and drawing balls is reminiscent of a class in probability and may have caused anxiety for some subjects. Urns also hold the ashes of dead people; so, the term bucket is less morbid. Finally, subjects with English as a second language were more likely to have bucket as part of their vernacular.

[11] There were also designs for twelve and sixteen participants if subjects did not show up.

Signal Endowment Correlations

We were still interested in testing the effect of different signal endowment combinations, but this time we controlled for aggregate risk. In doing so, we lost the ability to hedge two out of three signal groups (see the above example). The signal endowment combinations effectively became signal endowment correlations as Black-Black and White-White signal groups had complementary endowments. The correlation treatments are in appendix A.3.

Information Structure

After the first two pilots, we speculated that the structure of the two-urn design affected the results. Aggregate risk was likely not to be the only factor driving the high prices in the first two pilots. Without speculating, or thinking of behavioural biases to explain the off equilibrium dynamics, we regrouped and focused on the differences in structure between information aggregation experiments and theoretical models.

In the Grossman model presented above, the average individual signal is equivalent to the traded security's true value. It is also equal to the average individual signal conditional on the true value. In Grossman's model, signals are mean independent, and the errors in individual signals cancel when averaging.

In the two-urn design, private values (individual posterior means based on private signals) do not average to the true value of the asset. Private signals are also mean dependent. The expected private signal is not the same as the expected private signal conditional on the true urn. These are subtle features, but they imply a difference between the fully revealing and private value equilibria. [12] This difference could be playing an essential role in the lack of aggregation in earlier experiments.

To test this, we designed a one-urn treatment where signals were drawn with re-

[12] See Chapter 2 for more details on this conjecture.

placement from a bucket of ten balls. Again, balls were either black or white. The traded security paid the proportion of black balls in the urn. In this case, signals were also mean dependent, but private signals and private value equilibria were closer to the fully revealing equilibria than in the two-urn design.

We ran ten periods of the two-urn design, and four periods of the one-urn design. The treatments are summarized in the table below and presented in appendix A.3 with the updated instructions.

Table 2.3: Pilot 3 Design

Period	FRE Price	Perfect Foresight Mapping	Signal Endowment Correlation	Urn Design
1	\$9.60	Yes	Negative	Two-Urn
2	\$5.00	Yes	Zero	Two-Urn
3	\$0.40	Yes	Positive	Two-Urn
4	\$9.60	Yes	Positive	Two-Urn
5	\$5.00	Yes	Positive	Two-Urn
6	\$0.40	Yes	Zero	Two-Urn
7	\$9.60	Yes	Zero	Two-Urn
8	\$5.00	Yes	Negative	Two-Urn
9	\$0.40	Yes	Negative	Two-Urn
10	\$7.00	Yes	Positive	One-Urn
11	&3.00	Yes	Negative	One-Urn
12	\$3.00	Yes	Positive	One-Urn
13	\$7.00	Yes	Negative	One-Urn

2.4.7 Pilot 3 Results

We can see in Figure 2.5 across the entire two-urn treatment the vast majority of transactions and final period transactions occur above a price of eight dollars. The exceptions are periods eight and nine, which had a negative signal endowment correlations. In these periods, individuals' information signals implied values of the traded security that were inversely correlated with individual endowments. Those who had signals suggesting a

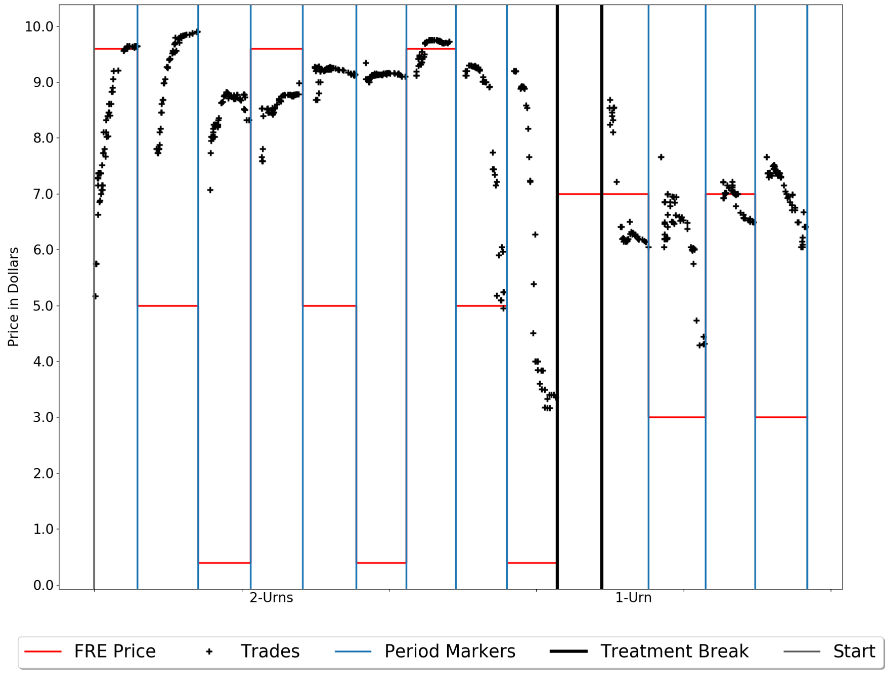

Figure 2.5: Pilot 3 Trade Prices

high value for the traded security had a low amount of the security in their initial endowments, and vice versa. Overall, these results did not present promising evidence for information aggregation for the two-urn case even after accounting for aggregate risk.

The one-urn prices were still high but rather promising in comparison. The final trade prices finished close to the fully revealing equilibrium in three out of the four rounds. These findings provided evidence of a meaningful difference between the one-urn and two-urn settings. As a result, we shifted our focus to testing the difference in structure between the two settings.

2.4.8 Pilot 4 Design

In the first half of the fourth pilot, we tested six periods of the two-urn design with zero signal endowment correlation and one period with a negative signal endowment correlation and low fully revealing equilibrium price. In the second half, we tested the same one-urn design from the third pilot, but with zero signal endowment correlation. All periods had zero aggregate risk.

We alternated trading stock-A and stock-B across periods to control for a potential behavioural bias toward holding stock-A.

We updated the instructions by including new diagrams. We created a web-based probability calculator for the posterior mean of urn-A being chosen given a set of signals.

2.4.9 Pilot 4 Results

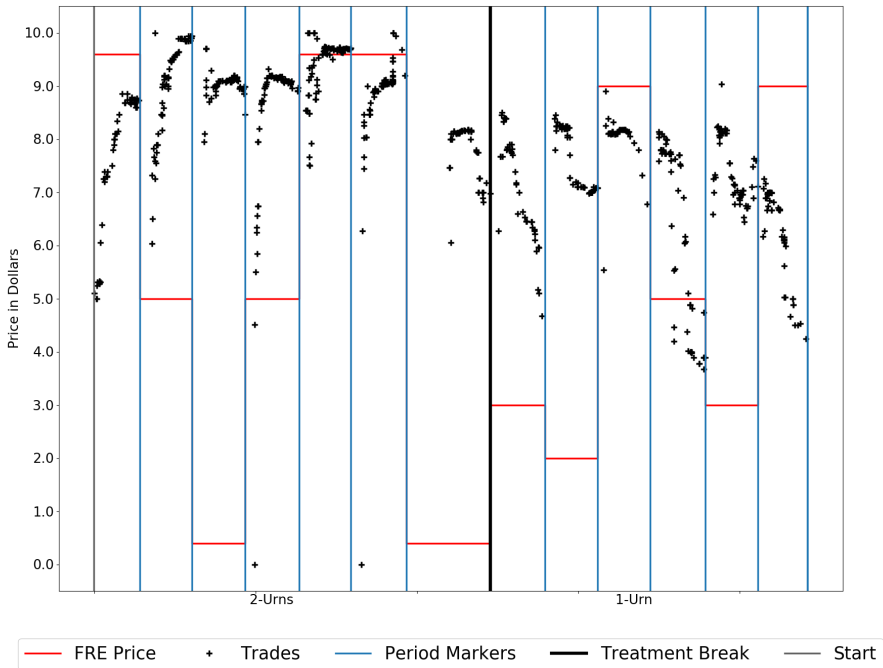

Figure 2.6: Pilot 4 Trade Prices

Again, in Figure 2.6 prices in the two-urn case were extremely high in all periods, including the case with a negative signal endowment correlation. Alternating the tradeable security made no difference.

In the one-urn setting, average transaction prices were still high but ended below five dollars in three out of six periods. The difference between the two environments is easy to see in the figure above. Finally, after four pilots with extremely high prices, we had something promising that was worth investigating.

2.4.10 Pilot 5 and 6 Design

To explore the difference between the two settings, we included a major addition to the design and finalized the experiment's structure. The final design will be explained in detail in the next chapter. I will provide a brief overview here.

The most radical change was the inclusion of a sequence of information auctions in each period. Auctions altered the concentration of information amongst traders in the market and led to competition between informed winners. They allowed us to control and observe a wider variety of settings within each urn design and created a richer environment for comparison.

Including information auctions during the experiment was a challenge. We created a web-based system that interacted with the flex-e-markets server and allowed the auction to run in a separate market. Auction information specific to each individual displayed on personalized webpages for each subject. The unified system allowed flex-e-markets software to compile data from trading and the auction into one source.

To ensure the auction and trading systems ran correctly in parallel, we used an algorithm to manage the sequence of periods and auctions. After the hands-on training, the algorithm was started, and the experiment ran on its own. [13]

Another major addition was the ability to short sell. In the previous pilots, subjects with zero initial holdings of the tradeable asset, and a low private value, could not sell any assets. In this case, they were restricted from taking action in accordance with their

[13] The web display and manager algorithm were implemented by Abhijeet Anand. Without him, the process would have been even more tedious.

private information. To remedy this, we allowed short selling of up to eight units of the traded asset. This option to short sell did not affect the lack of aggregate risk. A new unit of the tradeable was created when a short sale occurred. The short seller's holdings would also go negative, thereby neutralizing the addition in aggregate, and maintaining zero aggregate risk.

Most of the parameters remained the same in the new design, but the periods split into four separate ninety-second trading rounds with thirty-second buffers. We ran four periods of each setting.

We reduced the number of subjects to sixteen from twenty, to make the number closer to previous studies, ease the difficulty of recruitment, and lower the probability of no shows.

We simplified the language in the instructions to compensate for the additional complexity of the information auctions. We called the stocks: Stock-Black and Stock-White, and the buckets: Bucket-Black and Bucket-White.

We included an additional training period, which included information auctions.

In the one-urn setting, we increased the bucket size to two times the number of subjects and distributed all the balls. In aggregate, the distributed balls fully reveal the urn's composition, and there is no longer uncertainty based on the aggregate signal. When there is uncertainty, the posterior weights on the observed signal and the fifty-fifty prior (see Appendix A.7). In theory, eliminating aggregate uncertainty from the fully revealing equilibrium should not make a difference in our setting. There was zero aggregate risk in our experiment, and risk-neutral pricing should obtain. Simplifying the one-urn setup lowered the computational burden and created a linear state price mapping. Overall, it created a cleaner setting.

Aggregate Signals Across Treatments

When testing two different designs in an information aggregation setting, it may be unrealistic to keep all parameters constant across treatments. For example, the aggregate signals used could not be the same if we wanted to generate price variation in both treatments. Choosing signals to generate price variation in one setting would have led to the clustering of equilibrium prices in the opposing treatment. This occurs because of the difference in structure between the state price mappings. This should become clear in the second chapter.

To avoid price clustering, we used the previous aggregate signals from the two-urn treatments and independently selected aggregate signals for the one-urn treatment that generated price variation. Again, in theory, the aggregate signals selected should not affect aggregation performance.

2.4.11 Pilot 5 and 6 Results

Unfortunately, in the fifth pilot, there was an error in the informational display that meant our results had to be thrown out. Similarly, in the sixth pilot, the server where space was rented to run the market software kept crashing. The server was located on the east coast of the United States. The pilot occurred in the morning in Australia and the evening in the US when American web traffic is known to increase. The NBA playoffs were also on at the time of the pilot when the server was crashing. It is not clear if this was a factor, but it cannot hurt to take precautions in the future. [14] As a result, data from the sixth pilot could not be used.

The back to back mishaps led to extensive dry run testing of the experiment, which should have been conducted earlier. Doing a dry run each time a design change is implemented is necessary to avoid such costly mistakes. It also gives the experimenter

[14] A simple solution would be to watch all playoff games and schedule pilots during more appropriate time slots.

hands-on experience.

2.5 Conclusion

In this essay, I discussed the history of the theory of information aggregation, presented a classic model, reviewed the experimental literature and provided insight into the piloting process used for the next essay. Hopefully some of the ideas presented here will inspire future experimental and theoretical studies on information aggregation. There is still a lot left to explore.

Chapter 3

Efficiency and Stability of Information Aggregation in Markets

Information aggregation is the ability of financial markets to consolidate dispersed information into prices. Depending on the context, the implications of new information vary greatly and, hence, play an integral role for the aggregation process performed by markets. However, in the literature, the nature of information is often chosen for analytical tractability and is rarely the focal point of analysis. Thereby, a fundamental aspect of information aggregation has been left on the sidelines. We shall argue that the stability of information aggregation and the efficiency of resulting prices crucially depend on the underlying information structure.

The term "information structure" used throughout this paper encompasses the process that generates private information, the corresponding context of private information, and the Bayesian formulation used for aggregation. Whenever Bayesian updating is performed on the basis of aggregate information, the respective posterior distribution constitutes the informationally efficient benchmark. The underlying information structure determines the characteristics of the mapping from aggregate information into equilibrium prices.

We show that changes in information structure can lead to vastly different levels of performance with regards to price efficiency. This is consistent with empirical observations from the field. Indeed, there are situations where markets appear to efficiently aggregate information, with prices accurately revealing dispersed information. Examples include inflation (everyone knows what local price changes are) and industrial production (individual producers know what they have produced). Upon announcement of the aggregated data (inflation; industrial production), market prices (nominal bonds; stocks) hardly react anymore. For empirical evidence on inflation and bond prices, see, e.g., Huberman and Schwert (1985); Kandel e.a. (1993). We shall refer to this as the "inflation" case.

In contrast, there are situations where announcements of dispersed information often lead to big surprises, e.g., in the context of referenda (Brexit) or elections (Trump) (see, e.g., Bossaerts F. e.a., 2019; Wagner, Zeckhauser and Ziegler, 2018). This occurs despite the fact that everyone observes their own piece of information (they know how they are going to vote) and perhaps that of neighbours, family, and friends, but evidently, competitive markets can get the aggregation completely wrong. We call this the "Brexit" case.

We provide a simple theoretical argument for this discrepancy. For that, we introduce two new stability concepts to the study of *Rational Expectations* (RE) equilibria. First, we define "absolute stability" of a RE equilibrium. In a RE equilibrium, all agents know the mapping from states (aggregated information) to prices and use this mapping to infer what others know collectively. Absolute stability concerns the effect of inference among agents and the feedback effect of this inference on prices when a signal is lost or garbled, or someone makes a small mistake in trading on private information. The inflation case is stable: a garbled signal or a "fat finger" mistake by a trader does not fundamentally alter what agents infer from the price, and hence, hardly feed back into further price changes. The Brexit case is unstable in that equilibrium prices may change dramatically upon garbled signals or mistaken actions. As we will argue below, this is because its equilibrium mapping is highly nonlinear.

Second, we introduce the notion of "relative stability." Even rational agents cannot know which equilibrium obtains without making assumptions about the behavior of others. Are other agents playing the RE equilibrium, so that prices can be trusted to reveal the aggregated information? Or are other agents ignoring information and prices, and hence, are prices as in a private-information-only, i.e., a Walrasian equilibrium (PE)? Relative stability obtains when the mappings from aggregate information to prices are close across the two equilibria. In that case, the agent can trust prices (to infer information), no matter what she assumes about the information inference of others. Relative instability emerges when the two mappings are very different. Treating a price in a PE as if it were generated in a RE equilibrium would lead to a completely erroneous inference. As a result, prices cannot be trusted, and maladaptive trust may lead to entirely erroneous prices, i.e., to information "mirages" (Camerer and Weigelt, 1991).

An experiment allows us to assess the veracity of the theory. We find that prices are far less reliable in the Brexit case, and, in fact, frequently generate mirages. In the

Brexit case, winner-take-all contracts are traded. They pay \$10 if the majority of agents have private information in one particular direction, and \$0 otherwise. In contrast, in the inflation case, while biases can still emerge, price quality is far closer to revealing the correct information. There, the asset pays (a comparable multiple of) the simple average of the dispersed information.

To determine whether participants can sense when they can or cannot trust prices to reflect aggregate information, we elicit their willingness to pay to know whether prices are close to being fully revealing.[1] We find cumulative bids to be significantly higher in our Brexit relative to our inflation case. Hence, our participants sense the presence of relative instability. As time progresses, bids decrease gradually, correctly anticipating improved information aggregation. Consistent with Bossaerts, Frydman and Ledyard (2014), net of auction payments, the winners of this bidding process, on average, do not earn more than other traders.

Prices almost invariably end up to (partially) reveal aggregate information, regardless of the information treatment. We attribute such informative prices not solely to information aggregation, but also to competition among auction winners. The presence of insiders who merely have access to information about the quality of prices (and the direction in which they may be off) steers prices in the right direction. To borrow terminology from prediction markets, a robust way for the "wisdom of the crowd" to be reflected in prices is to have wisdom "in the crowd." Prediction markets in the field seem to rely on the same phenomenon (Bossaerts F. e.a., 2019).

The remainder of the paper is organized as follows. We develop our theoretical arguments in Section 3.1. Section 3.2 presents the design of our experiment and results. Finally, we conclude and provide avenues for future research in Section 3.3.

[1] An alternative would be to ask participants how much they are willing to pay to know the signals of others; however, we attempt to stay as close to the competitive equilibrium framework as possible, where agents are generally agnostic about individual idiosyncrasies, including signals. In a competitive equilibrium, all that matters is the price, and in a rational expectations setting, the mapping from states to prices.

3.1 Theoretical Framework

3.1.1 Equilibrium Concepts

We utilize two competitive equilibrium concepts. First, we rely on the Fully Revealing Rational Expectations Equilibrium (FRE), which constitutes the informationally efficient benchmark.[2] Here, prices correctly reflect the collection of all privately held information and can be inverted to infer the aggregate information.

We shall ignore risk attitudes by focusing, without loss of generality, on economies without aggregate risk. Therefore, in the FRE, prices equal posterior expectations conditional on *aggregate* information.

Second, we use the private-information (or Walrasian) equilibrium (PE), where agents trade solely on the basis of their private information; they do not utilize the information contained in prices. In the absence of aggregate risk, PE equilibrium prices equal *the average* of individual posterior expectations conditional on *private* signals only.[3]

3.1.2 Information Structures

The classic signal structure used in theoretical research on information aggregation is that of Gaussian, independent draws given the true (fundamental) value of the asset, which itself is a Gaussian variable (see, e.g., Grossman, 1976, 1978, 1981, or Hellwig, 1980). Specifically, agent i receives a signal s_i which equals the asset's true value v plus

[2] The FRE is informationally efficient, i.e., supports the Efficient Markets Hypothesis, but may not be Pareto optimal (Hirshleifer, 1971). When markets are complete, however, the Rational Expectations (RE) equilibrium is generically fully revealing – thus, also an FRE – and informational and Pareto efficiency coincide (Grossman, 1981).

[3] In between the FRE and PE is the noisy rational expectations equilibrium (NRE), which partially reveals aggregate information. A NRE is usually modeled with a noisy (unobservable) supply of the traded asset, which prevents prices from being fully revealing (see, e.g., Grossman and Stiglitz, 1980; Hellwig, 1980; Diamond and Verrechia, 1981).

noise ϵ_i:

$$s_i = v + \epsilon_i,$$

where ϵ_i is mean-independent of v and Gaussian with mean zero. Notice that, if somehow the price can be expressed as an unbiased function of the average signal across agents, by the law of large numbers, the price will reveal the true value v as the number of agents becomes large enough. To ensure this, one needs demand schedules that are linear in the signals, which is what one obtains by assuming Gaussian values and signals in combination with exponential utility. The above literature makes heavy use of the combination of Gaussian noise and exponential utility.

Here, we deviate from this Gaussian setting and instead consider an information structure where signals are composed of the most fundamental unit of information called 'bits.' A bit is simply an observation of a binary variable; a coin flip, a light switch, or, in our case, a black or white ball. Specifically, for our experimental implementation, we limit individual signals to two balls randomly drawn from an urn filled with black and white balls.[4] Importantly, using signals composed of two binary observations fixes the entropy of individual signals and allows us to focus on information aggregation. We consider two different signal generating distributions, and, hence, two different underlying information structures.

The One-Urn Problem ("Inflation")

The one-urn problem is motivated by the theoretical literature mentioned above. The problem starts with a single urn filled with black and white balls. There are enough balls in the urn such that, when drawing without replacement, every agent receives exactly two balls, which constitute her private signal. After receiving their signals, agents trade an asset that pays the initial proportion of black balls in the urn, times a fixed dollar constant ($10 in the experiment). The problem is illustrated in Figure 3.1. Similar to the standard theoretical setting, a simple average of individual signal proportions gives

[4] During the actual experiment, we refer to the "urns" as "buckets." We prefer this term since it is less technical.

the correct value of the asset.

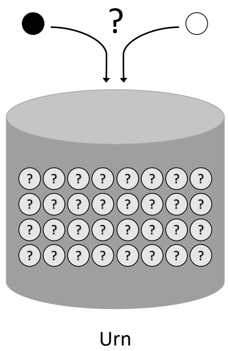

Figure 3.1: Example of one-urn problem

The Two-Urn Problem ("Brexit")

The two-urn problem is inspired by the traditional setting studied in the experimental literature on information aggregation (e.g., Plott and Sunder, 1988; Corgnet, DeSantis, and Porter, 2015; Page and Siemroth, 2018; Alfarano, Camacho, Morone, 2015) as well as by traditional prediction markets (e.g., Dreber e.a., 2015). The signals are again binary but the signal generating process is different. It starts with two urns filled with an even number of black and white balls. One urn, let us denote it by "Urn Black," contains a majority of black balls, while "Urn White" contains the mirroring majority of white balls. A coin is flipped to determine which urn signals are drawn from. Every agent receives an independent private signal about the selected urn that consists of two balls drawn with replacement. After receiving their signals, agents trade an asset ("Stock Black") that pays a fixed payoff ($10 in the experiment) if the underlying urn corresponds to "Urn Black" and zero otherwise. Figure 3.2 illustrates the two-urn problem with ten balls, of which six are black and four are white and vice versa.

In this setting, both the total count of black balls and the average signal (average number of bits indicating the draw of a black ball) are sufficient statistics for calculating the expected payoff of the traded asset. Importantly, however, the average signal no longer equals the traded asset's (scaled) FRE price. Instead, its FRE price is a non-

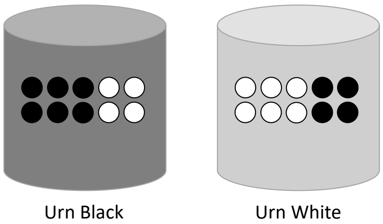

Figure 3.2: Example of two-urn problem

linear function of the average signal. To see this, let us first consider the functional form of both problems' respective Bayesian posteriors.

3.1.3 Bayesian Posteriors

Throughout this paper, we assume agents to obey Bayes' law when computing the conditional probability of event A conditional on event B:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}.$$

Let N denote the aggregate number of balls distributed across all agents, n the number of balls contained in the signal observed by an individual agent (in our case, $n = 2$), K the aggregate number of black balls, and $k \leq n$ the number of black balls observed by the individual agent.

In the one-urn case, an agent's posterior belief regarding the total number of black balls, given her individual signal of k black balls, equals (see Appendix B.2 for details on the Bayesian calculations):

$$P(K|k) = \frac{\binom{K}{k}\binom{N-K}{n-k}}{\binom{N}{n}} \frac{n+1}{N+1}.$$

Therefore, just based on her signal, the expected payoff of the traded asset paying $\$x$

times the proportion of black balls in the entire urn equals:

$$\begin{aligned}\mathbb{E}\left[\text{Payoff}|k\right] &= \$x \times \frac{\mathbb{E}\left[K|k\right]}{N} \\ &= \$x \times \sum_{K=k}^{N-(n-k)} P(K|k)\frac{K}{N} \\ &= \$x \times \sum_{K=k}^{N-(n-k)} \frac{\binom{K}{k}\binom{N-K}{n-k}}{\binom{N}{n}} \frac{n+1}{N+1}\frac{K}{N}. \end{aligned} \tag{3.1}$$

In the two-urn case with a proportion $p > 1/2$ of black balls in Urn Black and a proportion of $(1-p) < 1/2$ black balls in Urn White, the posterior probability that a n-ball signal with k black balls was drawn from Urn Black is given by:

$$P(\text{Black Urn}|k) = \frac{p^k(1-p)^{n-k}}{p^k(1-p)^{n-k} + p^{n-k}(1-p)^k} \tag{3.2}$$

$$= \left(1 + \left(\frac{p}{1-p}\right)^{(n-k)-k}\right)^{-1}. \tag{3.3}$$

Based on one signal only, the expected payoff of the traded asset paying $\$x$ if Urn Black is the underlying urn, therefore, simply equals:

$$\begin{aligned}\mathbb{E}\left[\text{Payoff}|k\right] &= \$x \times P(\text{Black Urn}|k) \\ &= \$x \times \left(1 + \left(\frac{p}{1-p}\right)^{(n-k)-k}\right)^{-1}, \end{aligned} \tag{3.4}$$

where $(n-k)$ is the number of white balls in the signal.

3.1.4 The Fully Revealing RE Equilibrium

If we condition the posterior probabilities not on the small signals received individually (n), but on the aggregate signal across the economy (N), we can recover the FRE.[5] The FRE equals the corresponding expectation scaled by the (relative) payoff amount.

[5] Here, we assume risk neutrality. In the experiment, this assumption can be justified because there will be no aggregate risk.

In the one-urn problem, according to Equation (1), for $n = N$ and $k = K$, the posterior expected payoff simply equals the ratio of K (total number of black balls) to N (total number of balls), multiplied by the payoff scaling factor $\$x$. Therefore, the FRE price, which equals the expected payoff, is a linear function in K. This is depicted by the blue straight line in Panel (a) of Figure 3.3.

For the two-urn problem, the expected payoff of the traded asset is a highly non-linear function in the difference of black and white balls. If, out of $n = N$ balls in total, the number of $k = K$ black balls increases by one, the exponent in Equation (2) decreases by two. This non-linear sensitivity is depicted by the red line in Panel (a) of Figure 3.3.

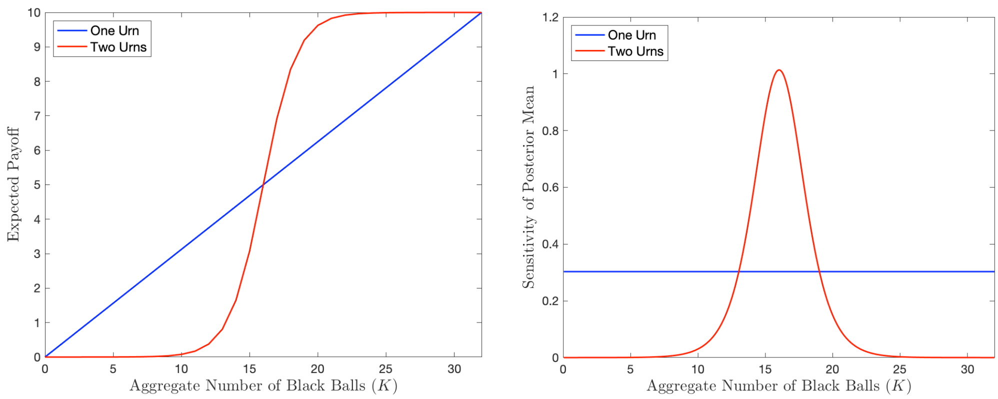

(a) Posterior Expected Payoffs (b) Sensitivities of Posterior Expectations

Figure 3.3: Posterior Expectations Based on Aggregate Information

Illustrations are based on a aggregate signal of $N = 32$ balls and a maximum payoff of \$10. For the two-urn case, it is assumed that both urns contain ten balls: Urn Black (White) consists of six black (white) and four white (black) balls (see Figure 3.2).

3.1.5 Absolute Stability

Absolute stability measures the sensitivity of the FRE price to small changes in the aggregate signal. Here, such a change corresponds to flipping a one-ball signal from black to white or vice versa. In terms of entropy, flipping a signal (bit) is identical in both urn settings, i.e., it is the smallest informational change possible [6]. More generally, one can think of flipped signals to arise from garbled messages, or trading mistakes. In this study, we are not interested in why signals (bits) get flipped, but what happens in equilibrium. Moreover, we do not focus here on the equilibrium (or equilibria) that may obtain when agents are conscious of the presence of garbled messages or mistakes; we save that for future work.[7]

When there is absolute instability, there is higher price sensitivity, which directly translates into a bigger effect of flipped signals on inference based on prices. This, in turn, feeds back into observed prices. For instance, in the two-urn case shown in Figure 3.3, starting from a balanced aggregate signal and flipping two black into to white bits results in a drop of the expected payoff from \$5 to less than \$2. Effectively, a corresponding price drop would make observers believe that the signals were drawn from Urn White. Contrary, in the one-urn case in Figure 3.3, such a mistake would cause the expected payoff to only drop to \$4. In the latter case, the inference by observers is less affected, thereby reducing any potential feedback effects.

Absolute instability can be measured locally via the derivative of the posterior expectation with respect to the aggregate signal. [8] This is illustrated in Panel (b) of Figure 3.3. For relatively balanced aggregate signals, the sensitivity of the two-urn case is significantly greater than for the one-urn case. Hence, if aggregate signals tend to be close to the unconditional expectation (uninformed prior), a small variation in individual signals may change aggregate (market) beliefs (as reflected in prices) dramatically. In the extreme case, where the aggregate signal is balanced, we will show such variations may even lead to "mirages," that is, to aggregate beliefs that incorrectly converge to an

[6] Shannon's Entropy is a measure of information in a random variable based on the frequency of possible outcomes. It is expressed as, $H(X) = \sum P(x) \cdot log_2(P(x))$, where $P(x)$ is the probability of event x. This gives a measure in bits (e.g., A fair coin, has 1 bit of information). In both our cases, each raw signal contains 32 bits. If the probability of flipping a given bit is the same across settings, then the Shannon's Entropy of flipping a bit is the same.

[7] In game theory, the resulting equilibrium is referred to as Quantal Response Equilibrium. See McKelvey and Palfrey (1995).

[8] There are different ways to measure absolute instability globally across the state space. In our setting, given the two boundary conditions, an option may be the integral over the aggregate signal space of the absolute difference between the FRE and the straight line connecting the boundary conditions (i.e, $\int_0^{max_K} |FRE(K) - (a-b) \cdot \frac{K}{max_K}|dK$, where a is the maximum boundary value and b is the minimum.

unrealized state (i.e., incorrect urn composition).

It also follows from Panel (b) in Figure 3.3 that the two-urn sensitivity decreases to very low values near the boundaries. For the corresponding aggregate signals, the one-urn case actually is less stable than the two-urn case. However, such extreme aggregate signals are generally unlikely to occur, and also less interesting from an information aggregation perspective.[9]

3.1.6 Relative Stability

Even if no signals are garbled and no trading mistakes are made, agents still need to know what equilibrium they are in. In FRE, everyone correctly reads information from prices; in PE, nobody does. Whichever equilibrium obtains, for the individual agent it always pays to read information from prices. However, in PE, the mapping from aggregate information to price levels can be very different relative to FRE. In this case, depending on which mapping an agent relies upon, her inference will diverge substantially. This leads to the concept of relative stability: how different are state-to-price mappings in PE relative to FRE? Relative instability leads to tension between equilibria, and hence, may erode trust in prices, even to the degree that a rational agent abstains from reading any information from prices. [10] [11]

To assess relative stability in our two settings, in Panel (a) and (b) of Figure 3.4, we plot the PE equilibrium mapping against the FRE equilibrium mapping. The difference in relative discrepancy between the two cases is immediate. While, in the one-urn case, both mappings are linear and fairly close to each other, the respective mappings are distinctly different in the two-urn case. Hence, in the former, when incorrectly treating

[9] In our experiment, the color proportions for the two-urn case are fixed at 60%:40% and 40%:60%, respectively. This makes extreme aggregate signals very unlikely. Specifically, the chance that, out of 32 balls, either less than ten or more than 22 have the same color is less than one in eight.

[10] Another possibility is that agents form priors over the two cases, and then infer information from prices using hierarchical Bayesian updating. Absent strong empirical evidence on the frequency with which the two types of equilibria obtain, agents' priors necessarily are entirely subjective.

[11] We can measure relative instability locally for a given aggregate signal as the difference between the FRE and PE. Globally relative instability can be expressed as the integral of the absolute difference between FRE and PE over the entire state space (i.e, $\int_0^{maxK} |FRE(K) - PE(K)|dK$).

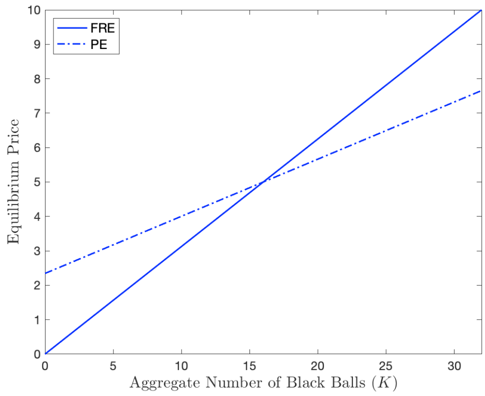

(a) One Urn

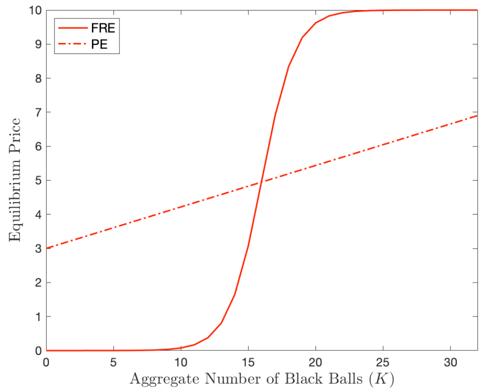

(b) Two Urns

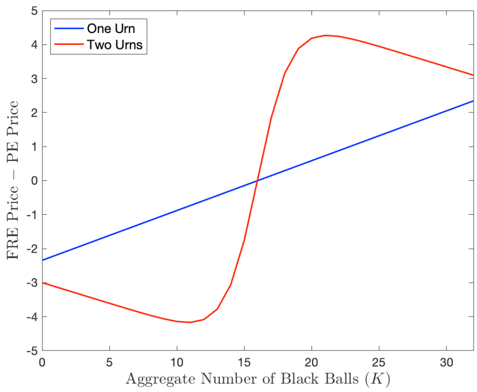

(c) Difference between Equilibria

Figure 3.4: Fully Revealing vs. Private-Information Equilibrium

Illustrations are based on a aggregate signal of $N = 32$ balls and a maximum payoff of \$10. For the two-urn case, it is assumed that both urns contain ten balls: Urn Black (White) consists of six black (white) and four white (black) balls (see Figure 3.2).

a price as if it were to come from a FRE, agents would infer, from the price, an aggregate number of black balls that is not far from the truth. In sharp contrast, corresponding inference mistakes are likely to be of much higher magnitude in the two-urn case. Panel (a) of Figure 3.4 plots the differences in equilibrium mappings for the two settings. Globally, the difference between equilibria is larger (in absolute value) for the two-urn case.

3.1.7 Absolute and Relative Stability

If we take the two notions of stability together, it is easy to see why "mirages" can emerge in the two-urn case, why they are less likely to occur in the one-urn case, and why prices in the former may be more volatile than in the latter. A *mirage* emerges if the price looks informative, i.e., it deviates substantially from the unconditional expected payoff, yet reveals the wrong information, i.e., information opposite to the truth (Camerer and Weigelt, 1991).

Absolute instability implies that small informational changes, caused by whatever reason, can translate into large price deviations from the FRE. For example, in the two-urn setting, if the true aggregate number of black balls is just one above the 50:50 distribution and two bits worth of information are flipped from black to white, then prices may drop below the unconditional expectation, potentially generating a mirage. However, if the price swing leads individuals to doubt that they are in an FRE, and relative instability makes them stop reading information from prices, prices are likely pushed to PE levels, i.e., closer to unconditional expected payoffs. If so, the mirage effect is contained but, ultimately, prices will be relatively uninformative.

Therefore, at the core of our analysis lies *trust in prices.* If agents trust prices (to be in FRE), absolute instability may lead to mirages, i.e., prices will look informative, but are entirely wrong. If agents do not trust prices, e.g., because of relative instability, prices will become relatively uninformative. In the most extreme case, e.g., if agents believe prices to be completely incorrect, they may even stop considering their private signal when trading. In this case, prices likely become devoid of information, eventually converting to unconditional expected payoffs. Clearly, trust in the nature of the equilibrium that agents face is also central to our predictions.

Trust is a fluid phenomenon. Importantly, the need to trust depends on the information setting. Our arguments are therefore fundamentally different from those that lead to a cursed equilibrium (Eyster, Rabin, Vayanos, 2019). There, a *fixed* proportion of "cursed" traders in the market never learn from prices. Trust is a more flexible concept: depending on the information structure, agents may or may not trust prices to reveal the right information, rather than always mistrust prices obstinately.

Consequently, to ascertain empirically whether trust is indeed crucial to understand the efficiency and reliability of information aggregation in the one-urn case vs. the two-urn case, we need to measure traders' trust in prices. We can quantify trust in two ways:

1. **Information auctions.** If agents do not trust prices, they are willing to bid for information as to whether prices are close to revealing aggregate information, and if not, in what direction they are biased.
2. **Agents' holdings.** If agents do not trust prices, they may trade to allocations that are not fully hedged, but are supported by their private signals. If agents do trust prices, we should see them holding hedged positions, as their risk aversion yields a preference for a balanced position near efficient prices (provided there is no aggregate risk, i.e., efficient prices equal conditional expectations).

We now turn to an experiment aimed at investigating whether the one-urn setting actually allows for better information aggregation than the two-urn setting, and, if so, whether this is because trust in prices is lower in the latter case.

3.2 Experiments

3.2.1 Experimental Design

We ran a series of real-time continuous double auctions using Flex-E-Markets, an online trading platform designed to perform like modern financial exchanges.[12] We ran a total

[12] See `http://www.flexemarkets.com`.

of eight *sessions*, seven with sixteen participants and one with fourteen participants. All sessions consisted of eight *periods*, each of which lasted for four *rounds* of simultaneous trading and bidding for information. There were four one-urn periods and four two-urn periods in each experimental session. In the two-urn periods, Urn Black (White) always contained six black (white) and four white (black) balls (see Figure 3.2).

At the start of each period, participants were endowed with eight shares of one of two complementary securities (Stock Black and Stock White), $80 in cash (in experimental dollars), and a two-bit signal; all of which expired at the end of the period. Stock Black and Stock White paid complimentary dividends such that one unit of both stocks held in tandem guaranteed $10 at the end of the period. In the one-urn treatment, one share of Stock Black paid the proportion of black balls in the aggregate signal and one share of Stock White paid the corresponding proportion of white balls, each multiplied by $10. In the two-urn treatment, one share of Stock Black paid $10 if Urn Black was chosen as the signal-generating urn and zero otherwise, while Stock White paid the complement.

The aggregate signals in both treatments were chosen to generate FRE price variation across the state space. In the two-urn case, the majority of aggregate signals were chosen close to the conditional mean of the underlying urn. However, we also included a counterfactual period in which half the balls were black. Overall our configurations of aggregate signals had high likelihood to occur in a random draw. This is important as otherwise participants could have questioned the the signal-generating process outlined in the instructions.[13]

All participants started with a risky portfolio so that risk aversion motivated trade. They were allocated sufficient cash to trade to a fully hedged position. We allowed trade in Stock Black but shut down the market for Stock White. This concentrated all hedging and information revelation into one market. In each period, there existed an equal

[13]For instance, in the two-urn case, if Urn Black with 60% black balls is chosen, aggregate signals consisting of at least as many white than black balls have only a 7% of occurring randomly. Thus, in two-urn periods, our most extreme configurations, i.e., aggregate signals with equally many black and white balls, represent rather rare outcomes. In the one-urn case, all possible configurations are equally likely, each with a chance of 3.125%. As such, our aggregate signals were no tail events, which prevented unnecessary confusion among participants (see Asparouhouva, Hertzel and Lemmon, 2009). Between cases, we used different signal configurations, as an identical range of aggregate signals across the two treatments would have imposed severe restrictions on FRE price variablity.

number of shares of both types of stocks, which completely eliminated aggregate risk. Furthermore, to control for the Hirshleifer effect, we counterbalanced signal allocations across participants, so there was no correlation between signals and endowments.

The absence of aggregate risk is a key feature of our experiment (see, e.g., Bossaerts, Frydman and Ledyard, 2014) that differs from previous studies using similar information structures (e.g., Page and Siemroth, 2018). Absent aggregate risk, efficient risk-sharing abrogates any risk premia and, thus, implies risk-neutral pricing. This is a crucial design feature when trying to assess markets' ability to aggregate information. If aggregate risk were present, the resulting risk premia would make the break down of observed prices into risk preferences and information difficult.

Another important design feature of our experiment is the inclusion of information auctions that ran in parallel to the securities market. As laid out above, these auctions are designed to provide an accurate real-time metric of trust in prices. We are not the first to include information auctions into experimental asset markets (Sunder, 1992; Copeland and Friedman, 1991 & 1992). However, to the best of our knowledge, we are the first to run such auctions in parallel with the underlying trading task.

Moreover, our auctions are novel in that they only provided loose information about the aggregate signal and did not alter the aggregate signal in any way. Specifically, the auctions concentrateded existing information amongst the winners, thus increasing the "wisdom *in* the crowd" while leaving the "wisdom *of* the crowd" unchanged. The sole purpose of the auctions was to obtain a measure of how much participants trust prices while gradually increasing the competition for profits among auction winners.

The information auctions were run in the first three out of the four trading rounds of each period and took the form of second-price auctions. After 90 seconds of trading and simultaneous bidding, a round ended and the auctioned information was sent to the highest bidder who was charged the *second*-highest bid (up to a maximum of $20). This timing is illustrated in Figure 3.5.

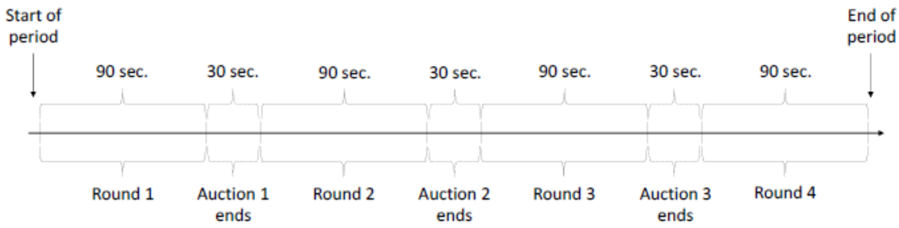

Figure 3.5: Timeline of One Period

Participants received their signals and auction information through a customized web-interface that interacted directly with the online markets system Flex-E-Markets. Screenshots of the interface are depicted in Figure 3.6. In a given period, the personalized display showed a participant's private signal and automatically updated between rounds to display the outcome of the latest auction. Only the auction winner received a "too high," "too low," or "within $1" signal, depending on whether the last trade price was above, below, or within $1 of the fully revealing equilibrium price (see Subfigure (a)). Those who did not win the auction saw an "X" appear instead (see Subfigure (b)). In the final round of a period, only trading and no more auction took place. Once the final round ended, the dividend of the traded stock was revealed through the same interface.

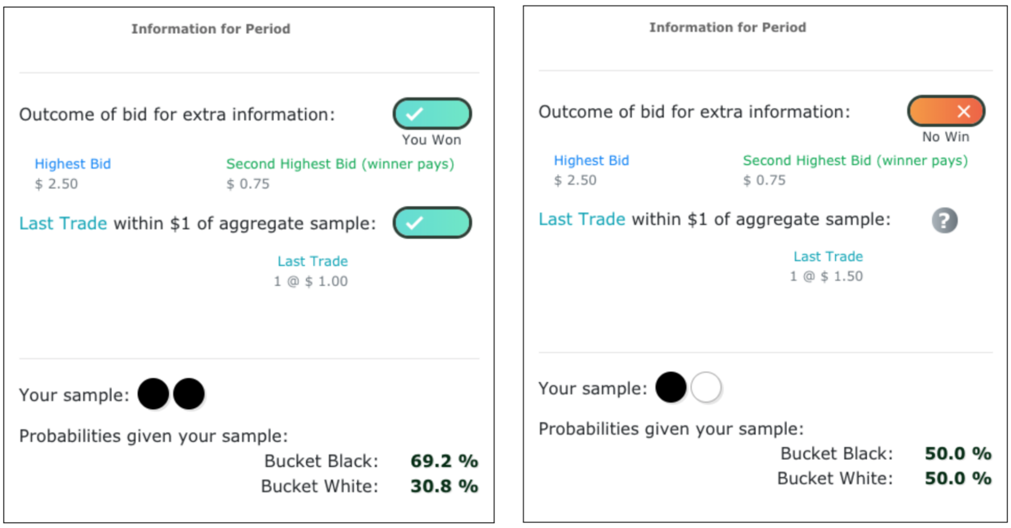

(a) Auction Winner (b) Auction Loser

Figure 3.6: Screenshots of Web-Interface

The private two-bit signals appeared as black and white balls. In the two-urn case, to simplify inference, participants were given the posterior probability, conditional on their private signal, that the underlying urn was Urn Black. We also provided additional conditional probabilities in the instructions (see Appendix B.7) and highlighted the sensitivity of the posterior probabilities to changes in the aggregate signal. We did so to provide participants with intuition about the shape of the aggregate posterior expectation function, i.e., the mapping from aggregate information to FRE equilibrium price levels. Importantly, this eliminates the potential confounding factor caused by (some) participants' inability to perform Bayesian updating.

On the Flex-E-Markets trading screen, participants could follow the order book, submit buy and sell orders, and track their real-time holdings. Figure 3.7 shows a screenshot of the trading screen. Participants were endowed with either eight shares of Stock Black and zero shares of Stock White or vice versa, counter-balanced across participants, as well as $80 in cash. To increase liquidity for those who wished to sell Stock Black but were only endowed with Stock White (the non-traded security), we allowed short selling of up to eight shares of Stock Black. Short selling occurred automatically upon submission of a sell order when a participant's current balance equaled zero. In this case, an additional share of Stock Black got created and sold whenever a shorting transaction was completed. At the end of a period, the respective dividend was then deducted from the short-seller's earnings. The submission of sealed bids occurred via a separate "Information" marketplace.

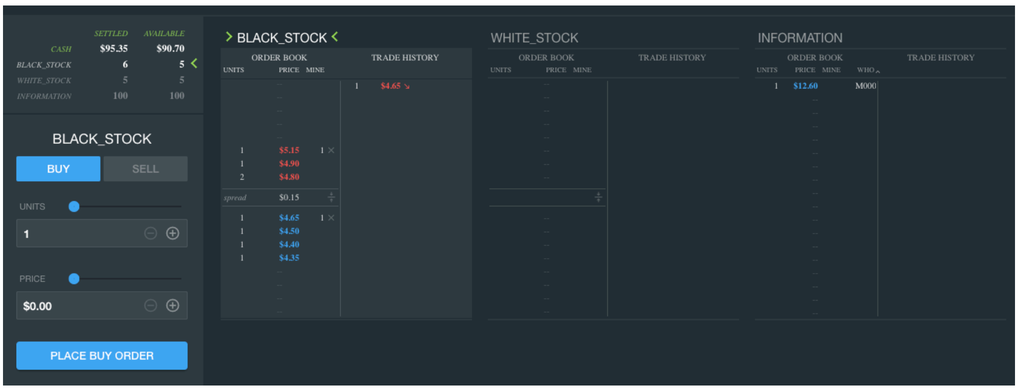

Figure 3.7: Screenshot of Trading Screen

All sessions were run at the Brain Mind and Markets Laboratory at the University of Melbourne, with undergraduate and graduate students as participants.[14] They were paid a fixed participation fee of 20 AUD plus their earnings from trading (converted to AUD at a 3:1 exchange rate). Payments from trading were determined by randomly choosing two periods and paying participants the end-of-period dividends plus cash accumulation, i.e., change in cash over the period. By not paying participants their final cash balance, we increased the salience of non-aggregate risk. On average, participants earned AUD 45, with a minimum of AUD 20 and a maximum of AUD 60.

Each session lasted three hours, with approximately one and a half hours of instructions, hands-on training, plus one practice period. The practice period was designed to ensure participants are comfortable with both the trading protocol and the auction mechanism. To control for order effects and learning, we alternated the treatment order and inverted colors across sessions, everything in completely counterbalanced manner.[15] As mentioned above, we also pseudo-randomized signals and endowments so that expected earnings were fair and there was no correlation between individual signals and outcomes across periods.

3.2.2 Results

Illustration

Figure 3.8 displays trade prices across the 8 periods in the first experimental session (the plots for the remaining sessions are in Appendix B.4).[16] This is a good illustration of our results in general. The first thing to notice is that the 2-urn information structure (in periods 1-4, delineated with black lines) causes much larger price movements per period than does the 1-urn treatment (periods 5-8), consistent with our hypotheses. The price drifts in both settings appear to be roughly monotonic within each round, and rather

[14] The study was approved by the University of Melbourne Ethics Committee (Ethics ID 1852076.2).

[15] In our analysis, we take care of within-participant effects by including fixed effects or by relying on first differences.

[16] Figure B.4 in Appendix B.4 shows session 4, where the information structures are introduced in reverse order (1-urn treatment first). The pattern generated is similar, suggesting that the order of the treatments has no effect.

steep in some sections. Prices sometimes move toward the fully revealing equilibrium (FRE; red level lines), but, in other periods, move toward the Private-Information equilibrium (PE), or even awy from either (period 3, where the two equilibria coincide). We formalize these observations below.

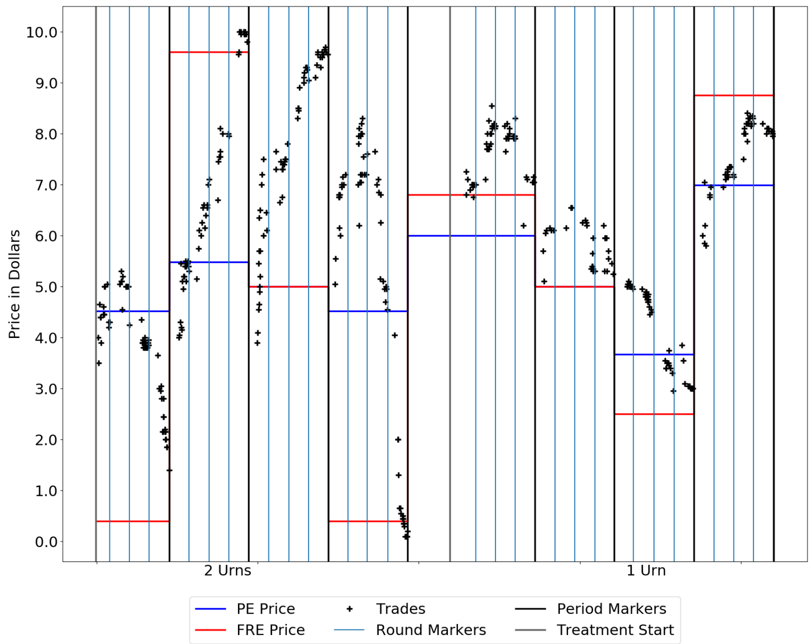

Figure 3.8: Evolution of Trade Prices, Session 1

In both treatments, the end-of-period prices are usually closer to FRE than where they started, which is a sign that markets at least partially aggregate information. If we focus on only the first round of each period (blue vertical lines delineate rounds between information auctions), we observe that trade prices start closer to the Private-Information equilibrium (PE). In subsequent rounds, many different things occur. Focusing on the 2-urn treatment, we observe mirages in the third and fourth periods. The collapse of the mirage by the fourth round of the fourth period can be explained by insider trading: by the fourth round, three participants knew that the price did not reflect aggregate information, and was too high; competition between these “insiders” lead to a collapse in prices towards the FRE. This is an excellent illustration of the effect of insider competition on information revelation: only if at least two insiders are present do prices become more informative (Bossaerts, Frydman and Ledyard, 2014; see also Copeland and Friedman, 1987 & 1991).

The effect of insiders is also apparent in the 1-urn treatment. While overall prices start out closer to FRE, prices at times (e.g., period 5) move away from FRE, but as more participants become aware of the mis-pricing (through the auction), prices revert back to the right level.

An important fact to recall is that the insiders never receive full information. They only receive a direction of mispricing (or confirmation of approximately correct pricing), yet their presence stimulates aggregation, i.e., leads to "amplification" (Copeland and Friedman, 1987; Bossaerts, Frydman and Ledyard, 2014). This can be seen quite clearly in the 2-urn treatment, where, in later rounds, prices begin to correctly aggregate information and suddenly surge toward the efficient price. But the effect is present in the 1-urn treatment as well. We will examine the effects of competition more formally towards the end of the next section.

Formal Analysis: Prices

End-of-Period Prices

Figure 3.9 shows histograms for the absolute differences between end-of-period prices and FRE. The y-axis is the percentage of total observations and the x-axis displays fifty-cent bins of price deviations. Nearly half of all observations are within a dollar (100 cents) of the fully revealing price. This is good news for information aggregation and the Efficient Markets Hypothesis, but it may not be surprising because, after 3 information auctions, there are 3 insiders who know whether the prices were right, and end-of-period pricing with insiders has been found to be informationally efficient (Copeland and Friedman, 1987; Bossaerts, Frydman and Ledyard, 2014).

There is, however, a big difference between the two treatments. The 2-urn treatment generates a long tail (in 3% of the observations, the mis-pricing is more than $8.50 (850 cents), while the 1-urn treatment generates a distribution that is concentrated below \$3.50, with the exception of a small number in the \$4.50-\$5 range. The outliers in the 2-urn case are examples of mirages: the prices look as if they are informative; only, they reveal the wrong information. Therefore, absolute instability and relative instability of the 2-urn case translate into heavy-tailedness in the distribution of deviations of end-of-period prices from FRE.

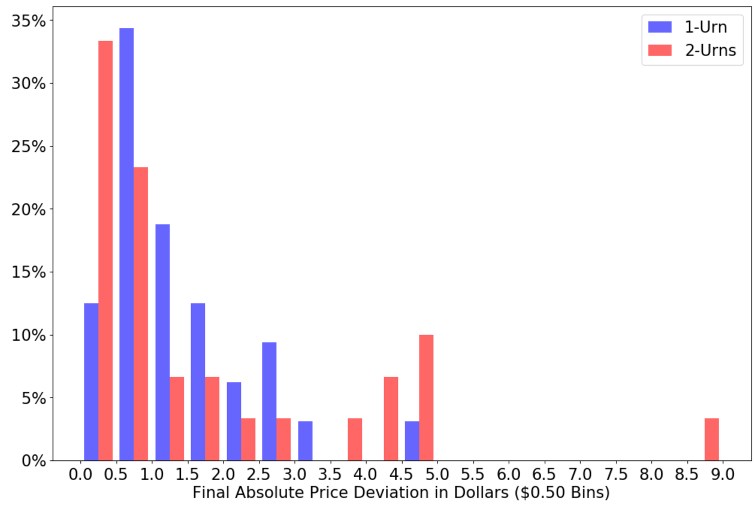

Figure 3.9: Histogram of Absolute Deviations of Final Prices from Fully Revealing Equilibrium

Figure 3.10 displays the same information for the PE equilibrium. Here, the difference between the two treatments is even more outspoken. The majority of the deviations in the 2-urn case are above the maximum obtained in the 1-urn case. This is, in part, a side-effect of the relative instability of the 2-urn case: if FRE obtains, prices in general will be very far from PE. Even when mirages occur, prices are still far from PE. Indeed, with mirages, prices reflect the wrong information, while PE will reveal some (correct) information.

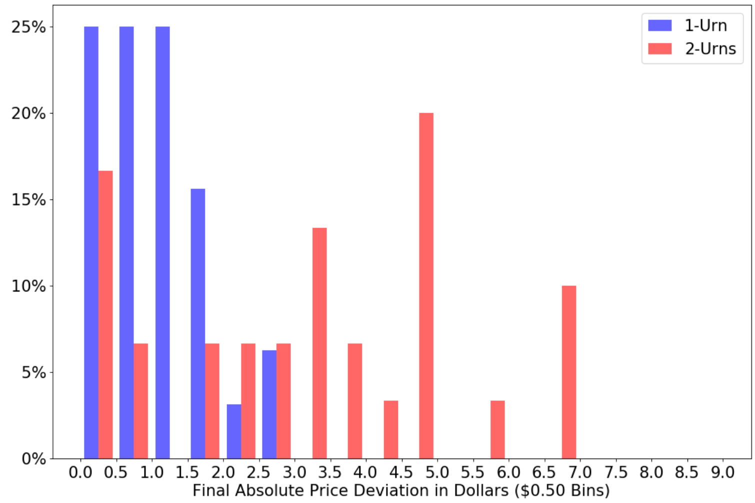

Figure 3.10: Histogram of Absolute Deviations of Final Prices from Private-Information Equilibrium

End-of-Round Prices

In Figure 3.11, we observe a gradual decrease in the average (absolute) deviation of end-of-round prices from FRE. On average, prices are closer to FRE in the (stable) 1-urn case; but in both treatments prices improve equally fast.[17] We attribute this to the competition between insiders that the between-rounds information auctions generate; the auction mechanism is common to the two treatments.

[17]* p<0.1, ** p<0.05, *** p<0.01

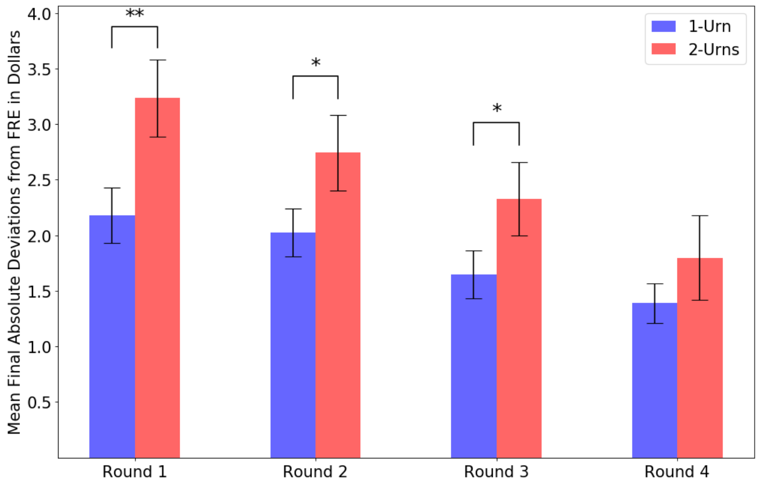

Figure 3.11: Average Absolute Deviations of Final Round Prices From Fully Revealing Equilibrium

Figure 3.12 displays the same information relative to the PE equilibrium. Here, we see the opposite effect between the (unstable) 2-urn treatment and the (stable) 1-urn treatment: prices tend to deviate more from the Private-Information equilibrium in later rounds in the 2-urn case, while they marginally move closer in the 1-urn case.[18] This is, of course, a spurious effect of relative stability: since the 2-urn case is relatively unstable, meaning that the FRE and PE are far apart, and prices generally converge to FRE after the 3rd information auction, prices necessarily have to move away from PE. The 1-urn case, instead, is relatively stable, implying that if prices converge to FRE, they may also move closer to PE as well.

[18] * $p<0.1$, ** $p<0.05$, *** $p<0.01$

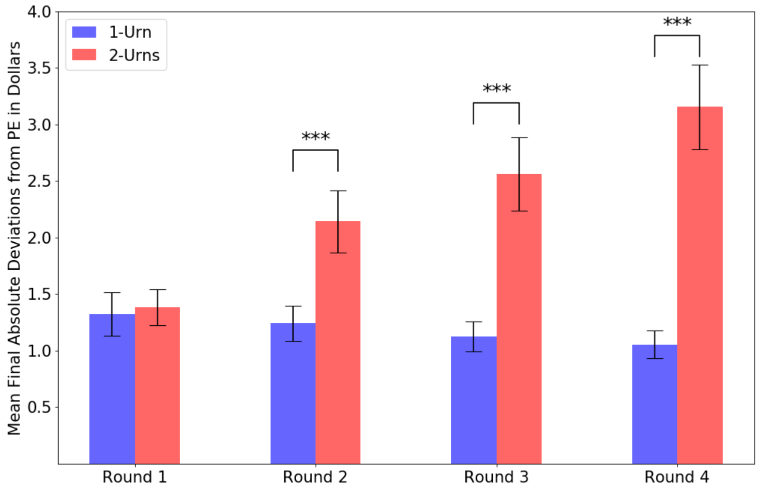

Figure 3.12: Average Absolute Deviations of Final Round Prices From Private-Information Equilibrium

Price Drift

Using trade-by-trade regression analysis, we now study to what extent price movements in a given round reflect attraction towards the FRE and/or PE. We graphically display the resulting 2-standard-error ($\approx$ 95%) confidence intervals around the estimated slope coefficients (confidence intervals are obtained by bootstrapping; see Appendix B.3 for details).

Figure 3.13 displays the results for the 1-urn treatment. Trading rounds are on the x-axis; the y-axis depicts price drift per trade. Positive values mean a drift towards, while negative values mean a drift away from, the corresponding equilibrium. The light blue shaded region depicts the evolution of the confidence interval for price drift towards FRE. The dark blue region depicts the same for PE, orthogonalized for the former. That is, it shows price drift towards towards PE not explained by movement towards FRE. Figure 11 shows how, for the (stable) 1-urn treatment, drift is continuously towards FRE and away from the PE. The drift towards FRE occurs early (drift of zero is outside the confidence interval) and is persistent through rounds 2 to 4; that is also the point at which the drift away from PE is the strongest.

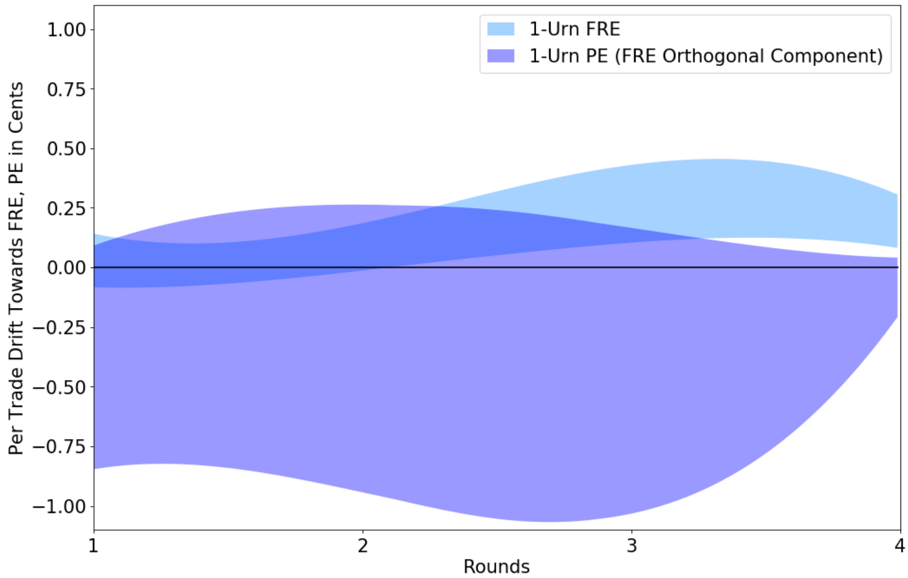

Figure 3.13: Ninety-Five Percent Confidence Intervals of Trade-by-Trade Price Drift Towards (Positive Values) or Away From (Negative Values) FRE (Light Blue) and PE (Dark Blue), Per Round, 1-Urn Treatment

Figure 3.14 displays the results for the 2-urn treatment. There, we see a different pattern. First, the push towards FRE is not as significant, even if the zero line is mostly in the lower part of the confidence intervals. Second, the push away from the orthogonalized PE is less pronounced as well, and is only significant in round 4. Altogether, this presents a picture of a market that has trouble equilibrating, consistent with the visual impression of Figure 3.11 (Periods 1-4).

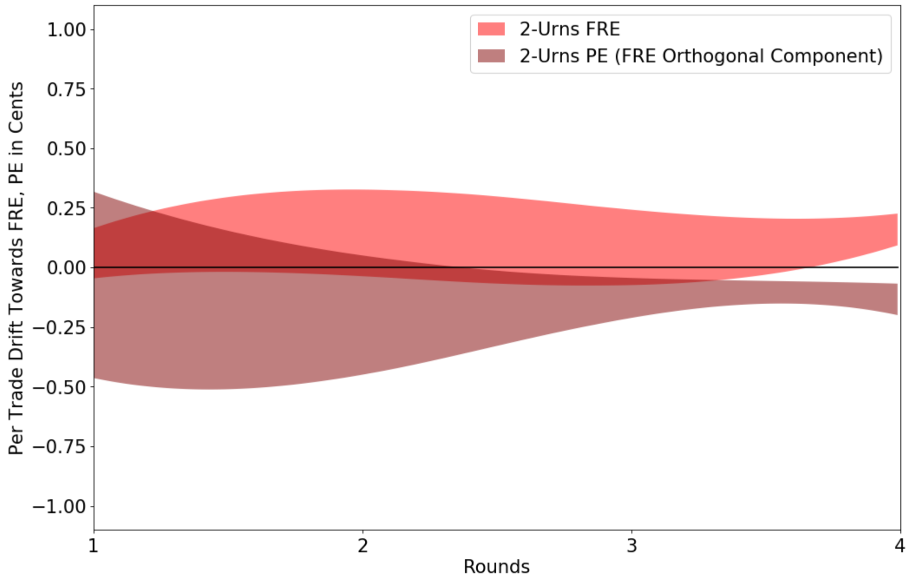

Figure 3.14: Ninety-Five Percent Confidence Intervals of Trade-by-Trade Price Drift Towards (Positive Values) or Away From (Negative Values) FRE (Light Blue) and PE (Dark Blue), Per Round, 2-Urn Treatment

Competition and Prices

In the previous sections, we documented how price efficiency improves in later rounds. One possible cause is the presence of a few individuals who, through the end-of-round auctions, know whether prices correctly reveal the aggregate signal. We now take a closer look at the signals that these insiders received and show that it is the increase in competition between the insiders that drives price efficiency, not just their presence.

In the regressions reported in Table 3.1, the dependent variable is the change, across rounds, in the absolute deviation between the final transaction price and the FRE. By taking the first difference, we remove session and period fixed effects. By focusing on changes in efficiency rather than levels, we remove the endogeneity associated with insiders receiving a “within $1” signal. To prevent losing an observation when taking the first difference, we include a “round 0,” and assign it a price equal to the unconditional expected payoff.

The independent variables of interest are indicators for the urn types, interacted with indicators for levels of insider competition. We deem there to be insider competition if at least two insiders received signals from the auction that point in the same direction. For example, if the last transaction prices in two subsequent rounds were both more than \$1 below the FRE, then the two auction winners would receive a "too low" signal. After two rounds, there would be "two competing insiders."

Errors are clustered by session, treatment and trade round. This is logical, as we expect heteroskedasticity across the different information structures, and across trade rounds as more insiders (auction winners) are present.

Table 3.1 shows that, in the 1-urn case, the presence of two or three competing insiders causes the price deviations (from FRE) to drop significantly ($p. < 0.05$). Similarly, in the 2-urn case, the coefficient on two competing insiders is negative and significantly different from zero at $p < 0.10$, while the coefficient on three competing insiders is negative but significant at only $p = 0.11$. Overall, therefore, competition works in both cases, but the effect is more significant in the 1-urn case.

The regression also shows that, when insiders received "within \$1" signals, price efficiency does not improve significantly. When more than one insider receives this signal type, the insiders are not necessarily competing as they are uncertain about the direction of price improvement; for this reason, we should not be surprised by the lack of significance in those cases either.

Table 3.1: Regression of Change in Distance of Final Trade Prices from FRE onto Measures of Competition Between Auction Winners ("Insiders")

	Change in Distance to FRE	p-value
Constant	9.584	(0.674)
One urn one insider	-34.227	(0.214)
One urn two competing insiders	-63.959**	(0.017)
One urn three competing insiders	-64.028**	(0.024)
Two urns	-21.110	(0.398)
Two urns one insider	-42.141	(0.127)
Two urns two competing insiders	-56.182*	(0.099)
Two urns three competing insiders	-75.533	(0.111)
One within $1 signals	25.210	(0.290)
Two within $1 insider signals	1.698	(0.941)
Three within $1 insider signals	36.434	(0.207)
Observations	251	
Clusters	64	
Adjusted R^2	0.056	

p-values in parentheses, * $p < 0.1$, ** $p < 0.05$.

Formal Analysis: Auction Outcomes

Auction outcomes should reveal to what extent participants trust the information revealed in prices. If clearing prices in the 2-urn case are higher, there is less trust at the margin. If more participants submit bids, then the lack of trust is widespread.

Figure 3.15 displays the total amount bid (sum of prices of all bids). It shows that bidding is more aggressive in the 2-urn treatment (two-sided t-stat = 2.61), suggesting more widespread lack of trust that prices are right.[19]

[19] * p<0.1, ** p<0.5, *** p<0.01

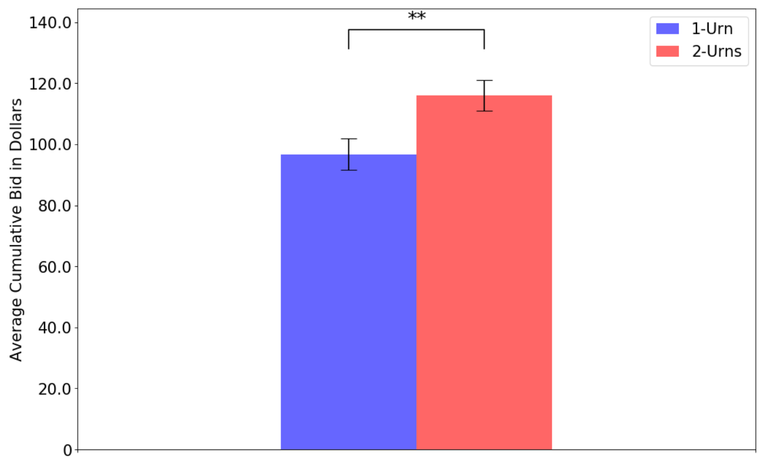

Figure 3.15: Sum of Bid Values in Information Auctions, Average Across Sessions, Stratified by Treatment

Figure 3.16 displays total amount bid by bidding round. Amount bid decreases across rounds, consistent with the increase in pricing accuracy (as documented in Figure 3.11).[20] In Round 1, there is no (significant) difference between the two treatments, but by the second round, bidding in the (relatively stable) 1-urn treatment is significantly less (two-sided t-stat $= -1.750213$). The amount bid in the third round is larger in the 2-urn case but no longer significant at $p = 0.10$.

Overall, the lack of trust in the 2-urn case is evident in each round, and significantly larger than in the 1-urn case. This suggests that the breakdown in trust is a result of the information structure.

[20]* p<0.1, ** p<0.5, *** p<0.01

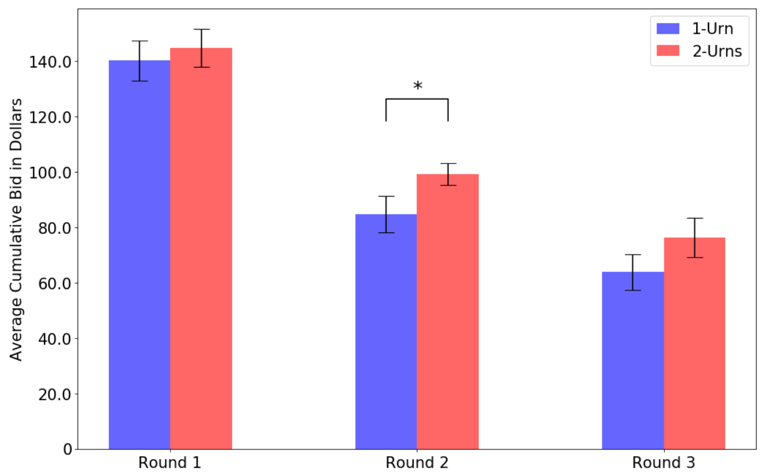

Figure 3.16: Sum of Bid Values in Information Auctions, Average Across Sessions, Per Period, Stratified by Treatment

Formal Analysis: Earnings and Holdings

Earnings

We know that prices in the marketplace behave very differently, with the 2-urn case sometimes leading to mirages. The question arises whether the higher auction bidding in the 2-urn case is compensated for in terms of trading profits. Figure 3.17 shows that, when comparing the two treatments, there is hardly any difference in average earnings after accounting for bid payments.[21] There is only a drop in net earnings for double winners (those who won two auctions in a single period). Winning two auctions is not an irrational strategy as it may be profitable to acquire additional information after information from the first win has been impounded into prices. However, the results show that those who won two auctions significantly overpay.[22] The extreme mistrust in prices revealed through aggressive bidding in the auction is costly, irrespective of information treatment.

The figure also shows that participants who do not win in the information auction

[21] * p<0.1, ** p<0.5, *** p<0.01

[22] The difference between the double winners and the non-winners is significant at the 1% level for the 1-urn treatment, and at the 5% level for the 2-urn treatment

earn no less than single auction winners. This is consistent with earlier evidence of the inability of insiders to make extraordinary profits when competing with each other (Bossaerts, Frydman and Ledyard, 2014).

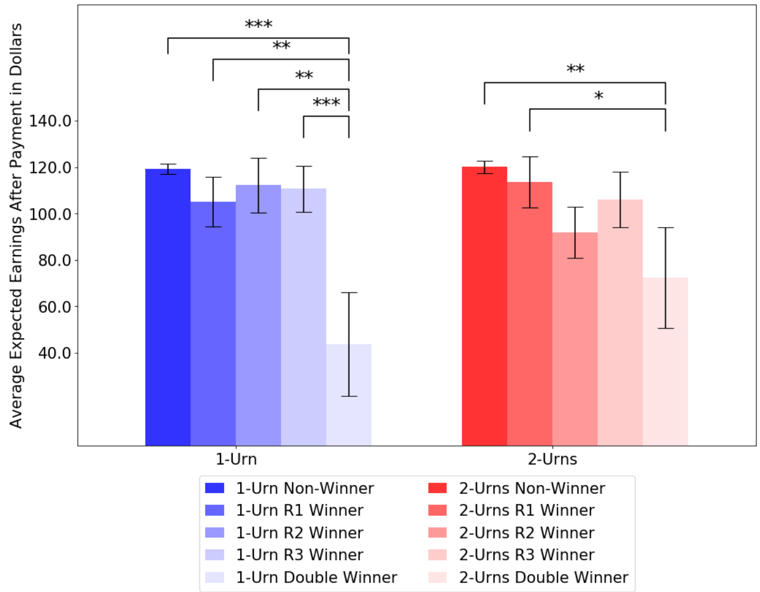

Figure 3.17: Average Expected Portfolio Earnings Net of Auction Costs, Stratified by Winning Status and Treatment

Appendix B.5 provides more details on the earnings within each trader group and additional post hoc testing.

Holdings

Let us now consider final holdings, measured as the *net* Black stock position, which is the number of Black stocks less the number of White stocks in a subject's portfolio. If the net position is zero, the participant is fully hedged. Figure 3.18 shows that the modal participant who does not win the auction (dark blue represents 1-urn treatment; red represents 2-urn treatment) is close to fully hedged, suggesting that the final price reflects conditional expectations. This is quite remarkable, since it implies that the modal participant agrees on the information that is out there, even if they all start with disparate information. Agreement obtains equally frequently in the two treatments, even if final prices in the 1-urn case better reflect true aggregate information.

Among auction winners, Figure 3.18 shows that the modal participant in the 1-urn treatment (lighter blue stems) also agrees with the information revealed in the final price. This is perhaps not surprising since final prices tend to be right (see Figure 3.8). In contrast, in the 2-urn case, the percentage of participants who hold a final balanced position, and hence, agree with the price, is substantially lower (see lighter red histograms). If we formally compare the holdings distributions of winners against non-winners in the 2-urn case using a Kolmogorov-Smirnov test, the distributions are significantly different with a p-value less than 5% (See Appendix B.6 for additional tests).[23] This is consistent with the finding that end-of-period prices deviate more from FRE in the 2-urn treatment (see Figure 3.8).

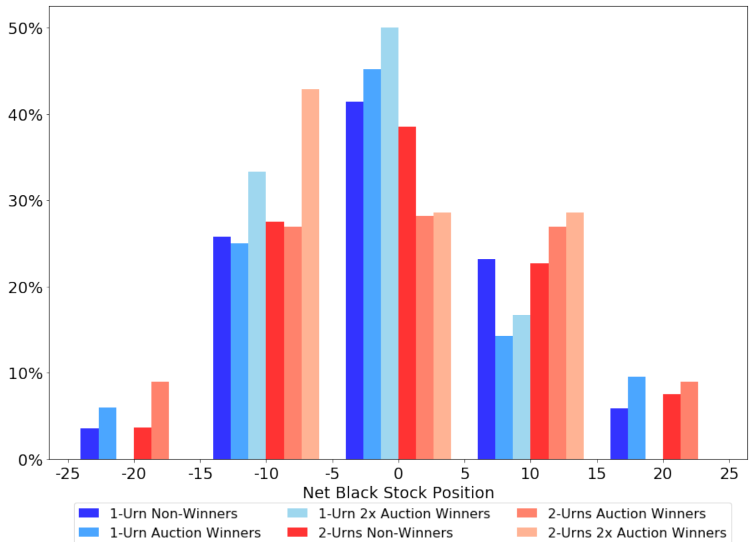

Figure 3.18: Period End Net Black Stock Holdings Histogram by Trader Type

Auctions, Earnings, and Holdings Summary

To recap, total auction bidding was predicted to be higher in the 2-urn setting because of the trust issues associated with the absolute and relative instability of the information

[23]Histogram shapes are sensitive to bin selection, so only a formal test using the empirical distribution function, such as the Kolmogorov-Smirnov test, is conclusive.

structure. The null hypothesis that bidding is the same in both settings is rejected at the 5% level. Our earnings results show that information obtained by single auction winners is reflected in market prices, such that we cannot reject the null hypothesis that profits are the same as that of non-winners. This is consistent with past experimental results on information revelation. Finally, holdings were predicted to be less balanced in the 2-urn case, as subjects rely more on their private information because of mistrust in prices.

3.3 Conclusion

This essay provides an explanation of why markets appear to aggregate information well when payoffs equal the average of the signals of agents in the economy, while mirages often emerge when payoffs are binary and depend on whether the majority of participants have evidence in one way or another. The former structure is the canonical one in theoretical work on asset pricing; the latter has been the information structure traditionally chosen in experimental work.

We introduced two notions of stability in order to explain the differences. We tested the theoretical predictions in a tightly controlled experiment and find evidence in their favor.

Future work should aim at testing the case of absolute stability and relative instability or vice versa. The cases we considered here either have both absolute and relative instability, or are stable in the two senses. The investigation of mixed cases should shed light on the eventual impact of garbled information on information aggregation (absolute stability), against the impact of trust in information aggregation (relative stability).[24]

[24] Related to relative instability, Hellwig (1981) points out that trust issues could also emerge when the number of agents increases. As the price aggregates more private information, it becomes a sufficient statistic, meaning that there is no longer a need to (also) condition on own, private signals; but then private signals no longer are impounded in the price, and the price loses its informativeness. In such an economy, a rational agent will question whether she can trust the price as relaying valuable information.

Our findings underscore the importance of understanding the underlying information structure to predict quality of information aggregation in markets. Ultimately, this is of importance for evaluating the empirical record of the Efficient Markets Hypothesis in the field.

Our results are also important for the design of prediction markets. It is common to choose a winner-take-all asset payoff structure. Our findings suggest that resulting prices cannot be relied upon as good forecasts of the future. In Dreber e.a. (2015), for instance, the market correctly predicted the outcome of scientific replications in 29 out of 41 cases, which is barely better than a rule that always predicts unsuccessful replication (which would have won 25 out of 41 cases). Insiders need to be relied on in order to improve forecasts (Bossaerts, F. e.a., 2019). Far smaller biases should emerge when prediction markets aim at directly estimating the average signal. Consistent with this prediction, Botvinik e.a. (2019) reports superior forecasting when predicting outcomes from scientific replications.

There is an emerging literature on design of mechanisms to extract the wisdom of the crowd. Baillon (2017), for instance, considers a case that maps into our 2-urn setting and constructs a mechanism for which the Nash equilibrium is such that all participants truthfully reveal their signal. Interestingly, the mechanism does not ask participants to reveal what they perceive to be the correct urn. Instead, they are to reveal their sample (the color of the balls they received). With this information, the true composition of the urn can be determined, and hence, the 2-urn problem can be resolved. Effectively, the mechanism turns a 2-urn problem into a 1-urn problem.[25] Our finding that the information structure in the 1-urn case is more robust should enhance revelation mechanisms that convert the 2-urn case into a 1-urn case.

[25]Baillon did not run into belief stability issues of the type we are interested in, because the participants were never told what others had revealed. That is, the participants did not have access to system-wide information such as a market price. There was a "market," but participants could only see personal trades.

Chapter 4

Modelling Asset Prices Under Heterogeneous Beliefs

4.1 Motivation

Consider an economy in which, among others, ice cream production decisions have to be made one period before consumption can take place. Consumption depends on the weather. If it is sunny, consumption of ice cream is high; otherwise consumption is low. Agents make savings and investment decisions based on their forecasts of future prices, including the price of ice-cream. Assume, as in Radner (1972, 1979) that agents have *perfect foresight* ("Rational Expectations" RE). They know that the ice-cream is going to be more expensive when the sun shines, but estimate the price not to be exorbitant. This could be the consequence of most agents' believing that the sun is likely to shine, and hence, expecting plenty of ice-cream production, which would then satisfy the high demand upon sunshine. Now take an agent who disagrees, and believes that sunshine is unlikely. This agent predicts low ice-cream production, and hence, should expect prices to be high if the sun happens to shine, since high demand meets low production. As such, her beliefs lead to price forecasts which are at odds with perfect foresight.

The question we ask here: is the above problem generic? That is: does there exist an inherent contradiction between perfect foresight and belief heterogeneity? We show that the problem is generic. However, a sufficient condition for perfect foresight about prices to be consistent with differences in beliefs is that higher-order beliefs (beliefs about other agents' beliefs) are correct.

We first show this in the context of forecasting future prices. The context is relevant for production economies, as in the model of Muth (1961), and for inter-temporal asset pricing, as in the models of Radner (1972) and Lucas (1978). The latter provides the foundation of modern dynamic asset pricing theory and dynamic stochastic general equilibrium (DSGE) models of the macro-economy. Next, we show that the same result is true in the context of asymmetric information, where market participants infer from prices what the true state could be, given perfect foresight of how prices relate to states. Models in this genre include those of Green (1975), Lucas (1972), Radner (1979) and Grossman (1976).

We study the case where priors are different, but there is no common knowledge of how priors differ. Traditional applications of the Radner's perfect foresight equilibrium (the "Rational Expectations Equilibrium" REE) follow Harsanyi's doctrine (Harsanyi, 1968) and assume a common prior. Aumann (1976) has shown that, if this is common knowledge, then the posterior beliefs must not differ except for differences in information, and hence, there cannot be belief heterogeneity. Asset pricing examples include Grossman (1976), Grossman and Stigliz (1980), Diamond and Verecchia (1981), Hellwig (1980) and Admati (1985).

We consider more recent versions of the REE, whereby agents start with irreconcilable differences in priors (differences that cannot be traced back to differences in information). Examples include Scheinkman and Xiong (2003), Basak (2000,2005), Gallmeyer and Hollifield (2008) and Osambela (2015). That is, agents "agree to disagree" (Aumann, 1976). Because they appeal to REE, the authors need not specify whether the differences in priors are common knowledge.[1] Indeed, agents take the mapping from states to prices as given, and plan their investment and consumption based on their beliefs about those states. As a matter of fact, agents need not even know that the mapping from states to prices is what it is because of differences in beliefs.

Here, we imagine that agents are aware of this (as were the authors), and, absent common knowledge of those beliefs, use their own beliefs about the beliefs of others, i.e., use their higher-order beliefs, check the veracity of the mapping from states to prices that, in accordance with REE, they are supposed to know.

Applied to the work of Basak, Gallmeyer, Hollifield and Osambela, for instance, this would mean the following. The authors specify beliefs about prices in terms of diffusion equations (see, e.g., Equations (10) and (11) in Basak, 2005); these diffusion equations are correctly specified, except for the drift coefficients, which are subjective. If agents know the beliefs of others, they could correctly derive these diffusion equations.[2]

[1] Scheinkman and Xiong differ in that agents are assumed to *know* differences in beliefs. Agent *derive the equilibrium pricing restrictions*, very much in the way we shall do. Hence, Scheinkman and Xiong impose far more restrictions on beliefs than is standard in a REE (Radner, 1972).

[2] This is what happens in Scheinkman and Xiong (2003).

If agents do not know, then the issue arises as to whether they would have come up with these equations. In particular, under what type of higher-order beliefs (beliefs about the beliefs of others) would they derive the right diffusion term? This is what concerns us here. We will study the issue in a more transparent way, though, since our setups will be far simpler.

In their model of differences of opinion, Banerjee, Kaniel and Kremer (2009) do not appeal to REE. Instead, they ask agents to come up with their own forecasts of prices, based on their own beliefs of other agents' beliefs. Effectively, we contrast Banerjee e.a.'s modeling approach with the traditional REE modeling of asset prices under belief heterogeneity. Our position is that agents in the REE should be capable to perform the forecasting exercise that Banerjee e.a.'s agents do. If they do so, will the predictions be consistent? More precisely: under what condition would the two approaches to forecasting future prices produce the same estimates?

On the other end of the spectrum is Guesnerie (1992), who stays with the traditional REE, but investigates ways to come up with the correct mapping from states to prices. In the spirit of the theory of games with incomplete information, Guesnerie's agents, knowing the constraints others face, eliminate dominated responses based on mental forecasts (including forecasts of others' forecasts). Guesnerie studies under what conditions perfect foresight, and hence, REE, can emerge through this type of "eduction." Instead, we assume that agents somehow know the mapping from states to prices, and study to what extent this knowledge conflicts with the beliefs they are holding about fundamentals in the economy, and about other agents' beliefs.

There is another approach in the literature on asset pricing under heterogeneous beliefs. There, differences in posterior beliefs emerge not because priors differ, but because the "model" (the likelihood function) with which data are evaluated differs across agents. See, e.g., Anderson and Sonnenschein (1985), Harris and Raviv (1993), Kandel and Pearson (2005), Ferguson (2015), Schraeder (2016). How clashes between prices and agents' (necessarily differing) models work out is studied, theoretically and experimentally, in Asparouhova e.a. (2015). Such clashes are bound to occur, since two rational agents cannot possibly agree to disagree on the choice of the model with which to interpret the data, since one, if not both, must be wrong. Effectively, the heterogenous-modeling approach is a special case of the one we consider here. Differences in modeling can be built into differences in priors.

Our use and definition of higher order beliefs goes back to Keynes' (1936) beauty contest. First-order beliefs are the priors about the fundamentals in the economy. Second-order beliefs are the beliefs about other agents' first-order beliefs. Third-order beliefs are the beliefs about other agents' second-order beliefs. Etc. Higher-order beliefs feature explicitly in the models of, e.g., Townsend (1983), Biais and Bossaerts (1998), Allen, Morris and Shin (2006), Banerjee, Kaniel and Kremer (2009), and Makarov and Rytchkov (2012).

"Perfect foresight" has its analog in game theory. In asset pricing, perfect foresight is the ability to correctly predict the right (read: equilibrium) prices in every future contingency. In game theory, the analog is subgame perfection: players are supposed to know how the game continues in every future node; continuation is restricted to Nash equilibrium play. See e.g., Fudenberg and Levine (1983). Unlike game theorists, asset pricing theorists have not yet studied which beliefs could possibly be consistent with subgame perfection, or its extension to games of incomplete information, sequential rationality. See, e.g., Fudenberg and Tirole (1993), where the resulting "perfect foresight" equilibrium is referred to as Perfect Bayesian Equilibrium. This paper can therefore be viewed as addressing an issue in asset pricing that game theorists have long studied in games. Unlike game theorists, however, we allow beliefs to differ across agents in ways that cannot be traced to a common prior. That is, we allow agents to "agree to disagree" (Aumann, 1976).

The remainder of the paper is organized as follows. In the next two sections we present two examples in which agents disagree and are equipped with perfect foresight information. The first example is in a production framework, and the second example is in an asymmetric information framework. Section 4 discusses the plausibility of the sufficient condition for perfect foresight not to clash with individual agents' beliefs under belief heterogeneity.

4.2 Example 1: Forecasting Future Prices

In order to make appropriate production decisions, agents need to forecast future market prices. The forecasts are assumed to be rational. Muth, and later, Lucas, interpreted this to mean that forecasts had to be consistent with equilibrium economic analysis. In other words the mapping from states to prices had to be suggested by economic theory.

Both Muth (implicitly) and Lucas (explicitly) also required agents' beliefs about the states to be correct. This example will illustrate that this assumption can be dispensed with. The main issue is: can disagreement be allowed for? The example illustrates that this cannot be done without the risk of introducing inconsistencies. The example is deliberately kept simple, in order not to distract from the main point.

Consider a perfectly competitive market for a good that is produced by two types of producers: $(i = 1, 2)$. Total demand, Q, is linear in the price, p:

$$Q = ap + u, \qquad (a < 0) \tag{4.1}$$

where u, is a random shock.

Both types of producers face the same cost function C:

$$C(q_i) = \frac{1}{2}\alpha q_i^2 + \beta q_i + \gamma, \tag{4.2}$$

where q_i is producer i's output $(i = 1, 2)$ and u is a random variable defined on the probability space $(\Omega, \mathcal{F})$. There is a true probability measure P^* that determines actual outcomes. There are also two additional probability measures on $(\Omega, \mathcal{F})$, $P^i(i = 1, 2)$, which may or may not differ from each other and/or P^*. P^i determines the belief of producers of type i about u. P^i is not common knowledge. This means that beliefs are random variables as well. In fact, beliefs about beliefs are also random variables. As will be beliefs about beliefs about beliefs etc... We will assume that the entire beliefs hierarchy exists on $(\Omega, \mathcal{F})$. To be accurate, we would have to be more specific. For instance, for agent i's belief to be *consistent*, we must make sure the P^i assigns unit weight to the value of her own belief that she actually holds. The ensuing notation and complexity would distract from the main point of this note, however. The reader who is interested in a complete specification of beliefs hierarchies is referred to Mertens and Zamir (1985).

Given, P^i is not common knowledge, instead, let us assume that it is common knowledge that:

$$E^i[E^j[u]] = E^i[u], \tag{4.3}$$

where E^i and E^j denote expectations with respect to P^i and P^j, respectively. (4.3) is an assumption about higher-order beliefs. $E^i[u]$ are determined by i's *first-order beliefs* ($i = 1, 2$). Let us use the shorthand notation:

$$\lambda^i = E^i[u], \tag{4.4}$$

$E^i[E^j[u]]$ is determined by i's *second-order beliefs* about j's first-order belief. Note that, (4.3) is the *average opinion rule* highlighted in Biais and Bossaerts (1997) and inspired by Muth (1961). Here, i thinks that the average j has the same belief that she has. Biais and Bossaerts (1997) demonstrate how rules such as the average opinion rule dramatically simplify the analysis of games where beliefs may differ but are not common knowledge ("beauty contests"). The addition of the average opinion rule in this case eliminates the need to fully specify the entire beliefs hierarchy.

The fact that P^i is not common knowledge, and that (4.3) is common knowledge, will be irrelevant for the derivation of REE. The common knowledge assumptions play a role only when checking the REE for internal consistency: if (4.3) is not common knowledge, the equilibrium pricing function of the REE (perfect foresight) may contradict one's conjecture about these prices as derived from equilibrium analysis under one's own beliefs.

Price taking producers maximize expected profits. Hence:

$$q_i = \frac{E^i[p] - \beta}{\alpha}, \tag{4.5}$$

Notice producers make production decisions before observing equilibrium prices. Consequently, production decisions are based on expectations about future prices.

In equilibrium, demand (Q in (4.1)) must equal supply. Assuming an equal number of producers of each type, we can write this equilibrium condition as follows:

$$Q = q_1 + q_2. \tag{4.6}$$

After substitution of the production decision, (4.5) , one obtains:

$$Q = \frac{E^1[p] - \beta}{\alpha} + \frac{E^2[p] - \beta}{\alpha}. \tag{4.7}$$

How do producers form the expectations $E^i[p]$ in (4.5)? In a REE, they are formed on the basis of an announced mapping from states to prices. In the present example, this mapping would be linear in λ^1, λ^2, and u:

$$p(\lambda^1, \lambda^2, u) = \delta_0 + \delta_1\lambda^1 + \delta_2\lambda^2 + \delta_3 u. \tag{4.8}$$

When each producer takes expectations given the mapping:

$$\begin{cases} E^1[p|p(\cdot)] = \delta_0 + (\delta_1 + \delta_2 + \delta_3)\lambda^1, \\ E^2[p|p(\cdot)] = \delta_0 + (\delta_1 + \delta_2 + \delta_3)\lambda^2, \end{cases} \tag{4.9}$$

where the average opinion rule (4.3) is used; but not the fact that it is common knowledge.

The coefficients of the price function in (4.8) reflect the true relation between states (λ^1, λ^2, u) and prices (p) when expectations are based on (4.9). In other words, they are obtained as the values that solve the equilibrium equation (4.5) for all possible values of λ^1, λ^2, u.

Substituting (4.9) for $E^i[p]$ in (4.7), and solving for p:

$$p(\lambda^1, \lambda^2, u) = \frac{2(\delta_0 - \beta)}{\alpha a} + \frac{(\delta_1 + \delta_2 + \delta_3)}{\alpha a}\lambda^1 + \frac{(\delta_1 + \delta_2 + \delta_3)}{\alpha a}\lambda^2 - \frac{1}{a}u. \tag{4.10}$$

Hence, $(\delta_0, \delta_1, \delta_2, \delta_3)$ are the solution to:

$$\begin{cases} \delta_0 = \frac{2}{\alpha a}(\delta_0 - \beta), \\ \delta_1 = \frac{1}{\alpha a}(\delta_1 + \delta_2 + \delta_3), \\ \delta_2 = \frac{1}{\alpha a}(\delta_1 + \delta_2 + \delta_3), \\ \delta_3 = \frac{1}{a}. \end{cases} \tag{4.11}$$

For the consistency check, we only need the result that:

$$\delta_0 = -\frac{2\beta}{\alpha a - 2}, \tag{4.12}$$

$$\delta_2(\alpha a) = -\frac{\alpha}{\alpha a - 2}. \tag{4.13}$$

Now consider a producer of the first type ($i = 1$) who knows that prices are set in equilibrium. Instead of computing expectations on the basis of an announced price function she could apply economic analysis herself and compute expectations directly from the equilibrium condition $Q = q_1 + q_2$. In this mental exercise, she assumes that producers of type 2 continue to use the REE price function (4.8) to calculate expectations. The result is:

$$E^1[p] = \frac{-2\beta}{\alpha a - 1} + \frac{\delta_0}{\alpha a - 1} + \frac{\delta_1}{\alpha a - 1} E^1[E^2[\lambda]] + \frac{\delta_2 + \delta_3}{\alpha a - 1} E^1[\lambda^2] - \frac{\alpha}{\alpha a - 1}\lambda^1. \quad (4.14)$$

Notice that $p(\cdot)$ is left out of the conditioning information. We did this to reflect that our producer does not use $p(\cdot)$ for her own inference. Yet, she still uses $p(\cdot)$ to determine the inference of producers of type 2.

We have used the average opinion rule (4.3) in the computation of the REE. We use it again, to conclude:

$$E^1[\lambda^2] = \lambda^1. \quad (4.15)$$

We need more, however, to determine $E^1[E^2[\lambda^1]]$, by assuming *common knowledge* of the average opinion rule, we conclude:

$$E^1[E^2[\lambda^1]] = E^1[\lambda^2] = \lambda^1. \quad (4.16)$$

Hence, (4.14) becomes:

$$E^1[P] = \frac{\delta_0 - 2\beta}{\alpha a - 1} + \frac{\delta_1 + \delta_2 + \delta_3 - \alpha}{\alpha a - 1}\lambda_1. \quad (4.17)$$

This coincides with the expectation of REE, $E^1[p|p(\cdot)]$ in (4.9). Using (4.12) and (4.13), it is straight forward (but tedious) to establish that:

$$\frac{\delta_0 - 2\beta}{\alpha a - 1} = \delta_0, \quad (4.18)$$

$$\frac{\delta_1 + \delta_2 + \delta_3 - \alpha}{\alpha a - 1} = \delta_1 + \delta_2 + \delta_3 \quad (= \delta_2 \alpha a). \quad (4.19)$$

Hence, producers of type $i = 1$ see the announced REE price function confirmed in their own calculations.

Still, this confirmation depends crucially on the assumption that the average opinion rule (4.3) is common knowledge. In other words, it obtains only because producers' third order beliefs (beliefs about second-order beliefs of others) are common and correct.

To see this, let us assume (4.3), but delete the requirement that it is common knowledge. Instead, producers of type $i = 1$ believe:

$$E^1[E^2[\lambda^1]] = E^1[\psi\lambda^2] = \bar{\psi}\lambda^1, \tag{4.20}$$

where $\bar{\psi} \in (0, \infty)$. The REE is unaffected by this amendment of beliefs; however, producers of type $i = 1$, can repeat their computation of $E^1[p]$ from the equilibrium condition, $Q = q_1 + q_2$, and infer:

$$E^1[P] = \frac{\delta_0 - 2\beta}{\alpha a - 1} + \frac{\delta_1\bar{\psi} + \delta_2 + \delta_3 - \alpha}{\alpha a - 1}\lambda_1. \tag{4.21}$$

As before,

$$\frac{\delta_0 - 2\beta}{\alpha a - 1} = \delta_0, \tag{4.22}$$

but,

$$\frac{\delta_1\bar{\psi} + \delta_2 + \delta_3 - \alpha}{\alpha a - 1} = \frac{-\alpha\bar{\psi}}{a\alpha - 2} \neq \frac{-\alpha}{a\alpha - 2} = \delta_1 + \delta_2 + \delta_3 \quad (= \delta_2\alpha a), \tag{4.23}$$

unless $\bar{\psi} = 1$. Consequently,

$$E^1[p|p(\cdot)] \neq E^1[p], \tag{4.24}$$

unless $\bar{\psi} = 1$: the inference from the REE price function differs from what can be recovered by inverting the equilibrium condition.

Producers' third-order beliefs are correct only if $\bar{\psi} = 1$. It would hold if (4.3) is common knowledge. In general, however, $\bar{\psi} \neq 1$: third-order beliefs are biased. First-order beliefs were allowed to differ ($\lambda^1 \neq \lambda^2$), so they will be biased in general

($\lambda^1 \neq E^*[u]$; $\lambda^2 \neq E^*[u]; E^*$ denotes the (conditional) expectation with respect to P^*). It would therefore be questionable to assume that third-order beliefs are correct. So, we must allow $\bar{\psi}$ to be different from 1.

But this causes the producers of type $i = 1$ to question the REE price function: it generates inference that conflicts with what they themselves invert from the equilibrium conditions. Either they have to reject their own beliefs (in fact, only their higher-order beliefs), or they reject the REE price function. The latter leads to a breakdown in the REE.

There is information in the price function that is announced in the REE: it reveals where producers' beliefs are incorrect, because it is based on the true equilibrium relation between states and prices. If $\bar{\psi} \neq 1$, the REE price function is impossible according to producers' beliefs (has zero likelihood). As true Bayesians, producers either will reject the price function as questionable, or they must scrap their beliefs and start from scratch.

It is easy to imagine variations on the above analysis. We assumed that producers of type $i = 1$ assumed that only producers of type $i = 2$ used the REE price function when checking its plausibility (see (4.14))). A producer of type $i = 1$ may assume that other producers of her own class also continue to use the REE price function. Assuming that she is "small," her production decision does not affect the total production of producers of her type, and, (4.14) changes to:

$$E^1[p] = -\frac{2\beta}{\alpha a} + \frac{2\delta_0}{\alpha a} + (\frac{\delta_1 + \delta_2 + \delta_3}{\alpha a} - \frac{1}{\alpha})\lambda_1 + \frac{\delta_1}{\alpha a} E^1[E^2[\lambda^1]] + \frac{\delta_2 + \delta_3}{\alpha a} E^1[\lambda^2]. \quad (4.25)$$

Using,

$$E^1[E^2[\lambda^1]] = E^1[\psi \lambda^2] = \bar{\psi}\lambda^1, \quad (4.26)$$

$$E^1[p] = \frac{2(\delta_0 - \beta)}{\alpha a} + \frac{\delta_1(1 + \bar{\psi}) + 2\delta_2 + 2\delta_3 - \alpha}{\alpha a}\lambda^1. \quad (4.27)$$

The intercept is directly from the set of REE price function parameter equations:

$$\frac{2(\delta_0 - \beta)}{\alpha a} = \delta_0. \quad (4.28)$$

The slope:

$$
\begin{aligned}
&\frac{\delta_1(1+\bar{\psi})+2\delta_2+2\delta_3-\alpha}{\alpha a} \\
&= \frac{2(\delta_1+\delta_2+\delta_3)-\alpha+(\bar{\psi}-1)\delta_1}{\alpha a} \\
&= \frac{(\delta_1+\delta_2+\delta_3)}{\alpha a}+\frac{(\delta_1+\delta_2+\delta_3)}{\alpha a}-\frac{\alpha}{\alpha a}+(\bar{\psi}-1)\delta_1 \\
&= \delta_1+\delta_2+\delta_3+(\bar{\psi}-1)\delta_1. \qquad (4.29)
\end{aligned}
$$

equals that of the REE price function $(\delta_1+\delta_2+\delta_3)$ only if $\bar{\psi}=1$.

4.3 Example 2: Inferring Information from Prices

This example is based on the model of Grossman and Stiglitz (1980).

Consider the market for a forward contract that pays a random $X-p$ to the long position, where X is normally distributed, and p is the forward price. There are 3 types of traders:
i) hedgers that take a random exogenous position size $S \sim N(0,1)$,
ii) speculators type H, that receive signal $H \sim N(0,1)$,
iii) speculators type N, that receive signal $N \sim N(0,1)$
H, N are mutually independent of S.

$$X = H + N. \qquad (4.30)$$

As in example 1, all random variables are on $(\Omega,\mathcal{F})$, with probability measure $P^*, \quad P^i, \quad i=N,H.$

Speculator i's preferences are:

$$U^i(D^i) = E^i[D^i(X - P)|\cdot] - \frac{1}{2}(D^i)^2, \tag{4.31}$$

where D^i denotes the size of the position established by speculator i. We have not filled out the conditioning information in E^i yet; this will follow shortly.

Let, $\lambda^H = E^H[S], \quad \lambda^N = E^N[S]$. These first order beliefs are not common knowledge. Instead, $\lambda^H \sim N(\lambda^N, 1)$ for $i = N$, $\lambda^N \sim N(\lambda^H, 1)$ for $i = H$, so the average opinion rule still holds,

$$\begin{cases} \lambda^H = E^H[S], \\ \lambda^N = E^N[S]. \end{cases} \tag{4.32}$$

Let us skip the case of common knowledge of (4.32), and directly proceed to:

$$E^N[E^H[\lambda^N]] = E^N[\psi\lambda^N] = \bar{\psi}\lambda^N, \tag{4.33}$$

where $\bar{\psi} \in (0, \infty)$. We will analyze the REE price function from the point of view of speculators of type N. In that case, we do not need to specify the higher-order beliefs of speculators of type H. The results hold for N as well, *mutatis mutandis*.

The utility function in (4.31) is one of risk neutrality with a quadratic position cost. The optimal position for speculator i is:

$$D^i = E^i[X - p|\cdot]. \tag{4.34}$$

In a REE, agents form conditional expectations on the basis of their own signals, as well as the information that is reflected in the price, extracted using the announced pricing function, $p(\cdot)$:

$$\begin{cases} D^H = E^H[X - p|H, p, p(\cdot)], \\ D^N = E^N[X - p|N, p, p(\cdot)]. \end{cases} \tag{4.35}$$

The following mapping from states to prices is a REE:

$$p = b(H + N) + c(\lambda^H + \lambda^N) + dS, \tag{4.36}$$

where (b,c,d) are the solution to:

$$\begin{cases} 2b = \frac{b^2-(b^2+c^2+d^2)}{b-(b^2+c^2+d^2)}, \\ 2c = \frac{b(2c+d)}{b-(b^2+c^2+d^2)}, \\ 2d = \frac{b^2+c^2+d^2}{b-(b^2+c^2+d^2)}. \end{cases} \tag{4.37}$$

In order to derive the equilibrium, we must project N and H onto p, conditional on $p(\cdot)$, and H and N, respectively. For speculator H, the projection is:

$$\begin{aligned} & E^H[X - p|H, p, p(\cdot)] \\ = \quad & H + E^H[N|H, p, p(\cdot)] - p \\ = \quad & H - \frac{b}{b^2 + c^2 + d^2}(bH + c\lambda^H + cE^H[\lambda^N|H, p(\cdot)] + d\lambda_H) + \frac{b}{b^2 + c^2 + d^2}p - p \\ = \quad & H - \frac{b}{b^2 + c^2 + d^2}(bH + (2c + d)\lambda^H) + \frac{b - (b^2 + c^2 + d^2)}{b^2 + c^2 + d^2}p. \end{aligned} \tag{4.38}$$

where the last equality follows after the application of the average opinion rule (4.32).

The REE price function then obtains from the equilibrium condition with an equal number of speculators of type H and N,

$$D^H + D^N = S. \tag{4.39}$$

Filling out the demands:

$$H + N - 2p + E^H[N|H, p, p(\cdot)] + E^N[H|N, p, p(\cdot)] = S. \tag{4.40}$$

Let us now take the position of a speculator of type N who investigates the plausibility of the announced price function against her knowledge that prices are set in equilibrium. To facilitate the analysis, she compares:

$$E^N[p|N, p(\cdot)], \tag{4.41}$$

to

$$E^N[p|N], \tag{4.42}$$

$E^N[p|N, p(\cdot)]$ is based on the announced price function (4.36). Applying the average opinion rule (4.32),

$$\begin{aligned} & E^N[p|N, p(\cdot)] \\ &= bN + cE^N[\lambda^H|N, p(\cdot)] + (c+d)\lambda^N \\ &= bN + cE^N[\lambda^H] + (c+d)\lambda^N \\ &= bN + (2c+d)\lambda^N. \end{aligned} \tag{4.43}$$

Contrast this with $E^N[p|N]$, which is obtained from the equilibrium condition $D^H + D^N = S$. Let us assume, as in example 1, that our speculator N, posits that all other agents of her type continue to use the REE price function to invert prices for information. If she is small, this leads her to infer that equilibrium prices are determined as follows:

$$H + N - 2p + E^H[N|H, p, p(\cdot)] + E^N[H|N, p, p(\cdot)] = S. \tag{4.44}$$

From (4.38),

$$\begin{aligned} H + N - \frac{b}{b^2+c^2+d^2} \big[& b(H+N) + c(\lambda^H + \lambda^N) \\ & + c(E^H[\lambda^N|H, p(\cdot)] + E^N[\lambda^H|N, p(\cdot)]) \\ & + d(\lambda^H + \lambda^N) \big] = S. \end{aligned} \tag{4.45}$$

Our speculator knows that $E^N[\lambda^H|N,p(\cdot)] = E^N[\lambda^H] = \lambda^N$, (see ((4.32)). Hence, p solves:

$$\begin{aligned} H+N- & \frac{b}{b^2+c^2+d^2}(b(H+N)+(2c+d)(\lambda^H+\lambda^N)) \\ & +2\frac{b-(b^2+c^2+d^2)}{b^2+c^2+d^2}p-\frac{bc}{b^2+c^2+d^2}(\lambda^H-E^H[\lambda^N|H,p(\cdot)]) \\ & = S. \end{aligned} \tag{4.46}$$

Rearranging, and using the equations that define (b,c,d), we can simplify (4.46):

$$p = b(H+N) + c(\lambda^H+\lambda^N) + dS - \frac{1}{2}\cdot\frac{c^2}{2c+d}(\lambda^H - E^H[\lambda^N|H,p(\cdot)]).$$

The speculator's higher-order beliefs imply:

$$\begin{aligned} & E^N[p|N] \\ = \;& bN+(2c+d)\lambda^N - \frac{1}{2}\cdot\frac{c^2}{2c+d}E^N[(\lambda^H - E^H[\lambda^N|H,p(\cdot)])] \\ = \;& bN+(2c+d)\lambda^N - \frac{1}{2}\cdot\frac{c^2}{2c+d}E^N[(\lambda^H - \psi\lambda^H)] \\ = \;& bN+(2c+d)\lambda^N - \frac{c^2}{2(2c+d)}(1-\bar{\psi})\lambda^N. \end{aligned} \tag{4.47}$$

When comparing (4.43) with (4.47) it is easy to see that $E^N[p|N,p(\cdot)]$ will not equal $E^N[p|N]$ unless $\bar{\psi}=1$. In other words, our speculator's inferences from the REE price function and from knowledge that price is set in equilibrium will be consistent only if $\bar{\psi}=1$. Again, as seen in example 1, in all cases where $\bar{\psi}\neq 1$, our agent will either reject the perfect foresight information or her own beliefs.

4.4 Discussion

The examples illustrate how heterogeneity in first-order beliefs may conflict with perfect foresight (RE) if agents hold incorrect higher-order beliefs. Effectively, the REE mapping from states to prices provides information to the agents that clashes with their beliefs. A sufficient condition to avoid such an inconsistency is for higher-order beliefs to be correct.

Is this sufficient condition plausible? There are a number of objections. First of all, why would agents disagree about the fundamentals of the economy (first-order beliefs) while they not only agree on higher-order beliefs, but their beliefs about other agents' first-order beliefs are *correct*? Fundamentals of the economy (demands, dividends, supplies, ...) are readily observable, and hence beliefs can be verified. Other agents' beliefs are generally not observable, even at the first order, let alone at higher orders.

In the models, there is no way for higher-order beliefs to become transparent to all agents. The models are supposed to apply to real-world markets, however, and these generally entail far more communication than is accounted for in the theory. Not only can traders observe, and condition on, volume (in the model, only prices are conditioned on); traders generally observe ample order flow information. Volume and order flows may provide information with which market participants can resolve uncertainty about the beliefs of others.

Only a controlled experiment can determine whether this is the case. We leave this for future work.

Chapter 5

Conclusion

This thesis presented three essays which built on existing theoretical and experimental applications of Radner's perfect foresight approach to modelling equilibrium.

The first essay contrasted different strains of information aggregation research. It contained a detailed description of pilot studies used in developing the experimental study presented in the second essay. I reviewed past theoretical and experimental research with an emphasis on perfect foresight, information structure, and information context. I revealed the design procedure used to conduct pilot studies and included procedural details and insights that may be used to develop new experimental research.

The second essay introduced two new stability concepts to the analysis of rational expectations equilibria. Absolute stability is the sensitivity of equilibrium prices to bits of information being lost or garbled. Relative stability concerns the difference between rational expectations and private information equilibrium prices (the larger the difference, the less stable the equilibrium price is).

We used these concepts to dissect two canonical settings of information aggregation, namely: (i) winner-take-all contracts, where binary payment depends on the signals held by the majority of agents; (ii) assets whose payoffs can be easily inferred from the average private signal. The former is unstable under both concepts and was shown to cast doubt on the information implied by prices. The latter is robust, and prices were shown to aggregate information more reliably. Our findings have important implications for the design of prediciton markets, for historical evidence on the Efficient Markets Hypothesis, and for recent work on information aggregation.

In the third essay, we showed that when the perfect foresight approach is used to model differences in beliefs among market participants, agents beliefs will not clash with perfect foresight when higher-order beliefs are correct. In particular, we provided two theoretical examples to illustrate this point: (i) a production-based model inspired by Muth (1961); (ii) a security market model inspired by the asymmetric information/noisy rational expectations equilibrium framework (Grossman and Stiglitz (1980)).

Overall this thesis has made several theoretical contributions by demonstrating how information structure and context can affect the occurrence and feasibility of a perfect foresight rational expectations equilibrium. Hopefully, it has provided the necessary stimulus to promote thought and motivate future research. Thank you for reading.

Appendix A

Appendix for Chapter 2

A.1 Pilot 1 and 2

Period 1

Signal	Number of Subjects	Posterior
BB	2	0.6923
BW	4	0.5
WW	4	0.3077
	Aggregate Net Signal:	WWWW

	Endowment		
Signal	A	B	Cash
WW	0	10	100
WW	0	10	100
WW	0	10	100
WW	0	10	100
BW	10	0	100
BW	10	0	100
BW	0	10	100
BW	0	10	100
BB	10	10	100
BB	10	10	100
Agg 40A:80B:1000Cash			

Period 2

Signal	Number of Subjects	Posterior
BB	3	0.6923
BW	4	0.5
WW	3	0.3077
	Aggregate Net Signal: BW	

	Endowment		
Signal	A	B	Cash
BW	0	10	100
BW	0	10	100
WW	0	10	100
BB	0	10	100
WW	10	0	100
BB	10	0	100
BW	0	10	100
BB	0	10	100
BW	10	10	100
WW	10	10	100
Agg 40A:80B:1000Cash			

Period 3

Signal	Number of Subjects	Posterior
BB	4	0.6923
BW	4	0.5
WW	2	0.3077
	Aggregate Net Signal: BBBB	

	Endowment		
Signal	A	B	Cash
BB	5	0	100
BB	5	0	100
BB	5	0	100
BB	5	0	100
BW	5	0	100
BW	5	0	100
BW	0	5	100
BW	0	5	100
WW	5	5	100
WW	5	5	100
Agg 40A:20B:1000Cash			

Period 4

Signal	Number of Subjects	Posterior
BB	2	0.6923
BW	4	0.5
WW	4	0.3077
	Aggregate Net Signal:	WWWW

	Endowment		
Signal	A	B	Cash
BW	0	10	100
BW	0	10	100
WW	0	10	100
BB	0	10	100
WW	10	0	100
BB	10	0	100
BW	0	10	100
WW	0	10	100
BB	10	10	100
WW	10	10	100
Agg 40A:80B:1000Cash			

Period 5

Signal	Number of Subjects	Posterior
BB	4	0.6923
BW	4	0.5
WW	2	0.3077
	Aggregate Net Signal:	BBBB

Signal	Endowment A	B	Cash
BB	0	10	100
BB	0	10	100
BB	0	10	100
BB	0	10	100
BW	10	0	100
BW	10	0	100
BW	0	10	100
BW	0	10	100
WW	10	10	100
WW	10	10	100
Agg 40A:80B:1000Cash			

Period 6

Signal	Number of Subjects	Posterior
BB	4	0.6923
BW	4	0.5
WW	2	0.3077
	Aggregate Net Signal:	BBBB

Signal	Endowment A	B	Cash
BW	0	10	100
BW	0	10	100
WW	0	10	100
BB	0	10	100
BB	10	0	100
BB	10	0	100
BW	0	10	100
BB	0	10	100
BB	10	10	100
WW	10	10	100
Agg 40A:80B:1000Cash			

Period 7

Signal	Number of Subjects	Posterior
BB	2	0.6923
BW	4	0.5
WW	4	0.3077
	Aggregate Net Signal:	WWWW

	Endowment		
Signal	A	B	Cash
WW	0	10	100
WW	0	10	100
WW	0	10	100
WW	0	10	100
BW	10	0	100
BW	10	0	100
BW	0	10	100
BW	0	10	100
BB	10	10	100
BB	10	10	100
Agg 40A:80B:100Cash			

Period 8

Signal	Number of Subjects	Posterior
BB	3	0.6923
BW	4	0.5
WW	3	0.3077
	Aggregate Net Signal:	BW

	Endowment		
Signal	A	B	Cash
BW	0	10	100
BW	0	10	100
WW	0	10	100
BB	0	10	100
WW	10	0	100
BB	10	0	100
BW	0	10	100
BB	0	10	100
BW	10	10	100
WW	10	10	100
Agg 40A:80B:1000Cash			

Period 9

Signal	Number of Subjects	Posterior
BB	2	0.6923
BW	4	0.5
WW	4	0.3077
	Aggregate Net Signal:	WWWW

	Endowment		
Signal	A	B	Cash
WW	5	0	100
WW	5	0	100
WW	5	0	100
WW	5	0	100
BW	5	0	100
BW	5	0	100
BW	0	5	100
BW	0	5	100
BB	5	5	100
BB	5	5	100
Agg 40A:20B:1000Cash			

Period 10

Signal	Number of Subjects	Posterior
BB	4	0.6923
BW	4	0.5
WW	2	0.3077
	Aggregate Net Signal: BBBB	

	Endowment		
Signal	A	B	Cash
BB	0	10	100
BB	0	10	100
BB	0	10	100
BB	0	10	100
BW	10	0	100
BW	10	0	100
BW	0	10	100
BW	0	10	100
WW	10	10	100
WW	10	10	100
Agg 40A:80B:1000Cash			

A.2 Pilot 1 and 2: Instructions

Information, Risk and Prices
Instructions

Summary

You will be participating in an online trading game. You will be paid based on your performance in that game, plus a fixed sign-up reward. The goal of the experiment is to determine how information affects prices of financial assets. There will be two parts to the experiment. After the first part, in which you trade based on the instructions below, you will be given further hints so you can trade more effectively in the second part.

1 Setting

The experiment consists of a number of replications of the same setting, referred to as ***periods.*** At the beginning of each period, you will be given ***securities and cash***. Markets will open and you will be free to trade securities. You buy securities with cash and you get cash if you sell securities. At the end of the period, the securities expire, after paying ***dividends*** as specified below.

Your ***period earnings*** have two components: (i) the ***dividends*** on the securities you are holding after the market closes, plus (ii) your final ***cash balance*** after trading. At the end of the experiment, ***one period*** will be chosen ***at random*** and your earnings of that period are yours to keep, in addition to a fixed sign-up reward of $20 Australian Dollars.

During the experiment, accounting is done in *Experimental Dollars.* At the end of the experiment, the amount you earned in the randomly selected period will be converted to Australian Dollars at an exchange rate of 5:1.

2 The Urn Problem

At the beginning of each period, ***one out of two*** urns, ***Urn A or Urn B***, gets chosen with equal chance. The chosen urn is ***the same*** for everyone and it is not disclosed until the end of the period. In order to earn money from trading, you have to determine which urn has been chosen.

Both urns contain exactly 10 balls of different colors. ***Urn A*** contains ***6 black*** and ***4 white*** balls, whereas ***Urn B*** contains ***4 black*** and ***6 white*** balls. This is illustrated below.

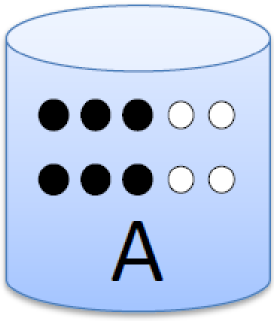

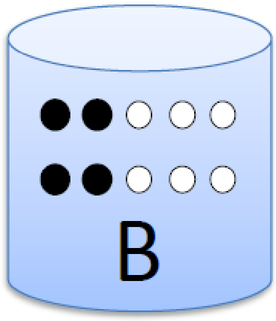

Information regarding the true urn is given to every participant in the form of ***2 balls*** randomly drawn, with replacement, from the ***chosen*** urn. Importantly, for each participant, these 2 balls will be drawn ***independently***.

3 Securities

You will be given two types of securities: ***Stock A*** and ***Stock B***. In each period, their dividends depend on which of the two urns has been chosen. In particular, the dividends per unit are as follows:

- One unit of ***Stock A*** pays ***$10*** if ***Urn A*** has been chosen and nothing if Urn B has been chosen.
- One unit of ***Stock B*** pays ***$10*** if ***Urn B*** has been chosen and nothing if Urn A has been chosen.

You will be able to freely trade ***Stock A***. However, you ***cannot*** trade your holdings in ***Stock B***. Notice that this does not limit you in managing your dividend risk, since the dividends of the two securities are complementary: when Stock A pays, Stock B does not, and vice versa. So, for example, if you are only holding shares of Stock B, which you cannot trade, you can reduce your exposure to risk by ***buying*** the same number of Stock A shares. In contrast, if you are only holding shares of Stock A, you can reduce your risk exposure by merely your ***selling*** Stock A shares.

4 Information

From the drawn balls, you can infer the likelihood of the true urn as follows: Urn A has more black balls than Urn B. So, the more black balls you observe, the more likely Urn A is the chosen urn. Conversely, the more white balls you see, the more likely urn B corresponds to the true urn. Remember, draws will be ***independent*** across participants.

For each period, you will be provided with a sheet that lists your personal ball draws along with the respective implied probabilities that urn A is the chosen urn. Keep this for yourself and do not discuss it with anyone else.

It is to be expected that many of you trade, at which point your collective information may affect prices. If two of you trade, you and your counter-party will jointly determine the trading price. Hence, prices may behave as if your samples were taken together. For instance, as both of you will have received a sample with 2 balls, prices could reflect a collective sample consisting of *4* balls. With 10 traders, the potential sample size eventually increases to 20.

5 Period Earnings

At the end of each period, the chosen urn will be revealed and your earnings will be calculated. In addition to your final cash holdings, your earnings of course depend on the number of units of the correct security you end up holding. Remember, your payment will be based on your earnings from ***one*** randomly drawn period.

6 The Trading Platform

Trading takes place through an electronic trading platform called “Flex-E-Markets”. In Flex-E-Markets you submit limit orders, which are orders to buy or sell at a price you determine, or, if possible, at any better price. Transactions take place from the moment a buy order with a higher price crosses a sell order with a lower price or vice versa. Orders remain valid until you cancel them or the marketplace closes. You will be given ample opportunity to train yourself in submitting and canceling orders.

You can access Flex-E-Markets as follows: use your login information sheet (e-mail and password) and log into the account named ?? at http://adhocmarkets.com/. You should then navigate to the marketplace called ??. Each period, the marketplace will be open for a pre-determined time period (approximately 5 minutes). You will be notified at halftime, as well as one minute before markets close.

Information, Risk and Prices
Instructions

Part 2: Likelihood Table & Examples

The computation of the likelihoods that Urn A is the chosen urn can be rather complicated. From now on, you do not have to compute these probabilities yourself anymore. The following Table lists all possible distributions of the collection of individual draws together with their respective likelihoods. Importantly, the Table lists the likelihood of Urn A as a function of the ***difference in the number of black and white balls*** only. This is because the probability that Urn A is the chosen urn ***solely*** depends on this difference, and not on the total number of balls drawn.

For instance, if 10 balls are drawn, and 6 thereof are black and 4 are white, then there are 2 more black than white balls. According to the Table, the chance that the chosen urn has a majority of black balls (i.e., that it corresponds to Urn A) is 0.6923, i.e., 69.23%. This chance is the same even if only 4 balls had been drawn, and 3 thereof were black while 1 was white. The difference between black and white balls would still be the same, and that is the only thing that matters.

In addition to the Table, we provide you below with specific examples which demonstrate that likelihoods indeed only depend on the difference between the number of black and white balls, and not on the total number of balls drawn.

Difference in the Number of Black and White Balls	Probability of Urn A
10 White	0.0170
9 White	0.0254
8 White	0.0376
7 White	0.0553
6 White	0.0807
5 White	0.1164
4 White	0.1649
3 White	0.2286
2 White	0.3077
White	0.4000
Equal and	0.5000
1 Black	0.6000
2 Black	0.6923
3 Black	0.7714
4 Black	0.8351
5 Black	0.8836
6 Black	0.9193
7 Black	0.9447
8 Black	0.9624
9 Black	0.9746
10 Black	0.9830

Example 1: 10 Ball Signal

1 White Balls
9 Black Balls
8 Black Balls Net

Probability of Urn A: 0.9624

Example 2: 10 Ball Signal

8 White Balls
2 Black Balls
6 White Balls Net

Probability of Urn A: 0.807

Example 3: 20 Ball Signal

10 White Balls
10 Black Balls
Equal White and Black

Probability of Urn A: 0.5

Example 4: Price of Stock A = 0.1164
Implies there are 5 more white balls than black balls

Example 5: Price of Stock A = 0.3077
Implies there are 2 more white balls than black balls

Example 6: Price of Stock A = 0.8351
Implies there are 5 more black balls than white balls

A.3 Pilot 3

Period 1

Signal	Number of Subjects	Individual Posterior
BB	8	0.6923
BW	8	0.5
WW	4	0.3077
	Aggregate Net Signal:	BBBB BBBB

	Endowments		
Signal	Stock A	Stock B	Cash
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
WW	10	0	100
WW	10	0	100
WW	10	0	100
WW	10	0	100
Aggregate 80A:80B:2200Cash			

Period 2

Signal	Number of Subjects	Individual Posterior
BB	6	0.6923
BW	8	0.5
WW	6	0.3077
		Aggregate Net Signal: BW

Endowments

Signal	Stock A	Stock B	Cash
BB	5	0	125
BB	5	0	125
BB	10	0	100
BB	0	5	125
BB	0	5	125
BB	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
WW	5	0	125
WW	5	0	125
WW	10	0	100
WW	0	5	125
WW	0	5	125
WW	0	10	100

Aggregate 80A:80B:2200Cash

Period 3

Signal	Number of Subjects	Individual Posterior
BB	4	0.6923
BW	8	0.5
WW	8	0.3077
	Aggregate Net Signal:	WWWW WWWW

	Endowments		
Signal	Stock A	Stock B	Cash
WW	0	5	125
WW	0	5	125
WW	0	5	125
WW	0	5	125
WW	0	5	125
WW	0	5	125
WW	0	5	125
WW	0	5	125
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BB	10	0	100
BB	10	0	100
BB	10	0	100
BB	10	0	100
Aggregate 80A:80B:2200Cash			

Period 4

Signal	Number of Subjects	Individual Posterior
BB	8	0.6923
BW	8	0.5
WW	4	0.3077
	Aggregate Net Signal:	BBBB BBBB

	Endowments		
Signal	Stock A	Stock B	Cash
BB	5	0	125
BB	5	0	125
BB	5	0	125
BB	5	0	125
BB	5	0	125
BB	5	0	125
BB	5	0	125
BB	5	0	125
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
WW	0	10	100
WW	0	10	100
WW	0	10	100
WW	0	10	100
Aggregate 80A:80B:2200Cash			

Period 5

Signal	Number of Subjects	Individual Posterior
BB	6	0.6923
BW	8	0.5
WW	6	0.3077
		Aggregate Net Signal: BW

Endowments			
Signal	Stock A	Stock B	Cash
BB	5	0	125
BB	5	0	125
BB	5	0	125
BB	5	0	125
BB	10	0	100
BB	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
WW	0	5	125
WW	0	5	125
WW	0	5	125
WW	0	5	125
WW	0	10	100
WW	0	10	100
Aggregate 80A:80B:2200Cash			

Period 6

Signal	Number of Subjects	Individual Posterior
BB	4	0.6923
BW	8	0.5
WW	8	0.3077
	Aggregate Net Signal:	WWWW WWWW

Endowments			
Signal	Stock A	Stock B	Cash
WW	0	5	125
WW	0	5	125
WW	0	5	125
WW	0	5	125
WW	5	0	125
WW	5	0	125
WW	5	0	125
WW	5	0	125
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BB	0	10	100
BB	0	10	100
BB	10	0	100
BB	10	0	100
Aggregate 80A:80B:2200Cash			

Period 7

Signal	Number of Subjects	Individual Posterior
BB	8	0.6923
BW	8	0.5
WW	4	0.3077
	Aggregate Net Signal:	BBBB BBBB

Endowments

Signal	Stock A	Stock B	Cash
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	5	0	125
BB	5	0	125
BB	5	0	125
BB	5	0	125
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
WW	0	10	100
WW	0	10	100
WW	10	0	100
WW	10	0	100

Aggregate 80A:80B:2200Cash

Period 8

Signal	Number of Subjects	Individual Posterior
BB	6	0.6923
BW	8	0.5
WW	6	0.3077
	Aggregate Net Signal:	WWWW WWWW

Endowments			
Signal	Stock A	Stock B	Cash
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	10	100
BB	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
WW	5	0	125
WW	5	0	125
WW	5	0	125
WW	5	0	125
WW	10	0	100
WW	10	0	100
Aggregate 80A:80B:2200Cash			

Period 9

Signal	Number of Subjects	Individual Posterior
BB	4	0.6923
BW	8	0.5
WW	8	0.3077
	Aggregate Net Signal:	WWWW WWWW

	Endowments		
Signal	Stock A	Stock B	Cash
WW	5	0	125
WW	5	0	125
WW	5	0	125
WW	5	0	125
WW	5	0	125
WW	5	0	125
WW	5	0	125
WW	5	0	125
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BB	0	10	100
BB	0	10	100
BB	0	10	100
BB	0	10	100
Aggregate 80A:80B:2200Cash			

Period 10

Signal	Number of Subjects
BB	10
BW	8
WW	2
Aggregate Ratio:	7 B to 3 W

Endowments			
Signal	Stock A	Stock B	Cash
BB	4	0	130
BB	4	0	130
BB	4	0	130
BB	4	0	130
BB	4	0	130
BB	4	0	130
BB	4	0	130
BB	4	0	130
BB	4	0	130
BB	4	0	130
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
WW	0	20	50
WW	0	20	50
Aggregate 80A:80B:2200Cash			

Period 11

Signal	Number of Subjects
BB	2
BW	8
WW	10
Aggregate Ratio:	3 B to 7 W

Endowments

Signal	Stock A	Stock B	Cash
WW	0	4	130
WW	0	4	130
WW	0	4	130
WW	0	4	130
WW	0	4	130
WW	0	4	130
WW	0	4	130
WW	0	4	130
WW	0	4	130
WW	0	4	130
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BB	20	0	50
BB	20	0	50
Aggregate 80A:80B:2200Cash			

Period 12

Signal	Number of Subjects
BB	2
BW	8
WW	8
Aggregate Ratio:	3 B to 7 W

Endowments

Signal	Stock A	Stock B	Cash
WW	4	0	130
WW	4	0	130
WW	4	0	130
WW	4	0	130
WW	4	0	130
WW	4	0	130
WW	4	0	130
WW	4	0	130
WW	4	0	130
WW	4	0	130
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BB	0	20	50
BB	0	20	50

Aggregate 80A:80B:2200Cash

Period 13

Signal	Number of Subjects
BB	10
BW	8
WW	2
Aggregate Ratio:	7 B to 3 W

Endowments			
Signal	Stock A	Stock B	Cash
BB	0	4	130
BB	0	4	130
BB	0	4	130
BB	0	4	130
BB	0	4	130
BB	0	4	130
BB	0	4	130
BB	0	4	130
BB	0	4	130
BB	0	4	130
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
WW	20	0	50
WW	20	0	50
Aggregate 80A:80B:2200Cash			

A.4 Pilot 3: Instructions

Information, Risk and Prices
Instructions

Summary

You will be participating in a series of online trading games. You will receive a fixed sign-up reward and additional pay based on your performance. The goal of this experiment is to determine how information affects prices of financial assets. You will trade based on the instructions below. You will also be given a table and examples so you can trade more effectively.

1 Setting

This experiment is made of separate trading games called ***periods***. At the beginning of each period, you will be given ***stocks and cash***. The market will open and you will be free to trade stocks. You may buy stocks with cash and you will get cash if you sell stocks. At the end of each period, the stocks pay ***dividends*** as specified below and then cease to exist.

Your ***period earnings*** have two components: (i) ***dividends*** from stocks, plus (ii) your ***cash balance***. Both are determined by your holdings at the end of the period when the market closes. At the end of the experiment, ***one period*** will be ***chosen at random*** and your earnings from that period are yours to keep. In addition you will receive a fixed sign-up reward of $20 Australian Dollars.

During the experiment, accounting is done in *Experimental Dollars*. At the end of the experiment, the amount you earned in the randomly selected period will be converted to Australian Dollars at an exchange rate of 5:1.

2 The Bucket Problem

At the beginning of each period, ***one out of two*** buckets is chosen with equal chance: either ***Bucket A or Bucket B***. The choice of this bucket will determine the payoffs of the stocks. The chosen bucket is not revealed until the end of the period. In order to earn money, you should try to determine which bucket has been chosen.

Both buckets contain 10 balls. ***Bucket A*** contains ***6 black*** and ***4 white*** balls, whereas ***Bucket B*** contains ***4 black*** and ***6 white*** balls. This is illustrated below.

Before trading begins, each participant is shown a sample of ***2 balls.*** These balls are ***drawn with replacement*** from the chosen bucket. The balls are drawn independently across participants. This means that your sample is based on ***a different draw*** than that of others.

3 Stocks

In each period you will be endowed with units of two stocks: ***Stock A*** and ***Stock B***. The dividends these stocks pay depend on which of the two buckets was chosen for that period. The dividends per unit stock are as follows:

- One unit of ***Stock A*** pays ***$10*** if ***Bucket A*** has been chosen and nothing if Bucket B has been chosen.
- One unit of ***Stock B*** pays ***$10*** if ***Bucket B*** has been chosen and nothing if Bucket A has been chosen.

You will be able to ***freely*** trade ***Stock A***. However, you ***cannot*** trade ***Stock B***. If you start holding only Stock B, you can manage the risk of not receiving any dividends by buying Stock A because the two stocks are complementary: when Stock A pays a dividend, Stock B does not; when Stock B pays, Stock A does not. By holding an equal number of Stock A and Stock B, it is possible to know your dividends for sure. When you start out with only Stock A, the only way to reduce the risk of not receiving any dividends is to sell Stock A.

4 Information

Bucket A has more black balls than Bucket B. If you have two black balls this suggests that Bucket A is more likely the chosen bucket. Alternatively, if you have two white balls this suggests that Bucket B is more likely to be the chosen bucket.

You will be given a sheet that lists your samples of balls for each period, along with the probabilities that these samples imply that ***Bucket A*** is the chosen bucket. Other participants may have different implied probabilities based on their own balls.

Each trade involves two participants. One can imagine that the information these participants have will be reflected in the price at which they are willing to trade. If they both have the same sample, say two black balls, then they both think Bucket A is more likely to be the chosen one of that period. So, they are willing to trade Stock A at a higher price. Higher prices may thus reveal that there is evidence in favour of Bucket A. Lower prices suggest that trading parties have samples with less black than white balls. Subsequent traders will take this into account, combining the information gained from prior trades together with their own sample. If they too have samples with more black balls, prices will increase even more.

5 Likelihood Table & Examples

Calculating the likelihood that Bucket A is the chosen bucket is complicated. However, you do not have to compute these probabilities yourself. The provided table lists the likelihood of Bucket A as a function of the ***difference in the number of black and white balls*** observed. This is because the probability that Bucket A is the chosen bucket ***only*** depends on this difference and not on the total number of balls shown. If we could observe everyone's balls at the same time and take the difference between the total number of black and white balls, the table would provide us with the implied probability that ***Bucket A*** is the chosen bucket.

For example, if 10 balls are shown, with 6 thereof being black and 4 being white, then we can use the fact there are 2 more black than white balls. According to the Table, the chance that the chosen bucket has a majority of black balls (Bucket A) is 0.6923, or, equivalently, 69.23%. This chance is the same even if only 4 balls had been shown and 3 were black and 1 was white. This is because the difference between black and white balls would still be the same, and the difference is the only thing that matters.

In addition to the table, you are provided with examples which demonstrate that likelihoods depend only on the difference between the number of black and white balls and not on the total number of balls drawn.

6 Period Earnings

At the end of each period, the chosen bucket will be revealed and your earnings will be calculated. In addition to your final cash balance, your period earnings depend on the number of units of the dividend-paying stock you end up holding. Remember, your actual payment for the experiment will be based on your period earnings from only ***one*** randomly chosen period. Therefore, it is crucial that you pay full attention to every single period.

7 The Trading Platform

Trading takes place through an electronic trading platform called "Flex-E-Markets". In Flex-E-Markets you submit limit orders, which are orders to buy or sell at a price you determine, or, if possible, at any better price. Transactions take place from the moment a buy order with a higher price crosses a sell order with a lower price or vice versa. Orders remain valid until you cancel them or the marketplace closes. You will be given ample opportunity to train yourself in submitting and canceling orders.

You can access Flex-E-Markets as follows: use your login information sheet (e-mail and password) and log into the account named “cogent-ace” at http://adhocmarkets.com/. You should then navigate to the marketplace called “Bucket Market”. Each period, the marketplace will open for a pre-determined time period (approximately 5 minutes). You will be notified at halftime, as well as one minute before markets close.

Difference in the Number of Black and White Balls		Probability of Bucket A
10 White	○○○○○○○○○○	0.0170
9 White	○○○○○○○○○	0.0254
8 White	○○○○○○○○	0.0376
7 White	○○○○○○○	0.0553
6 White	○○○○○○	0.0807
5 White	○○○○○	0.1164
4 White	○○○○	0.1649
3 White	○○○	0.2286
2 White	○○	0.3077
White	○	0.4000
Equal	○ and ●	0.5000
1 Black	●	0.6000
2 Black	●●	0.6923
3 Black	●●●	0.7714
4 Black	●●●●	0.8351
5 Black	●●●●●	0.8836
6 Black	●●●●●●	0.9193
7 Black	●●●●●●●	0.9447
8 Black	●●●●●●●●	0.9624
9 Black	●●●●●●●●●	0.9746
10 Black	●●●●●●●●●●	0.9830

Example 1: 10 Ball Signal

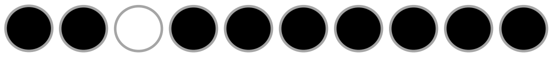

White Balls: 1
Black Balls: 9
Difference: 8 Black
Implied Probability of Urn A: 0.9624

Example 2: 10 Ball Signal

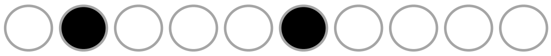

White Balls: 8
Black Balls: 2
Difference: 6 White
Implied Probability of Urn A: 0.0807

Example 3: 20 Ball Signal

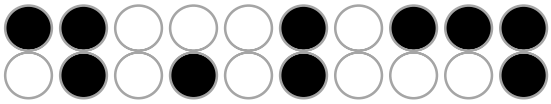

White Balls: 10
Black Balls: 10
Difference: 0
Implied Probability of Urn A: 0.5000

Example 4: Price of Stock A = 1.16
Implies there are 5 more white balls than black balls.

Example 5: Price of Stock A=3.08
Implies there are 2 more white balls than black balls.

Example 6: Price of Stock A=8.84
Implies there are 5 more black balls than white balls.

Information, Risk and Prices
Instructions - Part 2

1 New Game

Now, instead of two buckets, there will be ***only one bucket***. This bucket contains ***10 balls*** which can be either black or white. So, there could be anything between 0 and 10 black balls in the bucket.

You will again receive an independent signal about the number of black and white balls in the bucket. Similar as before, your signal will consist of ***2 balls.*** These balls are ***drawn with replacement*** from the bucket.

2 New Payoffs

As before, you will be endowed with units of two stocks: ***Stock A*** and ***Stock B***. However, the dividends these stocks pay now depend on the ***number*** of black balls in the bucket:

- One unit of ***Stock A*** pays ***$1*** for every ***black*** ball in the bucket. For example, if there are 6 black balls in the bucket, one unit of Stock A pays $6.
- One unit of ***Stock B*** pays ***$1*** for every ***white*** ball in the bucket. For example, if there are 6 black balls in the bucket, then it must also contain 4 white balls. Hence, one unit of Stock B would pay $4.

A.5 Pilot 4

Period 1: 2 Buckets: WWWW WWWW Aggregate Signal, Trade B, True State B

Signal	Number of Subjects	Individual Posterior
BB	4	0.6923
BW	8	0.5
WW	8	0.3077

Endowments			
Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	8	0	80
BB	0	8	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
Aggregate 80A:80B:1600Cash			

Period 2: 2 Buckets: BW Aggregate Signal, Trade A, True State A

Signal	Number of Subjects	Individual Posterior
BB	6	0.6923
BW	8	0.5
WW	6	0.3077

	Endowments		
Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80
WW	0	8	80

Aggregate 80A:80B:1600Cash

Period 3: 2 Buckets: BBBB BBBB Aggregate Signal, Trade B, True A

Signal	Number of Subjects	Individual Posterior
BB	8	0.6923
BW	8	0.5
WW	4	0.3077

Endowments

Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80

Aggregate 80A:80B:1600Cash

Period 4: 2 Buckets: BW Aggregate Signal, Trade A, True B

Signal	Number of Subjects	Individual Posterior
BB	6	0.6923
BW	8	0.5
WW	6	0.3077

	Endowments		
Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80
WW	0	8	80

Aggregate 80A:80B:1600Cash

Period 5: 2 Buckets: WWWW WWWW Aggregate Signal, Trade B, True B

Signal	Number of Subjects	Individual Posterior
BB	4	0.6923
BW	8	0.5
WW	8	0.3077

	Endowments		
Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	8	0	80
BB	0	8	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80

Aggregate 80A:80B:1600Cash

Period 6: 2 Buckets Case 1: BBBB BBBB Aggregate Signal, Trade A, True State A

Signal	Number of Subjects	Individual Posterior
BB	8	0.6923
BW	8	0.5
WW	4	0.3077

Endowments

Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80

Aggregate 80A:80B:1600Cash

Period 7: Correlation Case 1: BBBB BBBB Aggregate Signal and Negative Signal Endowment Correlation, Trade B, True State A

Signal	Number of Subjects	Individual Posterior
BB	8	0.6923
BW	8	0.5
WW	4	0.3077

	Endowments		
Signal	Stock A	Stock B	Cash
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BB	0	5	125
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	0	10	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
BW	10	0	100
WW	10	0	100
WW	10	0	100
WW	10	0	100
WW	10	0	100
Aggregate 80A:80B:2200Cash			

Period 8: 1 Bucket: 30 % Bucket Ratio 3:7, Total 12 Black 28 White, Trade A

Signal	Number of Subjects
BB	2
BW	8
WW	10

	Endowments		
Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80

Aggregate 80A:80B:1600Cash

Period 9: 1 Bucket Case 2: 80 % Black, Bucket Ratio 8:2, Total 32 Black 8 White, Trade B

Signal	Number of Subjects
BB	10
BW	8
WW	2

Endowments

Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	0	8	80

Aggregate 80A:80B:1600Cash

Period 10: 1 Bucket Case 1: 90 % Black, Bucket Ratio 9:1, Total 36 Black 4 White, Trade A

Signal	Number of Subjects
BB	16
BW	4
WW	0

Endowments

Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80

Aggregate 80A:80B:1600Cash

Period 11: 1 Bucket Case 3: 50 % Bucket Ratio 5:5, Total 20 Black 20 White, Trade B

Signal	Number of Subjects
BB	6
BW	8
WW	6

Endowments			
Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	8	0	80
BB	8	0	80
BB	0	8	80
BB	0	8	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80
WW	0	8	80

Aggregate 80A:80B:1600Cash

Period 12: 1 Bucket Case 4: 30 % Bucket Ratio 2:8, Total 8 Black 32 White, Trade A

Signal	Number of Subjects
BB	2
BW	8
WW	10

	Endowments		
Signal	Stock A	Stock B	Cash
BB	8	0	80
BB	0	8	80
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80

Aggregate 80A:80B:1600Cash

Period 13: 1 Bucket Case 5: 10 % Black, Bucket Ratio 1:9, Total 4 Black 36 White, Trade B

Signal	Number of Subjects
BB	0
BW	4
WW	16

Endowments			
Signal	Stock A	Stock B	Cash
BW	8	0	80
BW	8	0	80
BW	0	8	80
BW	0	8	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	8	0	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
WW	0	8	80
Aggregate 80A:80B:1600Cash			

A.6 Pilot 4: Instructions

Information, Risk and Prices
Instructions

Summary

You will be participating in a series of online trading games. You will receive a fixed sign-up reward and additional pay based on your performance. The goal of this experiment is to determine how information affects prices of financial assets.

1 Setting

This experiment is made of separate trading games called ***periods***. At the beginning of each period, you will be given ***stocks and cash***. The market will open and you will be free to trade stocks. You may buy stocks with cash and you will get cash if you sell stocks. At the end of each period, the stocks pay ***dividends*** as specified below and then stop to exist.

Your ***period earnings*** have two components: (i) ***dividends*** from stocks, plus (ii) your ***cash balance***. Both are determined by your holdings at the end of the period when the market closes. At the end of the experiment, ***one period*** will be ***chosen at random*** and your earnings from that period are yours to keep. In addition you will receive a fixed sign-up reward of $20 Australian Dollars.

During the experiment, accounting is done in *Experimental Dollars*. At the end of the experiment, the amount you earned in the randomly selected period will be converted to Australian Dollars at an exchange rate of 5:1.

2 The Bucket Problem

At the beginning of each period, ***one out of two*** buckets is chosen with ***equal chance***: either ***Bucket A or Bucket B*** with 50% probability each. The choice of this bucket will determine the dividends of the stocks. The chosen bucket is ***the same for everybody*** and is not revealed until the end of the period. In order to earn money, you should try to determine which bucket has been chosen.

Both buckets contain 10 balls. ***Bucket A*** contains ***6 black*** and ***4 white*** balls, whereas ***Bucket B*** contains ***4 black*** and ***6 white*** balls. This is illustrated below:

In each period, before trading begins, each participant is shown a sample of ***2 balls.*** These balls are ***drawn with replacement*** from the chosen bucket. The balls are drawn independently across participants. This means that your sample is based on ***a different draw*** than that of others. This is illustrated below:

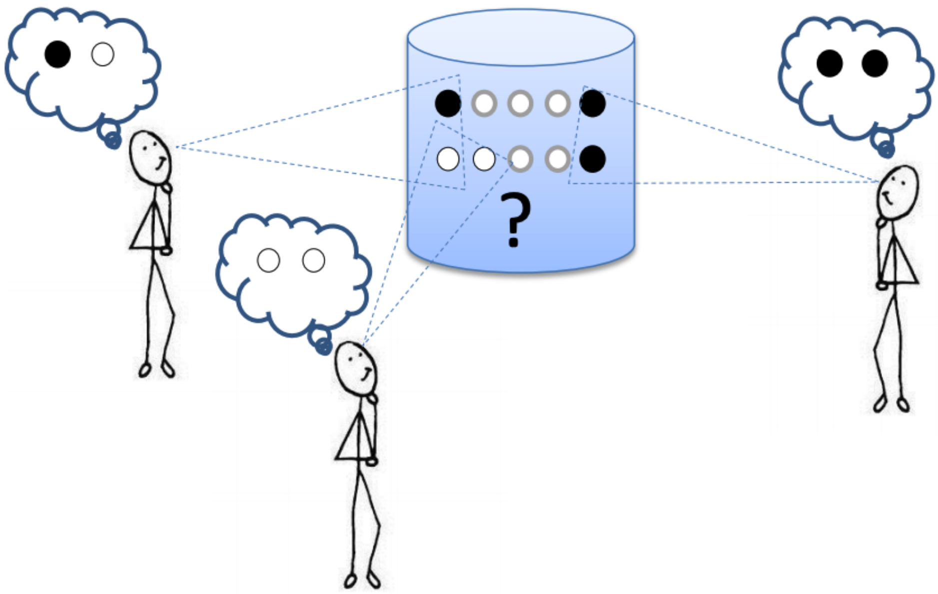

Calculating the likelihood that, for example, Bucket A is the chosen bucket is complicated. However, you do not have to compute these probabilities yourself. For each period, we will provide you with the following information:

1. The probability that Bucket A is the chosen bucket based on ***your private sample of 2 balls only***
2. A calculator that computes the probability that Bucket A is the chosen bucket based on the ***difference in the number of black and white balls*** across ***all independent samples in the whole room***

Notice, the first probability (1.) can be different for different participants.

3 Stocks

In each period you will be endowed with units of two stocks: ***Stock A*** and ***Stock B***. The dividends these stocks pay depend on which of the two buckets was chosen for that period. The dividends per unit stock are as follows:

- One unit of ***Stock A*** pays ***$10*** if ***Bucket A*** has been chosen and nothing if Bucket B has been chosen.
- One unit of ***Stock B*** pays ***$10*** if ***Bucket B*** has been chosen and nothing if Bucket A has been chosen.

You will be able to ***freely*** trade ***either*** Stock A ***or*** Stock B. This means in every round there is one stock which you ***cannot*** trade! However, you can still always manage your risk.

For example, if you start holding only the non-traded stock, you can manage the risk of not receiving any dividends by buying the traded stock because the two stocks are complementary: when Stock A pays a dividend, Stock B does not; when Stock B pays, Stock A does not. By holding an equal number of Stock A and Stock B, you always know your dividends for sure. When you start out with holdings in the traded stock only, the only way to reduce the risk of not receiving any dividends is to sell your holdings.

4 Information

Bucket A has more black balls than Bucket B. If you have two black balls this suggests that Bucket A is more likely the chosen bucket. Alternatively, if you have two white balls this suggests that Bucket B is more likely to be the chosen bucket.

Each trade involves two participants. One can imagine that the information these participants have will be reflected in the price at which they are willing to trade. If they both have the same sample, say two black balls, then they both think Bucket A is more likely to be the chosen one of that period. So, if Stock A can be traded during that period, they are willing to trade Stock A at a higher price. Higher prices may thus reveal that there is evidence in favour of Bucket A. Lower prices suggest that trading parties have samples with less black than white balls. Subsequent traders will take this into account, combining the information gained from prior trades together with their own sample. If they too have samples with more black balls, prices will increase even more.

5 Period Earnings

At the end of each period, the chosen bucket will be revealed and your earnings will be calculated. In addition to your final cash balance, your period earnings depend on the number of units of the dividend-paying stock you end up holding. Remember, your actual payment for the experiment will be based on your period earnings from only ***one*** randomly chosen period. Therefore, it is crucial that you pay full attention to every single period.

6 The Trading Platform

Trading takes place through an electronic trading platform called “Flex-E-Markets”. In Flex-E-Markets you submit limit orders, which are orders to buy or sell at a price you determine, or, if possible, at any better price. Transactions take place from the moment a buy order with a higher price crosses a sell order with a lower price or vice versa. Orders remain valid until you cancel them or the marketplace closes. You will be given ample opportunity to train yourself in submitting and canceling orders.

You can access Flex-E-Markets as follows: go to ***http://adhocmarkets.com/*** and log into the account named ***"cogent-ace"*** using your login information sheet (***e-mail*** and ***password***). You should then navigate to the marketplace called "Bucket Market". Each period, the marketplace will open for a pre-determined time period (approximately 5 minutes). You will be notified at halftime, as well as one minute before markets close.

You can access the information interface as follows: go to ...

Information, Risk and Prices
Instructions - Part 2

1 New Game

Now, instead of two buckets, there will be ***only one bucket***. This bucket contains ***10 balls*** which can be either black or white. So, there could be anything between 0 and 10 black balls in the bucket.

You will again receive an independent signal about the number of black and white balls in the bucket. Similar as before, your signal will consist of ***2 balls.*** These balls are ***drawn with replacement*** from the bucket.

2 New Payoffs

As before, you will be endowed with units of two stocks: ***Stock A*** and ***Stock B***. However, the dividends these stocks pay now depend on the ***number*** of black balls in the bucket:

- One unit of ***Stock A*** pays ***$1*** for every ***black*** ball in the bucket. For example, if there are 6 black balls in the bucket, one unit of Stock A pays $6.
- One unit of ***Stock B*** pays ***$1*** for every ***white*** ball in the bucket. For example, if there are 6 black balls in the bucket, then it must also contain 4 white balls. Hence, one unit of Stock B would pay $4.

A.7 Posterior Mean of the One-Urn Design With Replacement

In the fourth pilot we had one-urn treatments where stock-A paid out according to the true proportion of black balls in the single urn. Subjects' signals were two balls each drawn with replacement from an urn with ten balls. They did not know the true proportion.

I will now present the Bayesian representation of this signal processing problem as presented in example 205 in Liu and Wasserman, (2014). [1]

Let $\mathcal{D}_n = X_1, ..., X_n$ be the entire set draws distributed to subjects. Each ball is a Bernoulli draw, $X_1, ...X_n \sim Bernoulli(\theta)$. Stock-A pays $\theta \times \$10$, so subjects use their signals to estimate θ. Suppose we take the uniform distribution $\pi(\theta) = 1$ as the prior.

By Baye's theorem, the posterior based on the entire set of signals is,

$$p(\theta|\mathcal{D}_n) \propto \pi(\theta)\mathcal{L}_n(\theta) = \theta^{S_n}(1-\theta)^{n-S_n} = \theta^{S_n+1-1}(1-\theta)^{n-S_n+1-1}, \tag{A.1}$$

where $\mathcal{L}_n$ is the likelihood function, and $S_n = \sum_{i=1}^{n} X_i$ is the number of successes.

Recall that a random variable θ on an interval (0,1) has a *Beta* distribution with parameters α and β if its density is,

$$\pi_{\alpha,\beta}(\theta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}. \tag{A.2}$$

[1] I first numerically simulated the result presented here, and was then shown a similar version of the analytical proof by Felix Fattinger. He deserves credit for first finding the proof with the *Beta* distribution link.

We can write the posterior for θ as a Beta distribution with parameters S_{n+1} and $n - S_n + 1$.

$$p(\theta|\mathcal{D}_n) = \frac{\Gamma(n+2)}{\Gamma(S_n+1)\Gamma(n-S_n+1)}\theta^{(S_n+1)-1}(1-\theta)^{(n-S_n+1)-1}. \tag{A.3}$$

By re-arranging A.1 to resemble the functional portion of the *Beta* distribution, we recovered the corresponding normalizing constant from form of the *Beta* distribution.

The mean of a $Beta(\alpha, \beta)$ is $\alpha\backslash(\alpha + \beta)$, so the Bayesian posterior estimator is:

$$\bar{\theta} = \frac{S_n + 1}{n + 2}. \tag{A.4}$$

Which can also be expressed as a weighted average between the maximum likelihood estimate, $\hat{\theta} = S_n\backslash n$ and the prior mean, $\tilde{\theta} = 1\backslash 2$,

$$\bar{\theta} = \lambda_n\hat{\theta} + (1 - \lambda_n)\tilde{\theta}. \tag{A.5}$$

With weights, $\lambda_n = n\backslash(n + 2)$, the equation can be written as:

$$\bar{\theta} = (\frac{n}{n+2}) \times \frac{S_n}{n} + (\frac{2}{n+2}) \times \frac{1}{2}. \tag{A.6}$$

The $\tilde{\theta}$ represents a weight on the prior and in this case a bias towards the center. If there many signals drawn the bias will approach zero.

For the following pilots and final experiment, we decided to use a straight forward approach for the one-urn design that avoided this detail.

Appendix B

Appendix for Chapter 3

B.1 Private-Information Equilibrium and Average Private Value

Let, e_i be the initial endowment of agent i, x_i be the demand, and z be the payoff from the traded security. In a Private-Information equilibrium agents only utilize their private information, so agent i's demand can be expressed a function of the difference between their expected private value and the price:

$$x_i = a_i(E_i[z] - p),$$

where, a_i is a function of risk aversion $\mathcal{A}_i$ and posterior precision ρ_i:

$$a_i = \phi(\mathcal{A}_i, \rho_i).$$

Case 1: Gaussian signals

If we assume all agents have the same risk aversion, $\mathcal{A}_i = \mathcal{A}, \ \forall i$, and private signals are drawn from a Gaussian distribution which gives identical posterior variance for each agent, $\rho_i = \rho, \ \forall i$, then $a_i = a, , \ \forall i$.

We can express the equilibrium condition as an average and insert the demand functions rearrange and pull a in front of the summation:

$$\frac{1}{N}\sum_i^N (x_i - e_i) = 0,$$
$$a\frac{1}{N}\sum_i^N (E_i[z] - p) = \frac{1}{N}\sum_i^N e_i.$$

If agents' endowments are centered around zero, the summation on the right is zero, leaving price as a function of the average individual signal:

$$p = \frac{1}{N}\sum_i^N (E_i[z])$$

Case 2: Orthogonal Private Values and Precision

If the private posterior means are orthogonal to the posterior precisions, the expected product of the posterior mean and precision can be expressed as the product of the expected mean and the expected precision. Recall, if $X \perp Y$, then $E[XY] = E[X]E[Y]$. This allows the average demand to be expressed as:

$$\frac{1}{N}\sum_i^N a_i(E_i[z] - p) \approx \left(\frac{1}{N}\sum_i^N a_i\right)\left(\frac{1}{N}\sum_i^N (E_i[z] - p)\right).$$

If we substitute this into the equilibrium condition and use the same argument from case 1 for the summation on the right being zero, we can again recover the equilibrium prices as the average posterior mean:

$$\left(\frac{1}{N}\sum_{i}^{N} a_i\right)\left(\frac{1}{N}\sum_{i}^{N}(E_i[z]-p)\right)=\frac{1}{N}\sum_{i}^{N} e_i,$$
$$p=\frac{1}{N}\sum_{i}^{N}(E_i[z]).$$

The above derivations illustrate how the Private-Information equilibrium can be expressed as the average posterior mean. In our experimental set up, agents, on average, have a zero net difference in their holdings of Black and White stocks and the posterior means are approximately orthogonal to the posterior precisions, for these reasons we use the average posterior mean as an approximation for the private value equilibrium.

B.2 Posterior Means

In the one-urn case, let N be the aggregate number of balls distributed, n the number of balls in the observed signal, K the aggregate number of black balls, and k the number of black balls in the observed signal. Bayes' law can be expressed as:

$$\begin{aligned} P(K|k) &= \frac{P(k|K)P(K)}{P(k)}, \\ &= \frac{P(k|K)P(k)}{\sum\limits_{K=k}^{N-(n-k)} P(k|K)P(K)}. \end{aligned} \tag{B.1}$$

Assuming the uniform prior, it follows:

$$\begin{aligned} P(K|k) &= \frac{\frac{\binom{K}{k}\binom{N-K}{n-k}}{\binom{N}{n}}\frac{1}{N+1}}{\frac{1}{n+1}}, \\ &= \frac{\binom{K}{k}\binom{N-K}{n-k}}{\binom{N}{n}}\frac{n+1}{N+1}. \end{aligned} \tag{B.2}$$

In the two-urn case, after a coin flip, signals are drawn from one of two binomial distributions with complementary success probabilities. Thus, the probability of receiving a black ball from the Black Urn is the same as the probability of getting a white ball from the White Urn. As a result the Bayesian posterior function is a mixture of two binomial distributions, however, the complimentary nature of the probabilities helps to reduce the expression.

Let again n denote the number of balls in the observed signal and k its corresponding number of black balls. Moreover, let p denote the proportion of black (white) balls in the Black (White) Urn and $q = 1 - p$ the remaining proportion of white (black) balls. Bayes' law then implies:

$$
\begin{aligned}
P(\text{Black Urn}|k) &= \frac{P(k|\text{Black Urn})P(\text{Black Urn})}{P(k)}, \\
&= \frac{\binom{n}{k} p^k q^{n-k} \frac{1}{2}}{\frac{1}{2}\binom{n}{k} p^k q^{n-k} + \frac{1}{2}\binom{n}{k} q^k p^{n-k}}, \\
&= \frac{p^k q^{n-k}}{p^k q^{n-k} + q^k p^{n-k}}, \\
&= \frac{1}{1 + \left(\frac{p}{q}\right)^{n-2k}}.
\end{aligned} \tag{B.3}
$$

Let $m = n - k$ denote the number of white balls in the signal, then:

$$
P(\text{Black Urn}|k) = \left(1 + \left(\frac{p}{q}\right)^{m-k}\right)^{-1}. \tag{B.4}
$$

Hence, the probability that the Black Urn was selected as underlying urn can be written as a function of the signal's difference in the number of black and white balls. Intuitively, if $k = m$, the posterior is equal to the 50:50 prior.

B.3 Drift Diagrams

To generate the price drift diagrams, we regress trade prices within a round on the trade count and recover the drift per trade and the price level for each round:

$$p_i = \hat{\alpha} + \hat{\beta} \cdot i + \epsilon_i, \tag{B.5}$$

We then construct an estimate of the per trade direction toward the fully revealing and Private-Information equilibria using the fixed price level from the first regression and N the number of trades in the round:

$$\beta_{FRE} = \frac{p_{FRE} - \hat{\alpha}}{N}, \tag{B.6}$$

$$\beta_{PE} = \frac{p_{PE} - \hat{\alpha}}{N}. \tag{B.7}$$

To get the component of the PE slope that is orthogonal to the FRE slope, the FRE slope is regressed on the PE slope and the orthogoanl residuals are obtained:

$$\beta_{PE} = \hat{\gamma}_1 \beta_{FRE} + \epsilon_{PE}, \tag{B.8}$$

$$\beta_{PE}^{\perp} = \epsilon_{PE}. \tag{B.9}$$

The slope estimates from each round are then used in a cross-sectional generalized linear model (GLM) with dummies for treatment and round, and clusters at the session, treatment and round level, with the fully revealing slope and the orthogonolized private-information slope as independent variables:

$$\hat{\beta} = \sum_{j=1}^{2} \sum_{i=1}^{4} \hat{\gamma}_{FRE,b_j,r_i} \cdot \beta_{FRE} \cdot \mathbb{1}_{b_j,r_i} + \sum_{j=1}^{2} \sum_{i=1}^{4} \hat{\gamma}_{PE,b_j,r_i} \cdot \beta_{PE}^{\perp} \cdot \mathbb{1}_{b_j,r_i} + \eta. \tag{B.10}$$

The above regressions were estimated for each round and each treatment. The drift diagrams were then generated using interpolation and error bar simulation for the area between the rounds.

B.4 Session Trade Price Diagrams

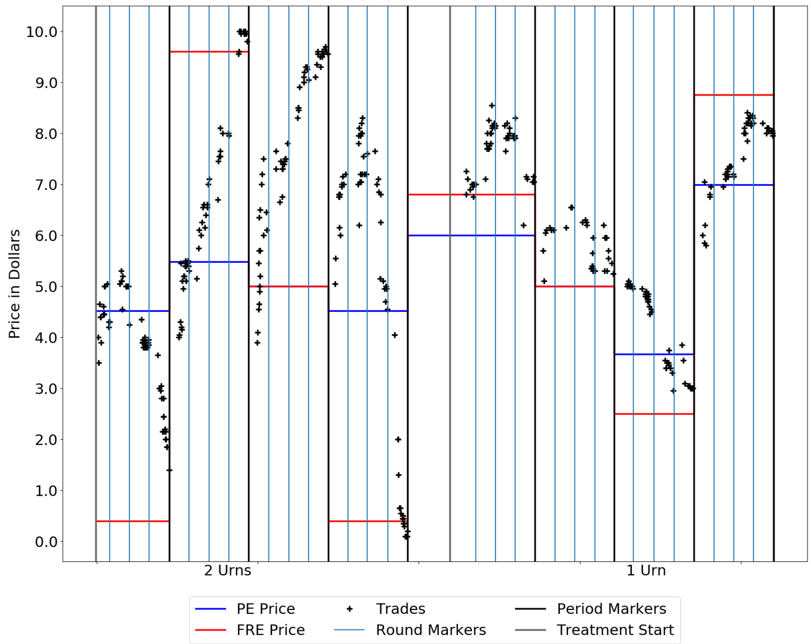

Figure B.1: Session 1 Trade Prices

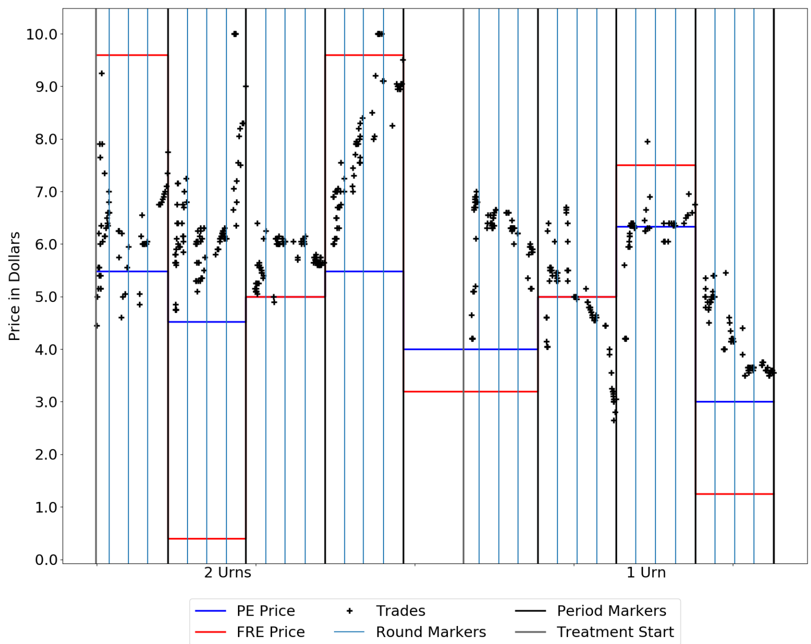

Figure B.2: Session 2 Trade Prices

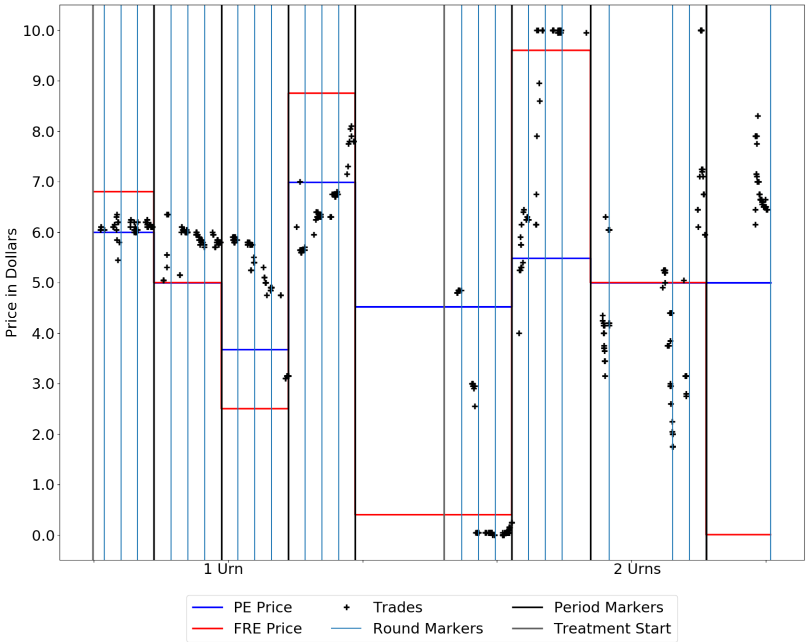

Figure B.3: Session 3 Trade Prices

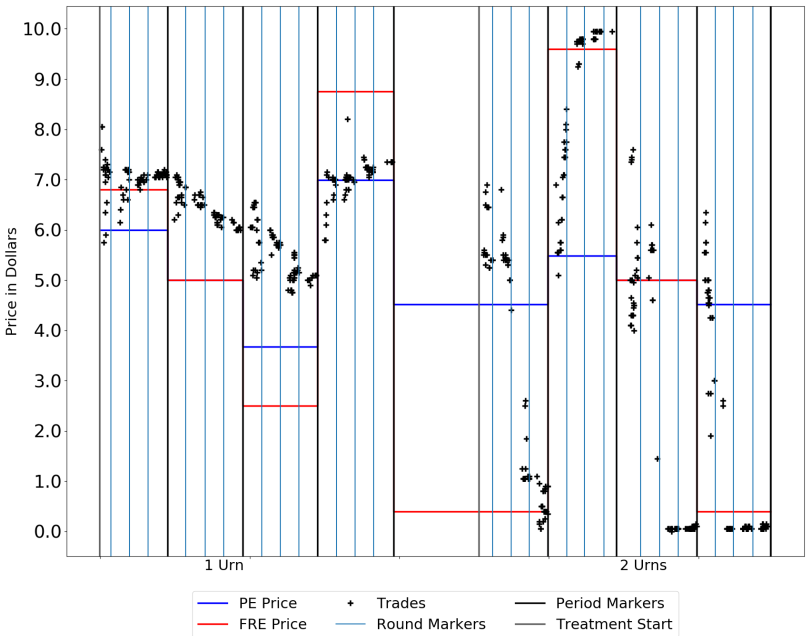

Figure B.4: Session 4 Trade Prices

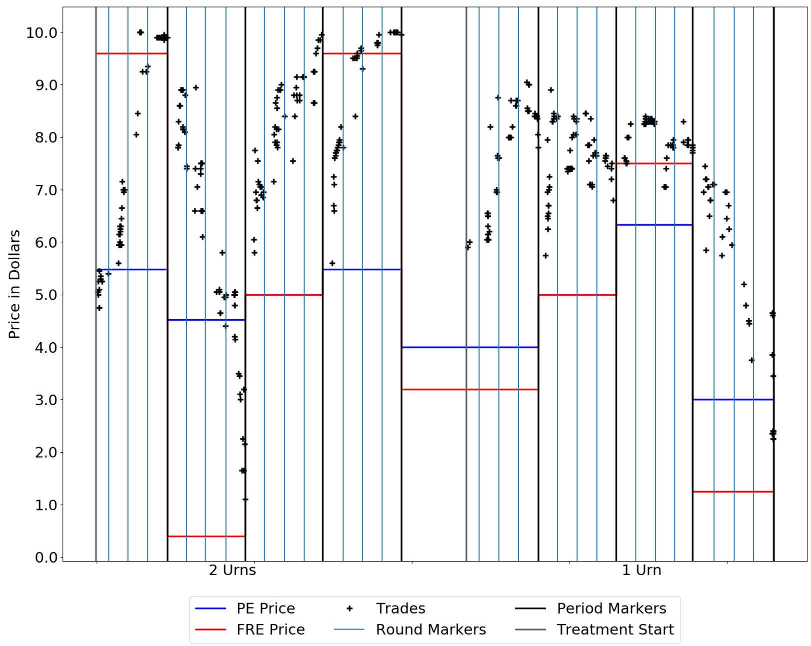

Figure B.5: Session 5 Trade Prices

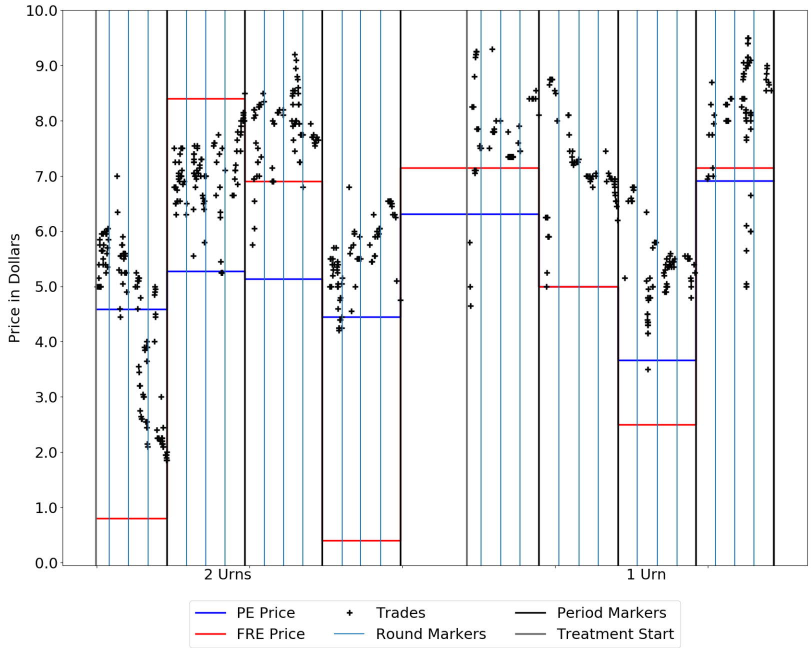

Figure B.6: Session 6 Trade Prices

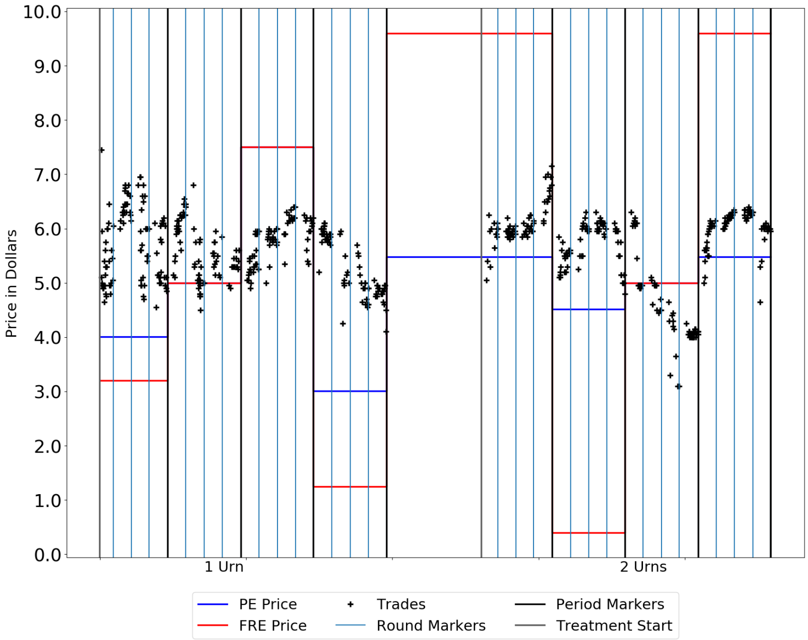

Figure B.7: Session 7 Trade Prices

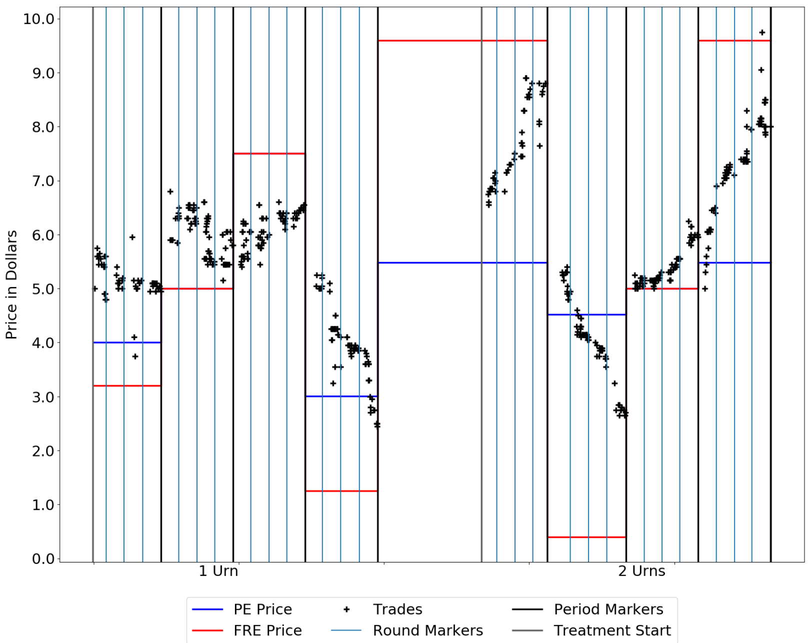

Figure B.8: Session 8 Trade Prices

B.5 Earnings

The additional tests and post hoc analysis presented below support the arguments discussed in the body of the paper. Tables B.1 and B.2 present Earnings ANOVA with Bonferroni corrections. We can see that average earnings are not significantly different before auction costs are deducted. When costs are deducted, those who win two auctions earn significantly less than the Non-Winners.

Table B.1: Mean Earnings Less Cost Comparison By Trader Type (Bonferroni Adjusted)

Row Mean -	One Urn			
Column Mean	R1 Auction Winners	R2 Auction Winners	R3 Auction Winners	Non-Winners
R2 Auction Winners	717.191			
	1			
R3 Auction Winners	557.005	-160.185		
	1	1		
Non-Winners	1430.36	713.166	873.352	
	1	1	1	
Double Winners	-6140.58	-6857.77	-6697.58	-7570.93
	0.397	0.132	0.194	0.017

Table B.2: Mean Earnings Before Auction Payments Comparison By Trader Type (Bonferroni Adjusted)

Row Mean -	Two Urn			
Column Mean	R1 Auction Winners	R2 Auction Winners	R3 Auction Winners	Non-Winners
R2 Auction Winners	-2176			
	1			
R3 Auction Winners	-746.689	1429.32		
	1	1		
Non-Winners	645.527	2821.53	1392.22	
	1	0.317	1	
Double Winners	-4123.65	-1947.64	-3376.96	-4769.17
	1	1	1	0.704

Table B.3 presents Kolmogorov-Smirnov tests for differences in distributions. Post hoc hypothesis testing, conditional on prices in the 2-urn setting not converging to the efficient price, we can see that 1-urn and 2-urn Non-Winners have significantly different distributions. If we look at Figure B.9, on the following page, we can see that the 2-urn Non-Winner's earnings have much fatter tails, which makes sense if prices fail to converge to the FRE.

In both settings, auction winner and Non-winner earnings distributions are significantly different when auction costs are subtracted from the winner's earnings. These distributions are plotted in Figures B.10 and B.11.

Table B.3: Expected Earnings Distributions Kolmogorov-Smirnov Tests (Bonferoni Adjusted)

Earnings			Earnings Less Cost		
One Urn Vs Two Urns					
Smaller group	D	P-value	Smaller group	D	P-value
One Urn Winners	0.0667	0.678	One Urn Winners	0.0654	0.688
Two Urn Winners	-0.0987	0.427	Two Urn Winners	-0.0876	0.511
Combined K-S:	0.0987	0.788	Combined K-S:	0.0876	0.891
One Urn Vs Two Urns					
Smaller group	D	P-value			
One Urn Non-Winenrs	0.0877	0.041			
Two Urn Non-Winners	-0.1101	0.006			
Combined K-S:	0.1101	0.013			
One Urn					
Smaller group	D	P-value	Smaller group	D	P-value
Non-Winners	0.083	0.359	Non-Winners	0.0351	0.833
Winners	-0.0643	0.541	Winners	-0.1986	0.003
Combined K-S:	0.083	0.686	Combined K-S:	0.1986	0.006
Two Urns					
Smaller group	D	P-value	Smaller group	D	P-value
Non-Winners	0.082	0.388	Non-Winners	0.017	0.96
Winners	-0.0812	0.395	Winners	-0.1928	0.005
Combined K-S:	0.082	0.731	Combined K-S:	0.1928	0.011

Figure B.9: Non-Auction Winner Expected Earnings

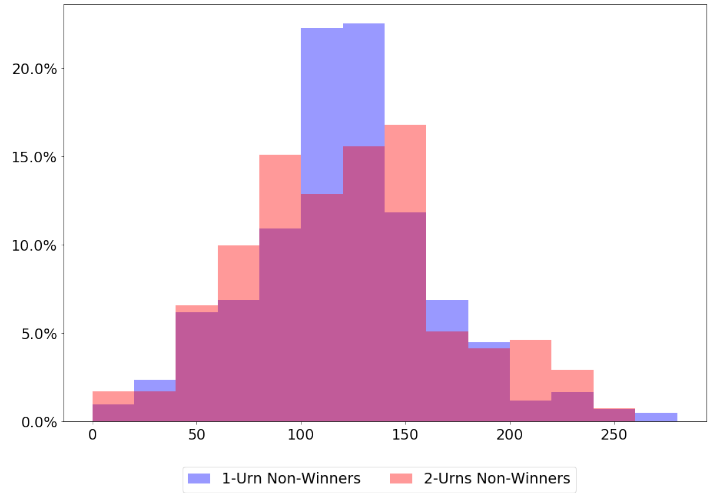

Figure B.10: 1-Urn Expected Earnings Less Auction Payments

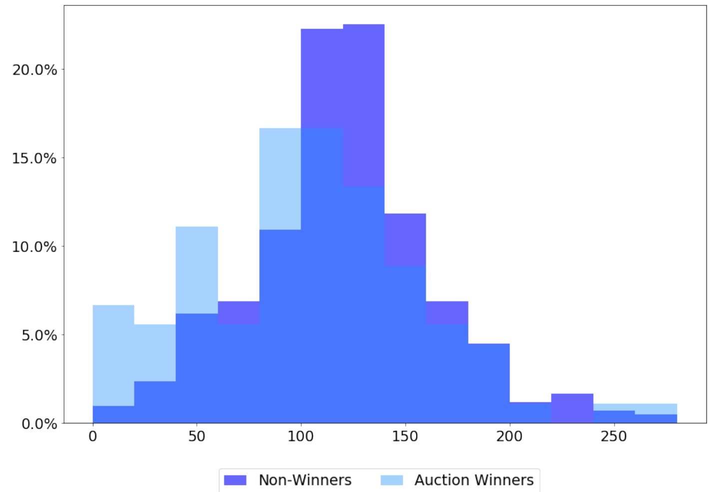

Figure B.11: 2-Urn Expected Earnings Less Auction Payments

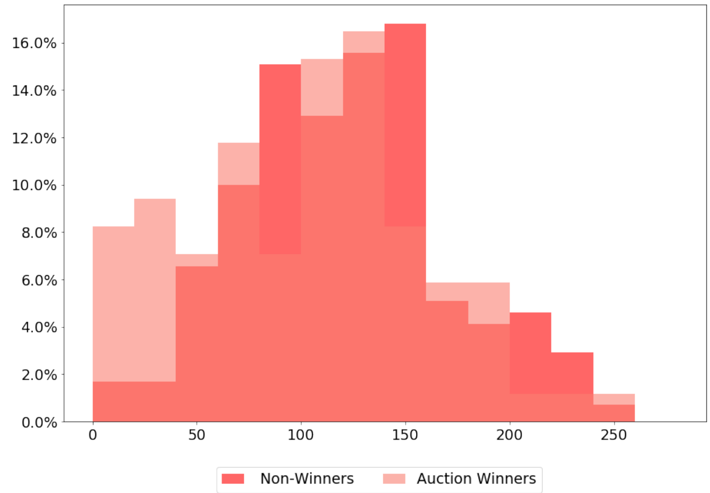

B.6 Holdings

Below are additional tests for holdings. Table B.4 shows that on average auction winners and losers do not hold significantly different portfolios; meaning, all groups have the same level of agreement about prices on average. Table B.5 shows that only the round 1 auction winners average portfolios are significantly different from that of the Non-Winners in the 2-urn setting. This is likely because prices in the first round are far away from the efficient price.

Table B.4: 1-Urn: Mean Absolute Holdings Deviation from Fully Hedged Comparison By Trader Type (Bonferroni Adjusted)

Row Mean -	One Urn			
Column Mean	R1 Auction Winners	R2 Auction Winners	R3 Auction Winners	Non-Winners
R2 Auction Winners	-2.26923			
	1			
R3 Auction Winners	-0.576923	1.69231		
	1	1		
Non-Winners	-1.30952	0.959716	-0.732592	
	1	1	1	
Double Winners	-3.26923	-1	-2.69231	-1.95972
	1	1	1	1

Table B.5: 2-Urns: Mean Absolute Holdings Deviation from Fully Hedged Comparison By Trader Type (Bonferroni Adjusted)

Row Mean - Column Mean	Two Urn R1 Auction Winners	R2 Auction Winners	R3 Auction Winners	Non-Winners
R2 Auction Winners	-3.31054			
	1			
R3 Auction Winners	-2.08593	1.22462		
	1	1		
Non-Winners	-3.5488	-0.238256	-1.46287	
	0.033	1	1	
Double Winners	-4.64021	-1.32967	-2.55429	-1.09141
	1	1	1	1

Table B.6 shows that only auction winners and losers in the 2-Urn setting have significantly different holdings distributions. This is not surprising when considering the result from table B.5.

Table B.6: Absolute Holdings Deviation from Fully Hedged Distributions Kolmogorov-Smirnov Tests

Absolute Holdings Deviation from Fully Hedged Distribution		
One Urn Vs Two Urns: Auction Winners		
Smaller group	D	P-value
One Urn Winners	0.1641	0.095
Two Urn Winners	0	1
Combined K-S:	0.1641	0.19
One Urn Vs Two Urns: Non-Winners		
One Urn Non-Winenrs	0.0501	0.351
Two Urn Non-Winners	-0.0173	0.882
Combined K-S:	0.0501	0.672
Two Urns Auction Winners vs Non-Winners		
Smaller group	D	P-value
Non-Winners	0.164	0.023
Winners	-0.0073	0.993
Combined K-S:	0.164	0.045
One Urn Auction Winners vs Non-Winners		
Smaller group	D	P-value
Non-Winners	0.0729	0.454
Winners	-0.0636	0.549
Combined K-S:	0.0729	0.825

B.7 Instructions

(First set: 2-Urn treatment followed by 1-Urn treatment; Second set: 1-Urn treatment followed by 2-Urn treatment.)

Information, Risk and Prices
Instructions

Summary

You will be participating in a series of online trading games. You will receive a fixed sign-up reward and additional pay based on your performance. The goal of this experiment is to determine how information affects prices of financial assets.

1 Setting

This experiment is made of separate trading games called ***periods***. At the beginning of each ***period***, you will be given ***stocks and cash***. The market will open and you will be free to trade stocks. You may buy stocks with cash and you will get cash if you sell stocks. At the end of each period, the stocks pay ***dividends*** as specified below and then stop existing.

Your ***period earnings*** have 2 components: (i) ***dividends*** from stocks, plus (ii) your ***change in cash***. Both are determined by your holdings at the end of the period when the market closes. At the end of the experiment, ***2 periods*** will be ***chosen at random*** and your earnings from those periods are yours to keep. In addition, you will receive a fixed sign-up reward of $20 Australian Dollars. If your earnings are negative in the chosen periods you will still receive the minimum payment of $40 AUD. The maximum you can earn in the experiment is $60 AUD.

During the experiment, accounting is done in ***Experimental Dollars***. At the end of the experiment, the amount you earned in the randomly selected periods will be converted to Australian Dollars at an exchange rate of 3:1.

2 The Bucket Problem

At the beginning of each period, ***a bucket*** will be filled with a large number of ***randomly drawn black and white balls***. All these balls will then be distributed among all participants. There will be enough balls to give each of you exactly ***2 balls***. The balls you receive are ***only visible to you***. At the end of each period, all balls are removed and the bucket will be refilled with ***new balls***. The filling of the bucket and the distribution of balls is illustrated below:

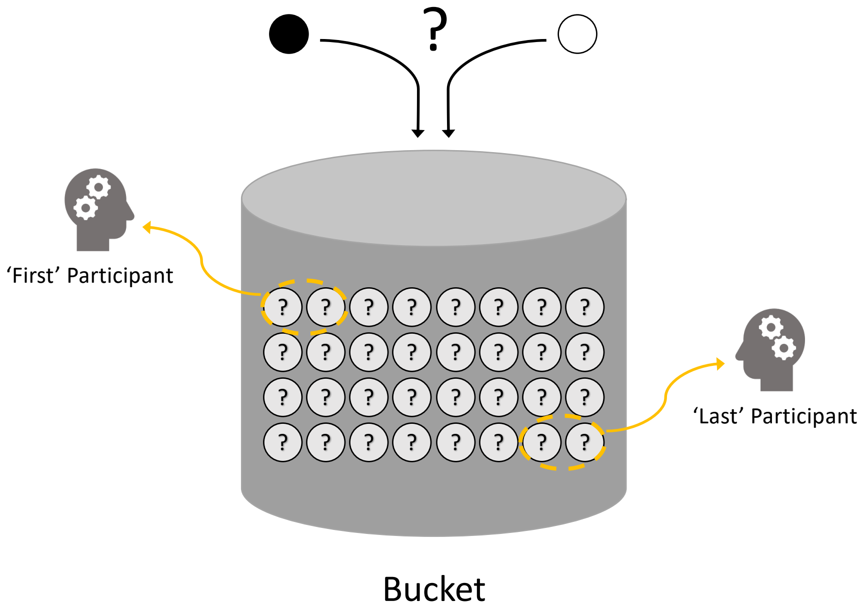

3 Stocks

In each period, you will be endowed with units of 2 stocks: ***Stock Black*** and ***Stock White***. The dividends these stocks pay depend on how many black balls were distributed during that period. The dividends per unit of stock are as follows:

- 1 unit of ***Stock Black*** pays ***the proportion of black balls times $10***. For example, if we distributed 26 black balls out of a total of 32 balls, then 81.3% are black, so 1 unit of Stock Black pays $8.13.

- 1 unit of ***Stock White*** pays ***the proportion of white balls times $10***. For example, if 81.3% of all balls are black, then 18.7% must be white and 1 unit of Stock White pays $1.87.

You will be able to **freely** trade ***Stock Black***. However, you ***cannot*** trade ***Stock White***. If you start holding only Stock White, you can manage the risk of receiving low dividends by buying Stock Black. This works because the 2 stocks are complementary! Whenever Stock Black pays a low dividend, Stock White pays a high dividend; whenever Stock White pays low, Stock Black pays high. By holding an equal number of Stock Black and Stock White, you will know your dividends for sure. When you start out with only Stock Black, the only way to reduce the risk of receiving low dividends is to sell Stock Black.

If you hold zero Stock Black at any moment, you can still sell Stock Black (and get cash). This is called ***short selling***. If you short sell 1 unit of Stock Black, you get to keep the sales price. However, ***the proportion of black balls times $10*** will be ***subtracted*** from your earnings at the end of the period. You will be able to short sell a maximum of 8 shares of Stock Black each period.

4 Information

At the beginning of each period, the bucket is refilled with new balls. If you receive 2 black balls, this of course suggests that there are more black than white balls in the bucket. Alternatively, if you receive 2 white balls, this suggests that there are more white than black balls in the bucket.

Each trade involves 2 participants. One can imagine that the information these participants have will be reflected in the price at which they are willing to trade. For example, when 2 people meet they have information about 4 balls, and if 3 of them are black, then Stock Black is expected to trade around $7.50. Higher prices may thus reveal evidence in favour of more black than white balls. In contrast, lower prices suggest that trading parties have samples with less black than white balls. Subsequent traders will take this into account, combining the information gained from the level of prior trades and their own sample. If they too have samples with more black balls, prices will increase even more.

5 Information Auctions

During each period, you will repeatedly be able to buy additional information which will tell you how close the trading price of Stock Black is with respect to the collated sample (that is the total number of balls collected from all traders in the room). Of course, this information is not free. To receive the information, you have to bid for it. The bidding will happen anonymously through a sealed auction, so that you will not see how much others are willing to pay. Whoever sends ***the highest bid***, receives the information. However, the winner only has to pay ***the second highest bid***. This is called a "second-price auction". It is not in your best interest to submit multiple bids, since you may pay the price of your lower bid rather than someone else's even lower price. Your holdings will not be affected by the auction, except that your cash will be reduced by the second highest bid in case you should win the auction.We decide randomly to determine who wins the auction if there is more than 1 bidder with the same winning price. The maximum you can bid is $20, which represents a significant proportion of your potential earnings.

The winner of the auction (the highest bidder) receives a "***too low***", "***too high***" or "***within $1***" answer. Whenever the last trade price for Stock Black ***is within $1*** (±$1) of the price implied by ***the collection of all balls***, the winner's screen will show this symbol "✓" and state "***within $1***". In contrast, whenever the last trade price for Stock Black ***is not within $1*** (±$1) of this fair price, the winner's screen will show this symbol "✗" and either state "***too low***", if the last trade price ***is below*** the fair price, or "***too high***", if it ***is above***. This is illustrated in Subfigures (c) and (d) on the last page of these instructions. Everyone will see the winning bid, the second highest bid, and the last trade price . However, ***only the winner*** will see the additional information (compare Subfigures (c) and (d) to Subfigure (b)).

In each ***period***, the additional information about the accuracy of the market price can be purchased ***3 times***. After 90 seconds of trading and simultaneous bidding, a ***round*** ends and the information is sent to the highest bidder. In summary, there will be 3 ***rounds*** of trading and bidding (90 seconds each, with 30 seconds in between) and 1 final ***round*** of trading only (90 seconds). This is illustrated below:

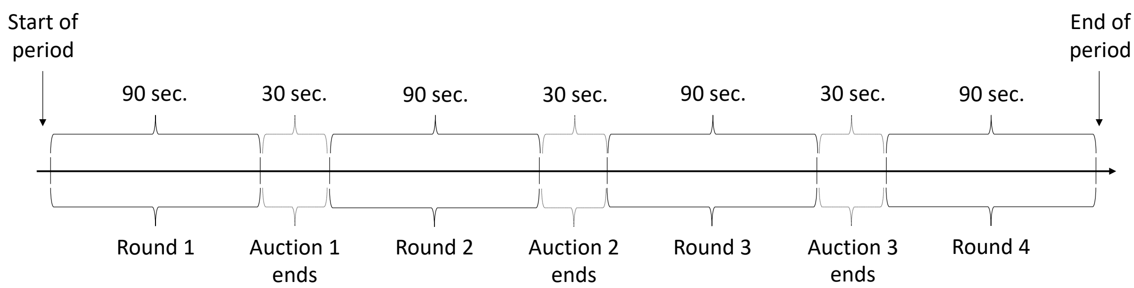

6 Earnings

At the end of each ***period***, that is after 4 ***rounds*** of trading, the actual proportion of balls will be revealed and your earnings will be calculated. Your period earnings depend on the ***change in cash*** from trading and ***dividends earned***. Change in cash is calculated as final cash minus starting cash. Dividends earned depend on the number of units of the dividend paying stock you end up holding. Remember, your actual payment for the experiment will be based on your earnings from only ***2 randomly drawn periods***. Therefore, it is crucial that you pay full attention to every single period.

7 The Trading Platform & Information Screen

Trading takes place through an electronic trading platform called "Flex-E-Markets". In Flex-E-Markets you submit limit orders, which are orders to buy or sell at a price you determine, or, if possible, at any better price. Transactions take place from the moment a buy order with a higher price crosses a sell order with a lower price or vice versa. Orders remain valid until you cancel them or the marketplace pauses or closes. You will be given ample opportunity to train yourself in submitting and canceling orders.

Flex-E-Markets also provides access to the information auction. There, you can submit your bid, if you want to participate in the auction. To do so, select the auctioneer, ***trader:M000***, from the list of "targets" in the order form. Do not submit orders to any other target; you will be penalised for doing so!

You can access Flex-E-Markets as follows: use your login information sheet (e-mail and password) and log into the account named ***cogent-ace*** at ***adhocmarkets.com***. You should then navigate to the marketplace called ***"Bucket-Market-4"***.

In addition to Flex-E-Markets, you are provided with a website that displays all additional information, including the outcome from the latest auction. To access it for the practice round, open a window in ***Google Chrome***. Once the window is open, go to ***etools.bmmlab.org*** and log in with your e-mail and password. After the practice round, you ***have to log out*** of this website by clicking the log out button at the bottom of the page. For the main experiment, you then need to sign in again with a ***new window***, otherwise you will not receive any information.

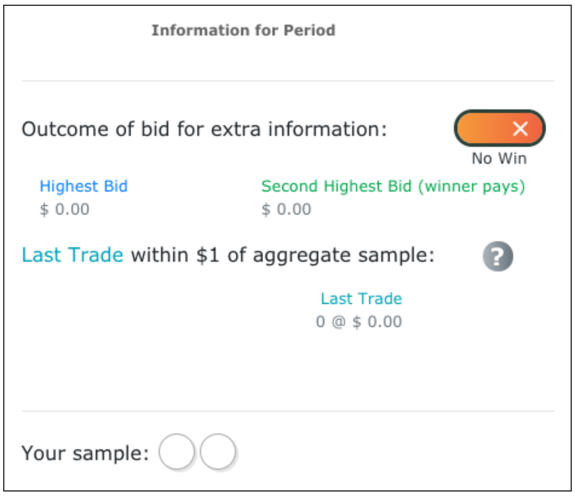

(a) Information Auction: Start

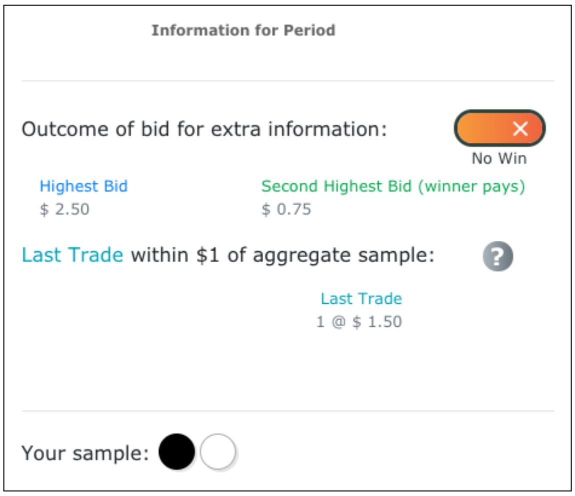

(b) Information Auction: Lose

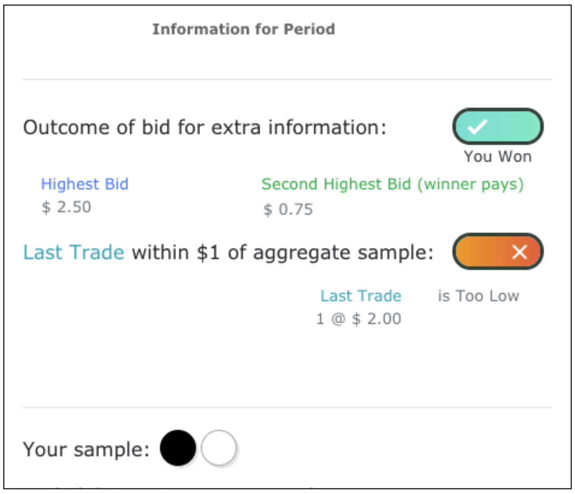

(c) Information Auction: Win/Price too low

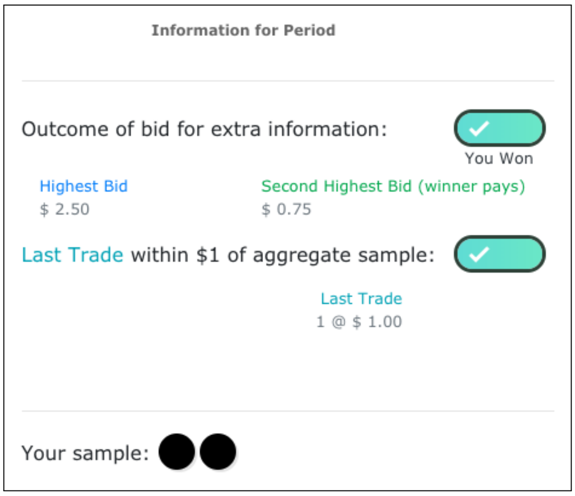

(d) Information Auction: Win/Price within $1

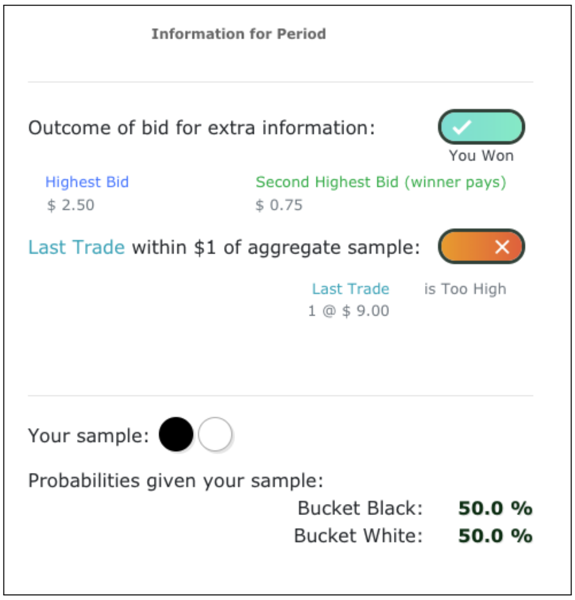

(e) Information Auction: Win/Price too high

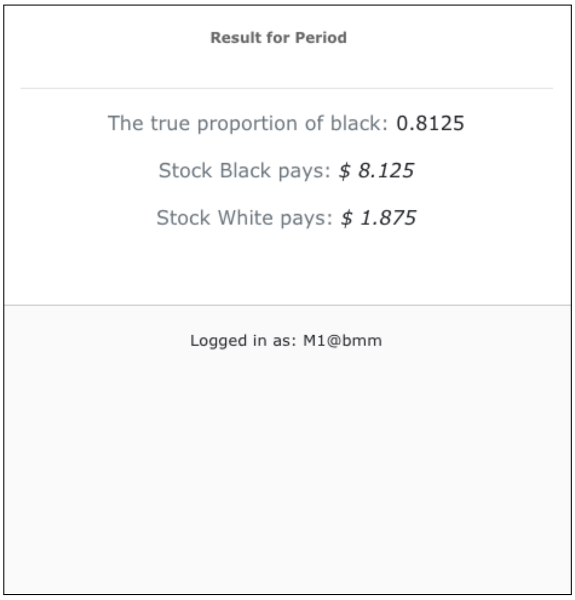

(f) End of Period: Final Dividends

Information, Risk and Prices
Instructions - Part 2

1 New Game

Now, instead of 1 bucket, there will be ***2 buckets***. At the beginning of each period, the computer chooses ***1 out of 2*** buckets with ***equal chance***: either ***Bucket Black or Bucket White*** with 50% probability each. The choice of this bucket will determine the dividends of the stocks. The chosen bucket is ***the same for everybody*** and is not revealed until the end of the period. In order to earn money, you should try to determine which bucket has been chosen. Both buckets contain 10 balls. ***Bucket Black*** contains ***6 black*** and ***4 white*** balls (majority black), whereas ***Bucket White*** contains ***4 black*** and ***6 white*** balls (majority white). This is illustrated below:

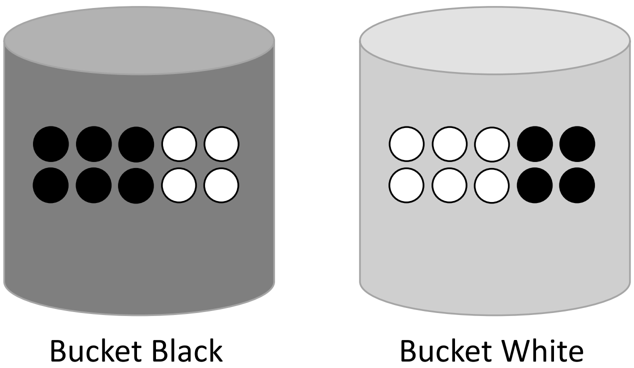

As before, in each period, every participant receives a sample of ***2 balls.*** These balls are ***drawn with replacement*** from the chosen bucket. The balls are drawn independently across participants. ***Hence, your sample is based on a different draw than that of others***.

2 New Payoffs

Again, in each period, you will be endowed with units of 2 stocks: ***Stock Black*** and ***Stock White***. The dividends these stocks pay now depend on which of the 2 buckets was chosen for the period. The dividends per unit of stock are as follows:

- 1 unit of ***Stock Black*** pays ***$10*** if ***Bucket Black*** has been chosen and nothing if Bucket White has been chosen.
- 1 unit of ***Stock White*** pays ***$10*** if ***Bucket White*** has been chosen and nothing if Bucket Black has been chosen.

3 Information

Bucket Black has more black balls than Bucket White. If you have 2 black balls this suggests that Bucket Black is more likely the chosen bucket. Alternatively, if you have 2 white balls this suggests that Bucket White is the chosen bucket.

Each trade involves 2 participants. One can imagine that the information these participants have will be reflected in the price at which they are willing to trade. For example, when 2 participants meet they have information about 4 balls, and if 3 of them are black (2 more black than white balls) there is a 69% chance of Bucket Black. Therefore, both participants are willing to trade Stock Black at a higher price. Higher prices may thus reveal evidence in favour of more black than white balls and thus in favour of Bucket Black. In contrast, lower prices suggest that trading parties have samples with less black than white balls. Subsequent traders will take this into account, combining the information gained from the level of prior trades and their own sample. If they too have samples with more black balls, prices will increase even more.

The ***most informative*** estimate of the probability that either Bucket Black or Bucket White is the chosen bucket is based on ***the total number of black and white balls*** distributed across ***all participants***. For instance, if there are 4 more white than black balls, Bucket White is likely to have been chosen with 84% probability. If there are 6 (8) more white than black balls, Bucket White is likely to have been chosen with 92% (96%) probability. As you can see, the probabilities of Bucket White or Bucket Black are very sensitive to the difference between the total number of white and black balls. In other words, the probabilities become quite extreme with only small differences between the aggregate number of white and black balls.

All other elements of the game remain the same. At the end of each period, it will be announced which of the 2 bucket has been chosen (see below).

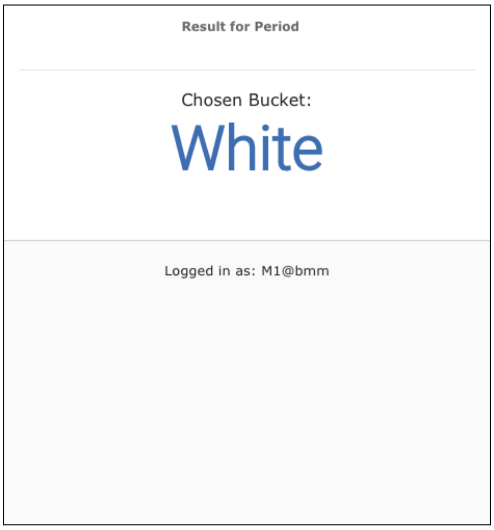

End of Period: Chosen Bucket

Information, Risk and Prices
Instructions

Summary

You will be participating in a series of online trading games. You will receive a fixed sign-up reward and additional pay based on your performance. The goal of this experiment is to determine how information affects prices of financial assets.

1 Setting

This experiment is made of separate trading games called ***periods***. At the beginning of each ***period***, you will be given ***stocks and cash***. The market will open and you will be free to trade stocks. You may buy stocks with cash and you will get cash if you sell stocks. At the end of each period, the stocks pay ***dividends*** as specified below and then stop existing.

Your ***period earnings*** have 2 components: (i) ***dividends*** from stocks, plus (ii) your ***change in cash***. Both are determined by your holdings at the end of the period when the market closes. At the end of the experiment, ***2 periods*** will be ***chosen at random*** and your earnings from those periods are yours to keep. In addition, you will receive a fixed sign-up reward of $20 Australian Dollars. If your earnings are negative in the chosen periods you will still receive the minimum payment of $40 AUD. The maximum you can earn in the experiment is $60 AUD.

During the experiment, accounting is done in ***Experimental Dollars***. At the end of the experiment, the amount you earned in the randomly selected periods will be converted to Australian Dollars at an exchange rate of 3:1.

2 The Bucket Problem

At the beginning of each period, the computer chooses ***1 out of 2*** buckets with ***equal chance***: either ***Bucket Black or Bucket White*** with 50% probability each. The choice of this bucket will determine the dividends of the stocks. The chosen bucket is ***the same for everybody*** and is not revealed until the end of the period (see Subfigure (f) on the last page of these instructions). In order to earn money, you should try to determine which bucket has been chosen.

Both buckets contain 10 balls. ***Bucket Black*** contains ***6 black*** and ***4 white*** balls (majority black), whereas ***Bucket White*** contains ***4 black*** and ***6 white*** balls (majority white). This is illustrated below:

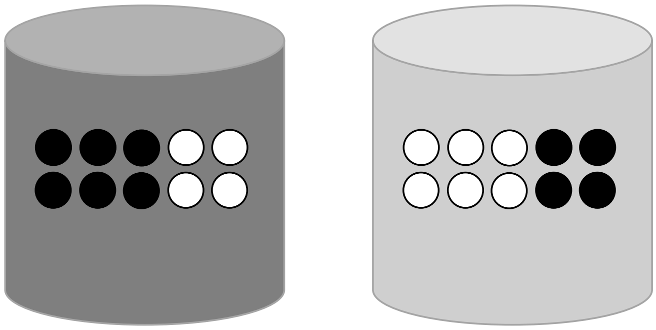

In each period, before trading begins, every participant receives a sample of ***2 balls.*** These balls are ***drawn with replacement*** from the chosen bucket. The balls are drawn independently across participants. ***Hence, your sample is based on a different draw than that of others***.

3 Stocks

In each period, you will be endowed with units of 2 stocks: ***Stock Black*** and ***Stock White***. The dividends these stocks pay depend on which of the 2 buckets was chosen for the period. The dividends per unit of stock are as follows:

- 1 unit of ***Stock Black*** pays ***$10*** if ***Bucket Black*** has been chosen and nothing if Bucket White has been chosen.
- 1 unit of ***Stock White*** pays ***$10*** if ***Bucket White*** has been chosen and nothing if Bucket Black has been chosen.

You will be able to **freely** trade ***Stock Black***. However, you ***cannot*** trade ***Stock White***. If you start holding only Stock White, you can manage the risk of not receiving any dividends by buying Stock Black. This works because the 2 stocks are complementary! Whenever Stock Black pays a dividend, Stock White does not; whenever Stock White pays, Stock Black does not. By holding an equal number of Stock Black and Stock White, you will know your dividends for sure. When you start out with only Stock Black, the only way to reduce the risk of not receiving any dividends is to sell Stock Black.

If you hold zero Stock Black at any moment, you can still sell Stock Black (and get cash). This is called ***short selling***. If you short sell 1 unit of Stock Black, you get to keep the sales price. However, if it turns out that Bucket Black was indeed the chosen bucket, $***10*** will be ***subtracted*** from your earnings at the end of the period. You will be able to short sell a maximum of 8 shares of Stock Black each period.

4 Information

Bucket Black has more black balls than Bucket White. If you have 2 black balls this suggests that Bucket Black is more likely the chosen bucket. Alternatively, if you have 2 white balls this suggests that Bucket White is the chosen bucket.

Each trade involves 2 participants. One can imagine that the information these participants have will be reflected in the price at which they are willing to trade. For example, when 2 participants meet they have information about 4 balls, and if 3 of them are black (2 more black than white balls) there is a 69% chance of Bucket Black. Therefore, both participants are willing to trade Stock Black at a higher price. Higher prices may thus reveal evidence in favour of more black than white balls and thus in favour of Bucket Black. In contrast, lower prices suggest that trading parties have samples with less black than white balls. Subsequent traders will take this into account, combining the information gained from the level of prior trades and their own sample. If they too have samples with more black balls, prices will increase even more.

The ***most informative*** estimate of the probability that either Bucket Black or Bucket White is the chosen bucket is based on ***the total number of black and white balls*** distributed across ***all participants***. For instance, if there are 4 more white than black balls, Bucket White is likely to have been chosen with 84% probability. If there are 6 (8) more white than black balls, Bucket White is likely to have been chosen with 92% (96%) probability. As you can see, the probabilities of Bucket White or Bucket Black are very sensitive to the difference between the total number of white and black balls. In other words, the probabilities become quite extreme with only small differences between the aggregate number of white and black balls.

5 Information Auctions

During each period, you will repeatedly be able to buy additional information which will tell you how close the trading price of Stock Black is with respect to the collated sample (that is the total number of balls collected from all traders in the room). Of course, this information is not free. To receive the information, you have to bid for it. The bidding will happen anonymously through a sealed auction, so that you will not see how much others are willing to pay. Whoever sends ***the highest bid***, receives the information. However, the winner only has to pay ***the second highest bid***. This is called a "second-price auction". It is not in your best interest to submit multiple bids, since you may pay the price of your lower bid rather than someone else's even lower price. Your holdings will not be affected by the auction, except that your cash will be reduced by the second highest bid in case you should win the auction. We randomly determine who wins the auction if there is more than 1 bidder with the same winning price. The maximum you can bid is $20, which represents a significant proportion of your potential earnings.

The winner of the auction (the highest bidder) receives a "***too low***", "***too high***" or "***within $1***" answer. Whenever the last trade price for Stock Black ***is within $1*** (±$1) of the price implied by ***the collection of all balls***, the winner's screen will show this symbol "✓" and state "***within $1***". In contrast, whenever the last trade price for Stock Black ***is not within $1*** (±$1) of this fair price, the winner's screen will show this symbol "✗" and either state "***too low***", if the last trade price ***is below*** the fair price, or "***too high***", if it ***is above***. This is illustrated in Subfigures (c) and (d) on the second to last page of these instructions. Everyone will see the winning bid, the second highest bid, and the last trade price . However, ***only the winner*** will see the additional information (compare Subfigures (c) and (d) to Subfigure (b)).

In each ***period***, the additional information about the accuracy of the market price can be purchased ***3 times***. After 90 seconds of trading and simultaneous bidding, a ***round*** ends and the information is sent to the highest bidder. In summary, there will be 3 ***rounds*** of trading and bidding (90 seconds each, with 30 seconds in between) and 1 final ***round*** of trading only (90 seconds). This is illustrated below:

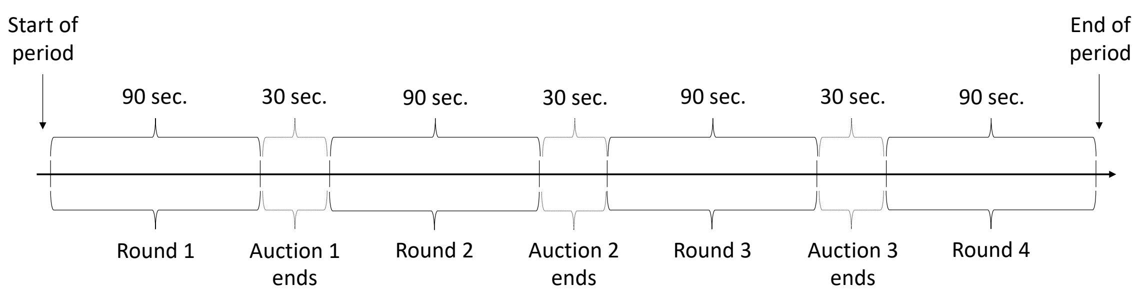

6 Earnings

At the end of each ***period***, that is after 4 ***rounds*** of trading, the chosen bucket will be revealed and your earnings will be calculated. Your period earnings depend on the ***change in cash*** from trading and ***dividends earned***. Change in cash is calculated as final cash minus starting cash. Dividends earned depend on the number of units of the dividend paying stock you end up holding. Remember, your actual payment for the experiment will be based on your earnings from only ***2 randomly drawn periods***. Therefore, it is crucial that you pay full attention to every single period.

7 The Trading Platform & Information Screen

Trading takes place through an electronic trading platform called "Flex-E-Markets". In Flex-E-Markets you submit limit orders, which are orders to buy or sell at a price you determine, or, if possible, at any better price. Transactions take place from the moment a buy order with a higher price crosses a sell order with a lower price or vice versa. Orders remain valid until you cancel them or the marketplace pauses or closes. You will be given ample opportunity to train yourself in submitting and canceling orders.

Flex-E-Markets also provides access to the information auction. There, you can submit your bid, if you want to participate in the auction. To do so, select the auctioneer, ***trader:M000***, from the list of "targets" in the order form. Do not submit orders to any other target; you will be penalised for doing so!

You can access Flex-E-Markets as follows: use your login information sheet (e-mail and password) and log into the account named ***cogent-ace*** at ***adhocmarkets.com***. You should then navigate to the marketplace called ***"Bucket-Market-4"***.

In addition to Flex-E-Markets, you are provided with a website that displays all additional information, including the outcome from the latest auction. To access it for the practice round, open a window in ***Google Chrome***. Once the window is open, go to ***etools.bmmlab.org*** and log in with your e-mail and password. After the practice round, you ***have to log out*** of this website by clicking the log out button at the bottom of the page. For the main experiment, you then need to sign in again with a ***new window***, otherwise you will not receive any information.

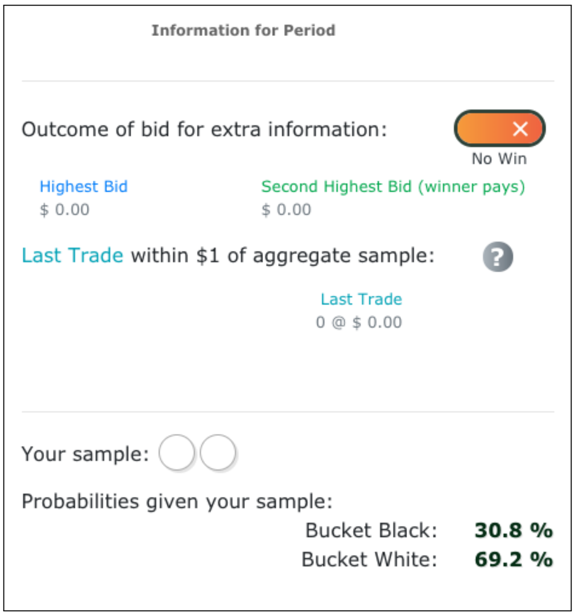

(a) Information Auction: Start

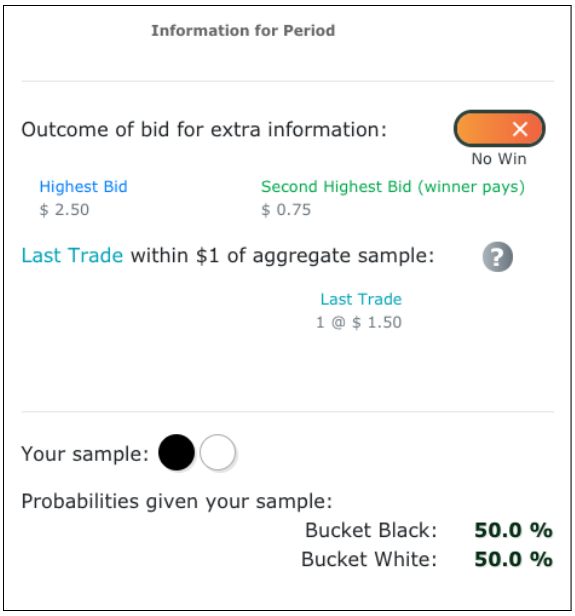

(b) Information Auction: Lose

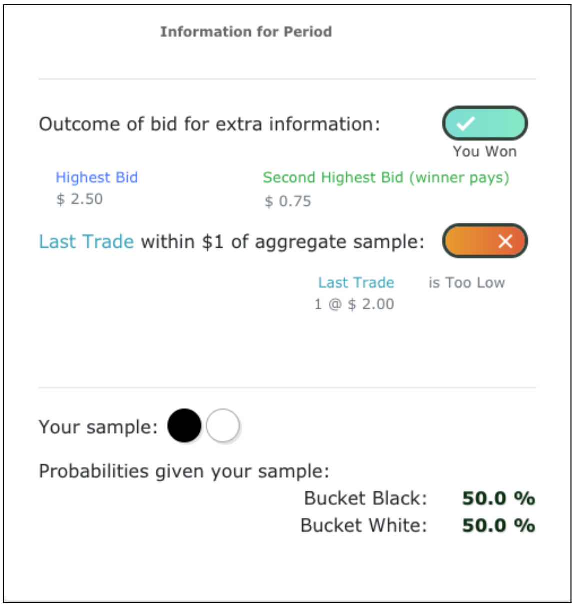

(c) Information Auction: Win/Price too low

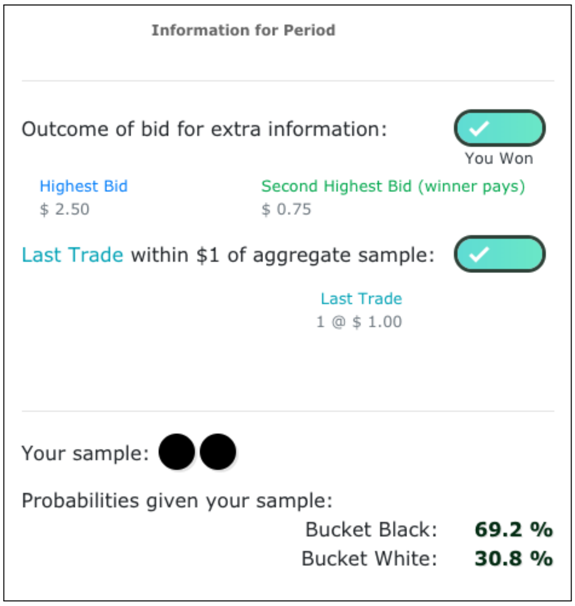

(d) Information Auction: Win/Price within $1

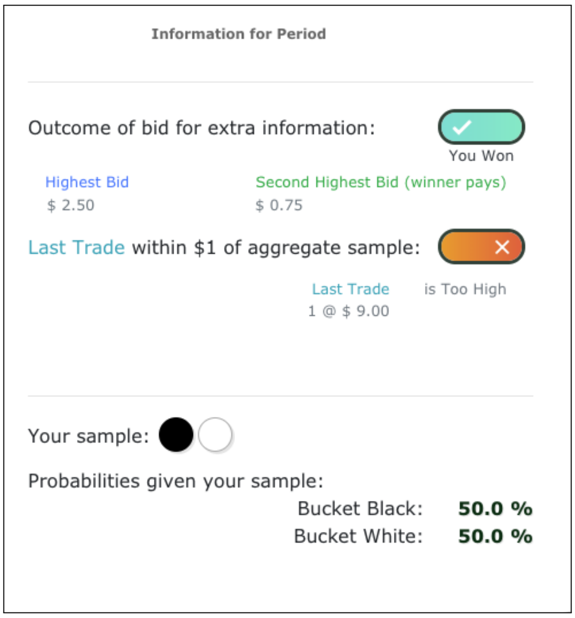

(e) Information Auction: Win/Price too high

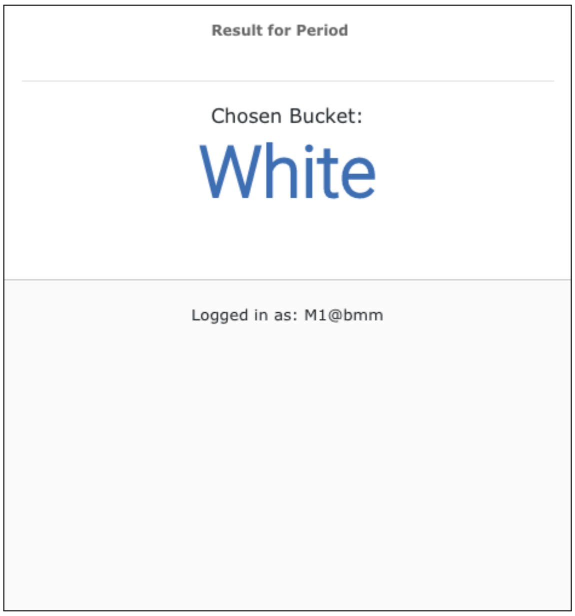

(f) End of Period: Chosen Bucket

Information, Risk and Prices
Instructions - Part 2

1 New Game

Now, instead of 2 buckets, there will be ***only 1 bucket***. This bucket contains ***a large number of randomly drawn balls*** which can be either black or white. All these balls will then be distributed among all participants. There will be enough balls to give each of you exactly ***2 balls***. As before, the balls you receive are ***only visible to you***. At the end of each period, all balls are removed and the bucket will be refilled with ***new balls***. The filling of the bucket and the distribution of balls is illustrated below:

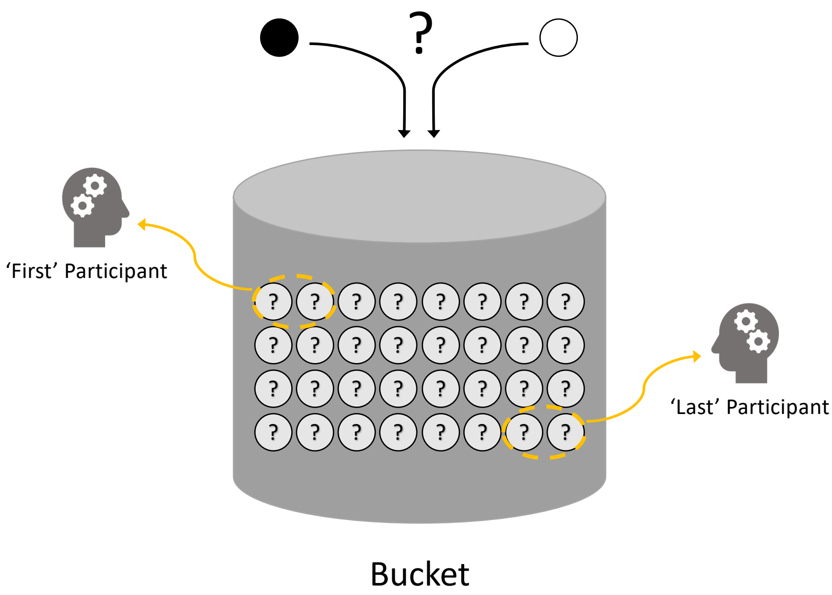

2 New Payoffs

As before, you will be endowed with units of 2 stocks: ***Stock Black*** and ***Stock White***. However, the dividends these stocks pay now depend on the ***proportion*** of black balls distributed among all participants.

- 1 unit of ***Stock Black*** pays ***the proportion of black balls times $10***. For example, if we distributed 26 black balls out of a total of 32 balls, then 81.3% are black, so 1 unit of Stock Black pays $8.13.
- 1 unit of ***Stock White*** pays ***the proportion of white balls times $10***. For example, if 81.3% of all balls are black, then 18.7% must be white and 1 unit of Stock White pays $1.87 (see below).

All other elements of the game remain the same.

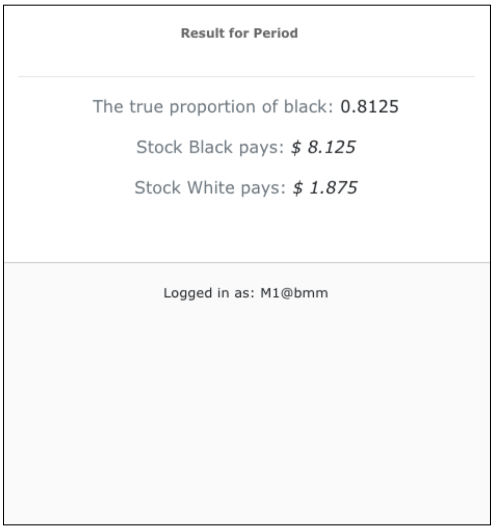

End of Period: Final Dividends

Bibliography

Admati, A.R., 1985, "A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets," *Econometrica* 53, 629-658.

Alfarano, S., Camancho, E., Morone, A. 2015. "Do investors rely too much on public information to be justified by its accuracy? An experimental study," FinMaP-Working Paper, No. 30.

Allen, F., S. Morris and A. Postlewaite, 1993, "Finite Bubbles With Short-Sale Constraints and Aysmmetric Information," *Journal of Economic Theory* 61, 206-229.

Allen, A., S. Morris and H.S. Shin, 2006, "Beauty Contests and Iterated Expectations in Asset Markets," *The Review of Financial Studies* 19, 719-752.

Anderson, L.R. and C.A. Holt, 1997, "Information Cascades in the Laboratory," *The American Economic Review*, Vol. 87, No. 5, 847-862.

Anderson, R.M. and H. Sonnenschein, 1985, "Rational expectations equilibrium with econometric models," *The Review of Economic Studies* 52, 359-369.

Ang, J.S., Shwarz, T., 1985, "Risk Aversion and Information Structure: an Experimental Study of Price Volatility in the Security Markets," *Journal of Finance,* Vol. 40, 824-44.

Asparouhova, E., Bossaerts, P., Eguia, J. and Zame, W., 2015. "Asset pricing and asymmetric reasoning," *Journal of Political Economy* 123, pp.66-122.

Asparouhouva, E., Hertzel, M., Lemmon, M. 2009. "Inference from Streaks in Random Outcomes: Experimental Evidence on Beliefs in Regime Shifting and The Law of Small Numbers," Management Science, Vol. 55, Issue 11, 1766-1782.

Arrow, K.J., 1953, "Le rôle des valeurs boursières pour la répartition la meilleure des risques," *Economètrie*, Paris, Centre National de la Recherche Scientifique, 41-48, 1953.

Arrow, K.J. and G. Debreu, 1954, "Existence of an Equilibrium for a Competitive Economy," *Econometrica*, 22, 265-290.

Aumann, R., 1976, "Agreeing to disagree", *The Annals of Statistics* 4, 1236-1239.

Bachelier, L., 1900, "Théorie de la Speculation," Paris: Gauthier-Villars.

Baillon, Aurélien 2017. "Bayesian markets to elicit private information," Proceedings of the National Academy of Sciences, Volume 114 (30) 7958-7962; DOI: 10.1073/pnas.1703486114.

Banerjee, S., R. Kaniel, and I. Kremer, 2009, "Price Drift as an Outcome of Differences in Higher-Order Beliefs," *The Review of Financial Studies* 22, 3707-3734.

Basak, S., 2000, "A model of dynamic equilibrium asset pricing with heterogeneous beliefs and extraneous risk," *Journal of Economic Dynamics and Control* 24, 63-95.

Basak, S., 2005, "Asset pricing with heterogeneous beliefs", *Journal of Banking and Finance* 29, 2849-2881.

Berg, J., Dickhaut, J. and McCabe, K., 1995. "Trust, reciprocity, and social history." Games and economic behavior, 10(1), 122-142.

Biais, B. and P. Bossaerts, 1998, "Asset Prices and Trading Volume in a Beauty Contest," *Review of Economic Studies* 65, 307-340.

Bossaerts, P., 1997. "Rational Expectations when Priors are Inconsistent," Caltech Social Science Working paper.

Bossaerts, F., Yadav, N., Nasg, C., Todd, T., Rudolf, T., Hutchins, R., Mattingly, K., Ponsonby, A.L. 2019. "When Information Aggregation Works: The Wisdom in the Crowd," Working paper, Dysrupt Labs.

Bossaerts, P., Frydman, C., Ledyard, J. 2014. "The Speed of Information Revelation and Eventual Price Quality in Markets with Insiders: Comparing Two Theories," Review of Finance, Volume 18, Issue 1, 1–22.

Bossaerts, P., Fattinger, F., van den Bogaerde, F. and W. Yang, 2019, "Asset Pricing in a World of Uncertainty," *Working Paper*

Botvinik-Nezer, R., Holzmeister, F., Camerer, C. F., Dreber, A., Huber, J., Johannesson, M., et al. 2019. "Variability in the analysis of a single neuroimaging dataset by many teams." bioRxiv. doi:10.1101/843193.

Camerer, C., Dreber, A., Forsell, E., Teck-Hua, H., Huber, J., Johannesson, M., Kirchler, M., Almenberg, J., Altmejd, A., Taizan, C., Heikensten, E., Holzmeister, F., Imai, T., Isaksson, S., Gideon, N., Pfeiffer, T., Razen, M. and W. Hang, 2016. "Evaluating Replicability of Laboratory Experiments in Economics," *Science,* Vol. 351, Issue. 6280, 1433-1436.

Camerer, C., Wiegelt, K., 1991. "Information Mirages in Experimental Asset Markets," *The Journal of Business,* Vol. 64, No.4, 463-493.

Copeland, T., Friedman, D., 1987. "The Effect of Sequential Information Arrival on Asset Prices: An Experimental Study," The Journal of Finance, Vol. 42, No. 3, 763-797.

Copeland, T., Friedman, D. 1991. " Partial Revelation of Information in Experimental Asset Markets," The Journal of Finance, Vol. 46, No. 1, 265 - 295.

Copeland, T., Friedman, D. 1992. "The Market Value of Information: Some Experimental Results," The Journal of Business, Vol. 65, Issue 2, 241 - 266.

Corgnet, B., DeSantis, M., Porter, D. 2015. "The Distribution of Information and Price Efficiency of Markets," Working Papers 18-09, Chapman University, Economic Science Institute.

Debreu, G., 1959, "Theory of Value," New York: Wiley, Cowles Foundation Monograph 17, Yale University.

Deck, C., Lin, L. and D. Porter, 2010, "Affecting Policy by Manipulating Prediction Markets: Experimental Evidence," Chapman University Digital Commons ESI Working Paper.

DeMarzo, P. and C. Skiadas, 1996, "Aggregation, Determinacy, and Informational Efficiency for a Class of Economies with Asymmetric Information," *Journal of Economic Theory* 80, 123-152.

DeGroot, M.H., 1970, "Optimal Statistical Decisions," New York, McGraw-Hill Book Company.

Diamond, D., Verrecchia, R. 1980. "Information Aggregation In A Noisy Rational Expectations Economy," University of Chicago, Chicago, IL 60637, USA.

Diamond, D.W. and R.E. Verrecchia, 1981, "Information Aggregation in a Noisy Rational Expectations Economy," *Journal of Financial Economics* 9, 221-235.

Dreber, A., Pfeiffer, T., Almenberg, J., Isaksson, S., Wilson, B., Chen, Y., Nosek, B.A. and Johannesson, M. 2015. "Using prediction markets to estimate the reproducibility of scientific research." Proceedings of the National Academy of Sciences, 112 (50), 15343-15347.

Duffie, D. and C.F. Huang, 1985, "Implementing Arrow-Debreu Equilibriy by Continuous Trading of Few Long-lived Securities," *Econometrica* Vol. 53, No. 6, 1337-1356.

Dutta, J. and S. Morris, 1996, "The Revelation of Information and Self-Fulfilling Beliefs," *Journal of Economic Theory* 73, 231-244.

Eyster, E., Rabin, M., Vayanos, D. 2019. "Financial Markets Where Traders Neglect the Informational Content of Prices," The Journal of Finance, Vol. 74, Issue 1, 371-399

Fama, E.F., 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work," The Journal of Finance, Vol. 25, No.2, 383-417

Ferguson, N., 2015, "Investor Information Processing and Trading Volume," *Asia-Pacific Journal of Financial Studies* 44, 322-351.

Ferri, G., Morone, A. 2014. "The Effect of Rating Agencies on Herd Behiour," Journal of Economic Interaction and Coordination, Vol. 9, Issue 1, 107 -127.

Forsythe, R., Lundholm, R., 1990, "Information Aggregation in an Experimental Market," *Econometrica,* 309-347.

Forsythe, R., Palfrey, T.R. and C.R. Plott, 1982, "Asset Valuation in Experimental Markets," *Econometrica* 309-347.

Fudenberg, D. and D. Levine 1983, "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory* 31, 251–268.

Fudenberg, D. and J. Tirole 1991, "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory* 53, 236–260.

Gallmeyer, M. and B. Hollifield, 2008, "An Examination of Heterogeneous Beliefs with a Short-Sale Constraint in a Dynamic Economy," *Review of Finance* 12, 323-364.

Glasner, D., 2020, "Hayek, Hicks, Radner and four equilibrium concepts: Perfect foresight, sequential, temporary, and rational expectations," *The Review of Austrian Economics*, Jan.

Green, J., 1975, "Information, Efficiency and Equilibrium," Harvard Institute of Economic Research, Discussion Paper.

Grossman, S.J., 1976, "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information," *Journal of Finance* 31, 573-585.

Grossman, S. 1977. "The Existence of Futures Markets, Noisy Rational Expectations and Informational Externalities," The Review of Economic Studies, Vol. 44, No.3, 431-449.

Grossman, S. 1978. "Further Results on the Informational Efficiency of Competitive Stock Markets," Journal of Economic Theory, Vol. 18, No. 2, 293-330.

Grossman, S.J. and J.E. Stiglitz, 1980, "On the Impossibility Of Informationally Efficient Markets," *American Economic Review* 70, 393-408.

Grossman, S.J. 1981. "An Introduction to the Theory of Rational Expectations under Asymmetric Information." Review of Economic Studies 48 (4), 541-559.

Giovanni, F., Morone, A. 2008. "The Effect of Rating Agencies on Herd Behaviour," Economics and Econometrics Research Institute, 2008-10-21.

Guesnerie, R., 1992, "An Exploration of the Eductive Justification of the Rational Expectations Equilibrium," *American Economic Review* 82, 1254-1278.

Halim, E., Riyanto, Y.E. and R. Nilanian, 2019, "Costly Information Acqusition, Social Networks, and Asset Prices: Experimental Evidence," *The Journal of Finance,* Vol. 74, Issue. 4, 1975-2010.

Hanson, R., Oprea, R. and D. Porter, 2006, "Information Aggregation and Manipulation in an Experimental Market," *Journal of Economic Behaviour and Organization,* Vol. 60, 449-459.

Harris, M. and A. Raviv, 1993, "The Review of Financial Studies," *The Review of Financial Studies* 6, 473-506.

Harrison, J.M. and M. Kreps, 1978, "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *The Quarterly Journal of Economics* 92, 323-336.

Harsanyi, J.C., 1968, "Games with Incomplete Information Played by 'Bayesian' Players, Part III. The Basic Probability Distribution of the Game," *Management Science* 14, 486-502.

Hayek, F.A., 1937, "Economics and Knowledge," *Economica New Series* 4(13), 33-54.

Hayek, F.A., 1945, "The Use of Knowledge in Society," *The American Economic Review* Vol. 35, No. 4, 519-530.

Hellwig, M.F., 1980, "On the Aggregation of Information in Competitive Markets," *Journal of Economics Theory* 22, 477-498.

Hellwig, M.F., 1982, "Rational Expectations Equilibrium with Conditioning on Past Prices: A Mean-Variance Example," *Journal of Economic Theory* 26, 279-312.

Hey, J., Morone, A. 2004. "Do Markets Drive out Lemmings: Or Vice Vesa?," *Economica*, New Serries, Vol. 71, No. 284, 637-659.

Hirshleifer, J. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity," *The American Economic Review*, Vol. 61, No. 4, 561-574.

Huberman, G., Schwert, G. 1985. "Information Aggregation, Inflation, and the Pricing of Indexed Bonds," *Journal of Political Economy*, vol. 93, issue 1, 92-114.

Kandel, E. and N. Pearson, 1995, "Differential Interpretation of Public Signals and Trade in Speculative Markets," *The Journal of Political Economy* 103, 831-872.

Kandel, S., Ofer, A., Sarig, O. 1993. "Learning from Trading," *The Review of Financial Studies,* Volume 6, Issue 3, August 1993, 507–526.

Keynes, J.M., 1936, *The General Theory of Employment, Interest, and Money*, Macmillan, London.

Kyle, A.S., 1985, "Continuous Auctions and Insider Trading," *Econometrica* 53, 1315-1336.

Lim, K.P. and R. Brooks, 2009, "The Evolution of Stock Market Efficiency Over Time: a Survey of the Empirical Literature," *Journal of Economic Surveys,* Vol. 25, No. 1, 69-108.

Liu, H. and L. Wasserman, 2014, "Statistical Machine Learning," Carnegie Mellon University, unpublished textbook.

Lucas, R., 1972, "Expectations and the Neutrality of Money," *Journal of Economic*

Theory 4, 103-124.

Lucas, R., 1978, "Asset Prices In An Exchange Economy," *Econometrica* 46, 1429-1445.

Mandlebrot, B., 1963, "The Variation of Certain Speculative Prices," *Journal of Business,* 36, 394-419.

Makarov, I. and O. Rytchkov, 2012, "Forecasting the Forecasts of Others: Implications for Asset Pricing," *Journal of Economic Theory* 147, 941-966.

McKelvey, R.D. and Palfrey, T.R., 1995. "Quantal response equilibria for normal form games." *Games and economic behavior,* 10(1), 6-38.

Mertens, J.F. and S. Zamir, 1985, "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory* 14, 1-29.

Merton, R.C., 1982, "On the Microeconomic Theory of Investment Under Uncertainty," *Handbook of Mathematical Economics* North-Holland Publishing, Vol. 2, 923 - 1006.

Milgrom, P. and N. Stokey, 1982, "Information, Trade and Common Knowledge," *Journal of Economic Theory* 26, 17-27.

Morgenstern, O., 1935 [1976], " Perfect Foresight and Economic Equilibrium," *Selected Writings of Oskar Morgenstern* [10], 169-183, Princeton University Press.

Morris, S., 1995, "Justifying The Common Prior Assumption in Economic Theory," *Journal of Economics and Philosophy,*11, 227-253.

Muth, J.F., 1961, "Rational Expectations and the Theory of Price Movements," *Econometrica* 29, 315-335.

Noeth, M., Camerer, C.F., Plott, C.R. and M. Webber, 1999, "Information Aggregation in Experimental Asset Markets: Traps and Misaligned Beliefs," Caltech Social Science Working Paper 1060.

Oprea, R., Porter, D., Hibbert, C., Hanson, R. and D. Tila, 2007, "Can Manipulatiors Mislead Market Observers?," Working Paper.

Osambela, E., 2015, "Differences of Opinion, Endogenous Liquidity, and Asset Prices," *The Review of Financial Studies* 28, 1914-1959.

Page, L. and R.T. Clemen, 2013, "Do Prediction Markets Produce Well Calibrated Forecasts?," *Economic Journal,* 123, 491-513.

Page, L., Siemroth, C. 2017, "An Experimental Analysis of Information Acquisition in Prediction Markets," *Games and Economic Behavior,* 101, 354-378.

Page, L., Siemroth, C. 2019, "How Much Information is Incorporated in Financial Asset Prices? Experimental Evidence," Working Paper.

Plott, C., Sunder, S. 1982. "Efficiency of Experimental Security Markets with Insider

Information: An Application of Rational-Expectations Models," *Journal of Political Economy,* Vol. 90, No. 4, 663-698.

Plott, C., Sunder, S. 1988. "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica,* Vol. 56, No. 5, 1085-1118.

Radner, R., 1967, "Equilibre des marchés à terme et au comptant en cas d'incertitude, " *Cahiers du Séminaire d' Economtrie,*4, 35-62.

Radner, R., 1968, "Competitive Equilibrium Under Uncertainty," *Econometrica* 36, 31-58.

Radner, R., 1972, "Existence of Equilibrium of Plans, Prices, and Price Expectations in a Sequence of Markets," *Econometrica* 40, 289-393.

Radner, R., 1974, "Market Equilibrium and Uncertainty: Concepts and Problems," *Frontiers of Quantitative Economics,* 2, Ch2., 43-105.

Radner, R., 1979, "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica* 47, 655-678.

Radner, R., 1982, "Equilibrium Under Uncertainty," *Handbook of Mathematical Economics* North-Holland Publishing, Vol. 2, 923 - 1006.

Rubinstein, A., 1989, "The Electronic Mail Game: Strategic Behavior Under"Almost Common Knowledge" ," *The American Economic Review*, 79, 385-391.

Samuelson, P.A., 1965, "Proof That Properly Anticipated Prices Fluctuate Randomly," *Industrial Management Review,* 6, 41-9.

Scheinkman, J.A. and W. Xiong, 2003, "Overconfidence and Speculative Bubbles," *Journal of Political Economy* 111, 1183-1219.

Schraeder, S., 2016, "Information Processing and Non-Bayesian Learning in Financial Markets," *Review of Finance*, 823-853.

Sunder, S. 1992. "Market for Information: Experimental Evidence," *Econometrica,* Vol. 60, No. 3, 667-695.

Tirole, J., "On the Possibility of Speculation under Rational Expectations," *Econometrica* 50, 1163-1181.

Townsend, R.M., "Forecasting the Forecasts of Others," *Journal of Political Economy* 91, 546-588.

Wagner, A., Zeckhauser, R., Ziegler, A. 2018. "Company stock price reactions to the 2016 election shock: Trump, taxes, and trade," *Journal of Financial Economics* 130, No. 2, 428-451.

Varian, H. R., 1985, "Differences of Opinion in Financial Markets," CREST working paper, 87–23.