On Information Aggregation in Financial Markets

Inauguraldissertation
zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
der Universität Mannheim

vorgelegt von

Christoph Siemroth

im Herbstsemester 2015

Abteilungssprecher:	Prof. Dr. Carsten Trenkler
Referent:	Prof. Dr. Hans Peter Grüner
Korreferent:	Prof. Klaus Adam, Ph.D.

Tag der Verteidigung: 28. April 2016

Acknowledgements

First and foremost, I would like to thank my supervisor Hans Peter Grüner for his support and guidance in the past five years. We had countless discussions, which even led to a joint research project. His profound impact on me was not limited to matters of research, however. I recall many instances of coaching regarding presentation style, how to sell ideas, or how to deal with other economists. Moreover, I am very grateful for the superb research environment he provided throughout the years, which played no small part in what I achieved in the research phase of my PhD.

I would also like to thank Klaus Adam for his comments and encouragement on the job market paper in particular. I am extremely grateful for his support and advice in the academic job market.

The two junior people that helped and influenced me the most during my time in Mannheim—Philipp Zahn and Pierre Boyer—also deserve my thanks. I had a lot of fun working alongside my office neighbors Johannes Bubeck and Felix Jarman. Their help and advice was very much appreciated (them eating all of my cookies not so much).

While in Mannheim, I have acquired some exceptional co-authors—Michael Gibbs, Hans Peter Grüner, Justus Inhoffen, Susanne Neckermann, Lionel Page, and Philipp Zahn—and I am privileged to work with them. Furthermore, I would like to thank Antonio Cabrales and UCL for their hospitality. I am also very grateful to Astrid Reich and Sylvia Rosenkranz for making academic life so much easier, as well as Sandro Holzheimer and Marion Lehnert from the CDSE for running a smooth operation. Finally, I would like to thank my parents and grandparents for their unconditional support, and all who I could not mention here.

Contents

List of Figures

List of Tables

General Introduction

This dissertation consists of three self-contained chapters, which are ordered from oldest to youngest project. The common theme is information aggregation in financial markets. In chapter 1 (job market paper), traders possess information about a state variable that both affects an asset value and an outcome that a policy maker cares about. Information aggregation in the financial market may not be achieved because the policy reaction to prices—which affects asset values—de facto punishes the informed for trading on information. In chapter 2 (joint work with Hans Peter Grüner), traders possess information about their own preferences regarding novel consumption products, and the return on investment in new firms depends on the future aggregate demand for those firms' products. Information aggregation about the preference distribution may not be achieved due to wealth constraints of consumers, so that the capital allocation reflects the preferences of the wealthy but not necessarily future demand. Chapter 3 (joint work with Lionel Page) is an experimental project, testing whether information acquired by traders is aggregated and incorporated into financial market prices.

Robustness checks and additional analyses of the chapters are delegated to the end of this dissertation. Brief summaries of the chapters follow.

Chapter 1: The Informational Content of Prices When Policy Makers React to Financial Markets

I presented this chapter under its previous title "The Impossibility of Informationally Efficient Markets When Forecasts are Self-Defeating" at several conferences. The starting point of this project is that financial market prices can reveal trader information to outsiders. However, what happens to trader incentives if policy makers start to use the forecasts implicit in asset prices for decision-making? For example, consider what happens if a central bank reacts to inflation forecasts implicit in asset prices. If asset prices reveal a high future inflation rate (absent policy change), but the central bank wants a low inflation rate, then the policy reaction to this information should prevent the forecast implicit in asset prices from coming true. Since traders stake money on their forecasts being right, the policy reaction to asset prices can seriously damage the incentives for traders to reveal the information.

Why should traders bet on high inflation, if they know that the central bank would just react and prevent high inflation in response?

The main objective of the chapter is to identify the situations where prices reveal trader information and policy makers use this information to implement their optimal policy in equilibrium. If there is such an equilibrium, then forward looking traders—who correctly anticipate the policy reaction—still have incentives to reveal their information by trading. The main result in the paper provides a sufficient and necessary condition on the existence of fully revealing equilibria, i.e., the condition tells us when policy makers can and cannot learn from financial market prices with forward looking traders. The model without noise generalizes and unifies several important applications from the literature, and shows there is a common cause to the equilibrium non-existence and non-revelation problems encountered there: self-defeating prophecies. The condition can be used to design or select assets that support information revelation by markets. In the model with noise, I adapt a novel solution approach to similarly find a necessary and sufficient condition for the existence of revealing equilibria. Interestingly, merely introducing noise can solve the problem of self-defeating prophecies in specific situations such as the one considered by Bernanke and Woodford (1997), because noise allows traders to retain some information rents, making it worthwhile for them to trade on information.

Chapter 2: Cutting Out the Middleman: Crowdinvesting, Efficiency, and Inequality

This chapter is joint work with Hans Peter Grüner. We investigate the impact of a mismatch between the wealth distribution and the income distribution of consumers on the funding decision of new firms with novel consumption goods. A mismatch of wealth and income distribution implies that there is a big share of individuals in society that have a positive income, but virtually no wealth, which is an empirical fact in most western countries. Consequently, these individuals consume goods and demand services, yet do not participate in the capital market and therefore do not influence which new technologies or products are funded. We show that this distribution mismatch implies an inefficient capital allocation, because the wealthy (rationally) tend to invest in projects and firms whose product they like, and not necessarily in those that are the most valuable to society. The main result shows that an efficient capital allocation and the welfare optimum is achieved if and only if the wealth and income distribution match in a sense precisely defined in the text.

We model the investment process as direct without financial intermediaries, similar to direct public offerings or crowdinvesting campaigns in practice. Thus, the model provides one of the first welfare analyses of the novel financing method of crowdfunding. In the main extension of the model, we investigate whether institutional investors with sufficient funds can fix the inefficiency due to wealth constraints among consumers. While institutional investors can improve welfare if there is a distribution mismatch, they cannot restore an efficient capital allocation. Thus, the main result continues to hold and an efficient capital allocation is possible if and only if the wealth and income distribution among potential consumers match.

Chapter 3: An Experimental Analysis of Information Acquisition in Prediction Markets

This chapter is joint work with Lionel Page and accepted at Games and Economic Behavior for publication. The main goal of the chapter is to test whether simple asset markets are effective in inducing information acquisition if information is costly. Moreover, we test several theoretical predictions about the drivers of costly information acquisition. Since information sets are not observable in field data, we ran a laboratory experiment to answer these questions.

In the experiment, the risky asset paid out either 0 or 10 at the end of the round, depending on a binary state of nature. The interpretation is that either the incumbent or the challenger is going to win in a political election, and the asset pays dividends if and only if the challenger wins (a neutral framing was used in the experiment). Subjects received two noisy signals about the state of nature for free, and had the chance to acquire more signals at a cost before trading in the asset market. More information made it easier to decide whether the asset is valuable. Trading was organized as a standard continuous double auction, as in many financial and prediction markets.

The standard rational expectations equilibrium predicts no information acquisition and no trade in this setting. However, we observe substantial information acquisition and trade. In fact, the subjects that chose to acquire more information received a lower profit than the subjects that did not acquire information, after accounting for the information costs. This result is very robust, and it does not vanish over the 13 rounds of the experiment.

The remaining findings are largely as expected. Subjects with larger endowment, less informative initial signals, less risk aversion, and less experience in financial markets tend to acquire more information. Moreover, more information in the market decreases the forecast error of market prices. Thus, the individually sub-optimal behavior of acquiring a lot of information is a novel explanation for the good forecasting performance of prediction markets.

Chapter 1

The Informational Content of Prices When Policy Makers React to Financial Markets

1.1 Introduction

Economists have long recognized that markets can aggregate and reveal diverse information among market participants via market prices (Hayek, 1945; Fama, 1970). The recent advance of prediction markets has extended the set of forecasting problems that allow for market-based forecasts to virtually all areas involving uncertain outcomes, such as election outcomes or armed conflicts. Thus, financial markets can potentially be a valuable tool helping policy makers to make decisions by providing information or forecasts.

For example, a central bank may use asset prices to infer information about inflation expectations or future demand shocks, and adapt policy in response (Bernanke and Woodford, 1997). A regulator may learn about the financial health of a bank from bond prices, and use this information for regulatory purposes (e.g., contingent capital with market trigger, Sundaresan and Wang, 2015). Or a company could use internal prediction markets—where asset values depend on the launch date of a new product—to predict whether deadlines can be met, and react if forecasts indicate major delays (e.g., Cowgill and Zitzewitz, 2015).

In all of these examples, one or several agents—which I shall call policy makers—react to the information contained in asset prices, and the reaction in turn affects asset values. This simultaneous feedback from prices to asset values and asset values to prices presents problems both in theoretical and practical terms. In practical terms, the (asset value) forecasts implicit in prices may be falsified by the policy reaction, which can diminish incentives to reveal information if traders are forward looking, thus making market prices less informative and reducing their usefulness for policy makers. A theoretical problem is that such policy maker-trader interactions may not have an equilibrium. The main goal of this paper is to identify if and when it is possible for markets to inform policy makers if traders correctly anticipate the policy maker reaction.[1]

In the first part, I develop a general model of trader and policy maker interaction in competitive financial markets without noise, and derive a sufficient and necessary condition where traders reveal their information by trading, and policy makers use this information for policy purposes in equilibrium. Thus, the condition identifies all situations where we can expect financial markets to work as policy maker tools without compromising the informational content of prices, and all situations where we cannot. Informally, the condition specifies whether the pricing problem is a self-defeating prophecy. With an uninformed policy maker, the condition is very simple and requires invertibility of the expected asset value given optimal policy reaction in a sufficient statistic of trader information. If the condition is not fulfilled, then traders anticipate that informative market prices will trigger a policy maker reaction that leads to trader losses, making market prices less informative. The model unifies several applications from the literature such as Bernanke and Woodford (1997) and Bond et al. (2010), generalizes their results, and shows that there is a common cause to the problem of uninformative prices: self-defeating prophecies. The theoretical results can be used to identify asset payoff functions and policy maker objectives that are more supportive of information revelation. For example, I show that 'deadline securities' in corporate prediction markets, which forecast whether a project will be completed on time, may provide improper incentives to share information about the project if management reacts to these forecasts, and I suggest alternatives with proper incentives.

In the second part of the paper, I develop a model with policy maker-trader interaction where prices are affected by noise. Adapting a new solution approach for noisy rational expectations equilibria (Breon-Drish, 2015), I can solve for equilibria with uninformed policy maker in closed-form even if the functional form of the equilibrium price function is unknown at first. The standard guess-and-verify approach, on the other hand, requires knowledge of the functional form, which depends on the policy reaction to prices. Similar to the no-noise model, I derive a sufficient and necessary condition for the possibility of revealing equilibria in a fairly general class of equilibria that includes and generalizes the usual linear equilibria. I show that noise can solve the problem of self-defeating prophecies in some settings such as the one considered by Bernanke and Woodford (1997): Because noise prevents full revelation of trader information, the informed retain incentives to trade on information, making prices informative while they would not be without noise.

If the policy maker is informed, i.e., receives private signals about the state, then policy maker preferences not only determine whether informative equilibria exist, but also how informative they are. Because asset values are affected by policy decisions, a policy maker

[1] Clearly, an unanticipated policy maker intervention does not diminish the incentives to trade on information. This paper instead focuses on the possibility of policy makers *systematically* using the information revealed by market prices.

with independent information introduces additional 'policy risk' in asset returns beyond the usual risk over asset fundamentals, which influences how aggressively the informed trade on information. Policy risk can induce strategic complementarity leading to multiple equilibria. Interestingly, a higher quality of the policy maker signal can decrease the overall information available to the policy maker, because the informed react by trading less due to policy risk.

I derive measures for price informativeness and policy maker welfare to determine the impact of information contained in financial market prices on real decisions. Comparative statics show that more information in the financial market, less risk aversion among informed traders, and less noise in the market all increase price informativeness and the policy maker welfare gain due to financial market information. The quality of independent policy maker information and prior information decrease price informativeness and the welfare gain. Less extreme policy maker intervention preferences make prices more informative, but also make information less important to the policy maker, so the effect on welfare is ambiguous.

These results can be useful to identify situations where policy makers can use financial market information, and in designing institutions/assets that allow for better information revelation. Moreover, the problem raised here is a more fundamental challenge for the possibility of informationally efficient markets. Financial market prices may not reflect trader information even if all traders have perfect information, are perfectly rational, and obtain their perfect information for free. Similar to the Grossman and Stiglitz (1980)-paradox and its resolution, prices may not be informative in noiseless markets due to self-defeating prophecies, and even if prices are informative in noisy markets, they cannot be fully revealing.

The paper is organized as follows. The next subsections describe two applications in more detail, give a simple example illustrating the problem of self-defeating prophecies, and review the related literature. The first main part analyzes a general model without noise. The second main part analyzes a model with noise. The last section concludes.

1.1.1 Applications

This section briefly discusses two applications that fit the model, which are by no means the only ones. The first is corporate prediction markets (e.g., Wolfers and Zitzewitz, 2004; Cowgill and Zitzewitz, 2015). Corporate prediction markets are designed to elicit information dispersed among employees about business-relevant future outcomes such as whether project deadlines can be met, what next quarter's demand for a product will be, or whether a competitor will enter a particular market segment. Prediction markets typically trade simple assets whose value depends on these outcomes. For example, a project deadline prediction market asset may pay \$1 if and only if the deadline is missed, otherwise it pays \$0. The price of this asset can be interpreted as a market forecast of the probability that the deadline is missed: If many traders think the deadline is not feasible, then they buy the asset, driving the price up.

The information revealed by these markets is most useful if it is used to improve corporate decisions: If a deadline is likely not going to be met, then the company can assign additional resources. If a competitor is about to enter a market segment, then a price reduction might deter him. However, this "policy maker" reaction to market prices is exactly what can create self-defeating prophecies, because it also affects asset values. If traders buy the asset to signal that the deadline cannot be met, and the company reacts by assigning additional resources, then the asset becomes worthless, diminishing the incentives to trade on information in the first place! The results in this paper help to design these markets properly, so that traders have incentives to use their information if they anticipate that the market information is used for real decisions.

One example where these prediction markets affected company policy is mentioned in Cowgill and Zitzewitz (2015): Ford decided against introducing several new products after prediction market forecasts revealed that these would not be popular among consumers. Indeed, improving decisions was the main reason for using these markets: "Ford Motor Company [turned to prediction markets] to improve their ability to make decisions that would be in line with customer interests" (HPC Wire, 2011).

One caveat is that the model uses a competitive equilibrium concept, where traders act as price takers. In very small corporate prediction markets, this may not be realistic: A single trader can affect prices, which possibly introduces incentives for price manipulation. A lot of these corporate markets at Google and Ford are quite large (Cowgill and Zitzewitz, 2015)—at Google, all employees were eligible to participate—so the price taker assumption is plausible in many cases.

A second major application is central bank reaction to market prices. It is well known that central banks monitor asset prices, which can reveal information about inflation expectations or future inflation shocks (e.g., Bernanke and Woodford, 1997). Moreover, a large literature on Taylor rules finds that central banks react to asset prices, housing prices, or oil prices (e.g., Rigobon and Sack, 2003; Castro and Sousa, 2012; L'œillet and Licheron, 2012; Finocchiaro and Heideken, 2013). While these empirical results do not always explain why the central bank reacts to these asset prices, they do establish that prices affect policy decisions. The example in the next subsection illustrates how a self-defeating prophecy can arise in the central bank-trader interaction.

1.1.2 Self-defeating prophecies

How do self-defeating prophecies arise in a financial market context? Consider a very simple example in the spirit of Bernanke and Woodford (1997) for illustration.

Example 1. Suppose the future interest rate π is a function of (random) inflation pressures

θ and interest rate i set by the central bank (CB), with $\pi = \theta - i$. Suppose $\theta \in \{0, 1\}$ with full support and $i \in \{0, 1\}$. The CB is inflation targeting and wants $\pi(\theta, i) = 0\ \forall\theta$. Suppose the CB does not have any information on the realization of θ, while traders know θ. The financial market trades one asset, which is worth 1 if $\pi = 0$ and 0 otherwise. Consider the situation where the realization is $\theta = 1$ and the current policy is $i = 0$, which will remain unless the CB receives new information about θ. Can traders profit from their information about θ, and how would they invest to do so?

Suppose traders buy the asset up to price 1, leading to a market price of 1. Then the CB infers from asset prices that the target rate is reached ($\pi = 0$), and does not change policy.[2] Without the policy change, $\pi = \theta - i = 1$, hence the target rate is missed, the asset is worthless, and the traders lose everything they spent buying the asset. Clearly, this is not a behavior that forward looking traders would engage in. Now, instead suppose traders sell the asset at any positive price, leading to a market price of 0. Then the CB infers from asset prices that there are strong inflation pressures $\theta = 1$, hence the CB changes policy to $i = 1$. The target rate is therefore reached, $\pi = 0$, assets have value 1, and since traders sold the asset below value, they again lose money.

It is easy to see that there exists no price in the example that equals the eventual asset value, because the CB reaction leads to value 0 if the price is 1 and to value 1 if the price is 0; the pricing problem is a self-defeating prophecy. This problem does not occur in standard models where the asset value is exogenously fixed. $\square$

1.1.3 Related literature

The paper probably closest to the first part is Bond et al. (2010). The authors consider the problem where a board of directors has no or only imperfect information about the quality of their agent, the company CEO, whereas traders have perfect information. A low quality CEO reduces the firm value, hence should be replaced to increase the firm value, whereas medium and high quality CEOs should not be replaced, since the intervention is costly. In this setting, there is a difficulty in inferring CEO quality from the company stock price, which is a function of the company value, if traders know that the board might react to it: If traders observe a low CEO quality and trade only at low prices, then the board infers low CEO quality from low prices, fires the CEO, increases the stock value, and effectively punishes traders for revealing the information. But then they have no incentive to reveal the information in the first place. The model without noise in this paper generalizes their setting to more flexible information structures and arbitrary policy maker preferences and policy variable spaces. The main addition provided here is the proof that derives a necessary and sufficient condition for the possibility of revealing equilibria, which also provides the link to

[2] Clearly, the CB expectations about how θ maps into prices is endogenous in equilibrium. In this example, the equilibrium candidate is that traders buy if $\theta = 1$, and sell otherwise, which generates the expectations described.

other applications describing similar problems. Moreover, the model with noise introduced here shows that similar problems can occur even if prices are not fully revealing.

Bernanke and Woodford (1997), in an extension of the Woodford (1994) model, consider a central bank (CB) that attempts to infer a state variable θ from private forecasts or forecasts implicit in asset prices to reach a constant inflation target. Forecasters directly observe θ, the CB does not. In their static model, there is no rational expectations equilibrium that fully reveals θ to the CB. This follows from the impossibility that a forecast simultaneously reveals the state and correctly forecasts inflation, which will not depend on θ if the state is revealed to the CB. They show, moreover, that no equilibrium exists in their setting. Again, their application is an example of a self-defeating prophecy. Their static model is a special case of the model without noise in this paper, and the model with noise in this paper shows that their conclusions do not carry over to a noisy setting: With noise, an inflation targeting central bank can learn from prices in equilibrium.

Birchler and Facchinetti (2007) address a similar problem in banking supervision, and give a nice description of self-defeating prophecies and the "double endogeneity" problem of asset values affecting prices and prices affecting asset values via policy. They model a kind of prediction market that predicts bank failure, and the banking supervisor can react to information contained in these asset prices. As in the noiseless models above, full revelation may fail to occur because forward looking traders take into account that the supervisor will react to prices.

The setting considered here is strongly related to the recent literature on contingent capital with market triggers. The idea of contingent capital with market triggers is that information revealed via prices (typically financial health of a bank) is used for real decisions (convert debt into equity, helping struggling banks raise equity), but this in turn affects asset values (returns to equity are diluted). The argument for market triggers is that they provide more current information than accounting measures, which tend to have a considerable lag. The contingent capital models can also suffer from equilibrium non-existence due to self-defeating prophecies (Prescott, 2012; Sundaresan and Wang, 2015). The main difference is that real decisions in these models are not made by a utility maximizing policy maker, but by a mechanical rule that reacts to market prices. Still, many of the problems encountered in the policy maker settings carry over to the contingent capital setting; this applies in particular to the problem of self-defeating prophecies. The main technical difference besides the mechanical decision rule rather than a policy maker in Sundaresan and Wang (2015) is that they consider a continuous time pricing problem. Both papers model a market without noise.

The paper probably closest to the noisy financial market model in the second part is Bond and Goldstein (2015). They also consider a CARA-normal noisy REE model where traders have information about a state θ that the government would like to have, and the government action affects asset values. However, in their model, the state does not directly affect asset values, only indirectly via the government action. Consequently, given their linear policy rule, even if there were no noise, there is no possibility of self-defeating prophecies in their model. Unlike here, Bond and Goldstein (2015) do not analyze the case of an uninformed government. Thus, they only consider linear policy rules whereas this paper also allows for nonlinear policy rules if the policy maker is uninformed. An interesting difference is one of their policy conclusions on transparency: In their model, the government should not disclose its information about θ, as it takes away all incentives to trade on information; in my model, since θ also affects the asset values directly, disclosure can help the government to get more information from the market. Bond and Goldstein (2015) also analyze the question whether the government should commit itself to using financial market information more or less, which is not the focus of this paper.

This paper contributes to the growing literature of the real effects of financial markets via an informational channel, which mostly consists of studies without self-defeating prophecies. In most of this literature, the 'real effect' is the financial market information impact on corporate decisions, as in Dow and Gorton (1997); Subrahmanyam and Titman (1999); Goldstein and Guembel (2008); Foucault and Gehrig (2008); Goldstein et al. (2013); Dow et al. (2015); Edmans et al. (2015).

1.2 The model without noise

1.2.1 Set-up

Consider a financial market with a single riskless asset with rate normalized to 1, and a single risky asset. The optimal policy and the risky asset value depend on a state θ, which is the realization of a random variable distributed according to a common prior distribution on support Θ, where Θ contains at least two elements. The policy maker sets policy $i \in I$. The value of the risky asset is a function $a : \Theta \times I \to \mathbb{R}$, determined by state θ and policy i. Consequently, asset values are directly affected only by policy maker actions (for outcome manipulation by traders, see Ottaviani and Sørensen, 2007). Throughout I assume $\Theta, I \subseteq \mathbb{R}$.

The financial market consists of a continuum of risk neutral traders with a common prior. Every trader j receives an informative i.i.d. signal s_j on the realization of the state variable θ, distributed according to density $f(s_j|\theta) \neq f(s_j|\theta') \ \forall \theta \neq \theta' \in \Theta$. Since different trader signal profiles can contain the same information, denote the summary statistic of the signal profile by $\boldsymbol{s}$, and the set of all possible unique realizations of $\boldsymbol{s}$ by $\boldsymbol{S}$, so that $\forall \boldsymbol{s} \neq \boldsymbol{s}' \in \boldsymbol{S} : h(\theta|\boldsymbol{s}) \neq h(\theta|\boldsymbol{s}')$, where h is the conditional probability density function of θ. $\boldsymbol{s}$ is a sufficient statistic for signal profile $\{s_j\}_j$ if and only if $h(\theta|\{s_j\}_j, \boldsymbol{s}) = h(\theta|\boldsymbol{s}) = h(\theta|\{s_j\}_j)$. The following are three examples of commonly used information structures with corresponding summary statistic that are consistent with this setup.

Example 2.

1. Traders receive perfect signals, i.e., $s_j = \theta$ for all j, as for example in Bernanke and Woodford (1997) or Bond et al. (2010). The summary statistic is $\boldsymbol{s} = \theta$.

2. State $\theta \in \{0, 1\}$ is binomially distributed, $s_j \in \{0, 1\}$, and traders receive imperfect signals, i.e., $1 > \Pr(s_j = 1|\theta = 1) = \Pr(s_j = 0|\theta = 0) > 1/2$. The summary statistic is $\boldsymbol{s} = \int s_j dj$.

3. The state space is the entire real line, $\theta \in \mathbb{R}$, $s_j \sim \mathcal{N}(\theta, \sigma^2)$, and traders receive imperfect signals, i.e., $\sigma^2 > 0$. The summary statistic is $\boldsymbol{s} = \int s_j dj$.

Let $p(\boldsymbol{s}) : \boldsymbol{S} \to A$ be the price function mapping trader information $\boldsymbol{s}$ into an asset price. For example, $a(\theta, i)$ may represent the value of a company, while $p(\boldsymbol{s})$ is the price of the publicly traded company stock. The timing of decisions is illustrated in Figure 2.1: first, trading among all j leads to a market price $p(\boldsymbol{s})$, then, observing the price, the policy maker sets i. The results would be the same for simultaneous trading and policy-making, since policy can condition on the price but traders (prices) cannot condition on policy. In order to keep the analysis focused, I only consider the market aggregate (market price function) and not individual trader strategies for now.[3]

The policy maker receives an imperfect signal $s_p \in S_p$ on the realization of θ. In a special case the policy maker receives a completely uninformative signal, i.e., $|S_p| = 1$, which is equivalent to no signal. The utility function u represents the rational preference ordering of the policy maker over the tuple (θ, i). Thus, if trader information $\boldsymbol{s}$ were known to the policy maker, she would choose policy

$$i(\boldsymbol{s}) \in \arg\max_i \mathbb{E}[u(\theta, i)|\boldsymbol{s}, s_p].$$

To simplify the exposition, I will assume throughout this paper that i exists and is single-valued. Results can be adapted for multiple solutions and mixed policy strategies in a straightforward manner.

In a backward-induction-like step, define $v(\boldsymbol{s}) : \boldsymbol{S} \to A$, the expected asset value if $\boldsymbol{s}$ were known to the policy maker, who then implements her optimal policy $i(\boldsymbol{s})$,

$$v(\boldsymbol{s}) := \mathbb{E}[a(\theta, i(\boldsymbol{s}))|\boldsymbol{s}].$$

The objects $i(\boldsymbol{s})$ and $v(\boldsymbol{s})$ are defined *assuming* $\boldsymbol{s}$ is known to the policy maker, even though it is not. The reason is that once prices are fully revealing (see definition 2 below), then $\boldsymbol{s}$ is known to the policy maker. Hence, the policy maker will implement policy $i(\boldsymbol{s})$ leading

[3] Section 1.3, Appendix A.1, and Appendix A.2 provide explicit microfoundations in terms of trader endowments, strategy spaces etc.

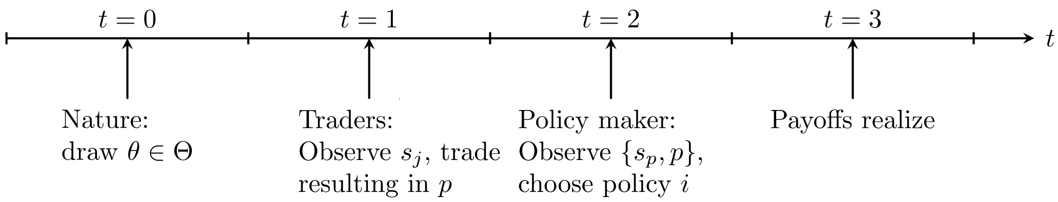

Figure 1.1: Timeline with traders and policy maker.

to expected asset value $v(\boldsymbol{s})$. These are reactions that forward looking traders are going to anticipate if prices are fully revealing.

Given this information structure, policy i cannot be conditioned on $\boldsymbol{s}$ directly, only on $p(\boldsymbol{s})$, so the resulting asset value is $a(\theta, i(p(\boldsymbol{s}), s_p))$. Equilibrium (to be formally defined below) will require policy given beliefs to be optimal. Policy maker behavior, in particular u, is common knowledge. For non-triviality, I assume the optimal policy depends on $\boldsymbol{s}$, so the policy maker is interested in additional information which traders have.

The next definitions introduce two properties of price functions.

Definition 1. *A price function $p(\boldsymbol{s})$ is accurate if and only if*

$$p(\boldsymbol{s}) = \mathbb{E}[a(\theta, i(p, s_p))|p(\boldsymbol{s}) = p, s_j] \text{ for all } j.$$

In words, an accurate price function requires that asset prices equal asset values from the perspective of all traders, where the information set of trader j is both the information contained in prices and his private information s_j. The condition can be interpreted as requiring no systematic mispricing, which becomes clearer in Lemma 1 below.

In the present setting, $\boldsymbol{s}$ can only be indirectly revealed to the policy maker via price $p = p(\boldsymbol{s})$ in combination with the policy maker signal s_p. If the policy maker knows the price function (knows trader behavior), as she does for example in the perfect Bayesian Nash or rational expectations equilibrium concept, then $\boldsymbol{s}$ can always be inferred from the tuple $\{p = p(\boldsymbol{s}), s_p\}$ if the Bayesian posterior probability is positive for at most one $\boldsymbol{s} \in \boldsymbol{S}$.

Definition 2. *A price function $p(\boldsymbol{s})$ and policy maker signals are jointly fully revealing if and only if*

$$|\{t \in \boldsymbol{S} : \Pr(\boldsymbol{s} = t|p(\boldsymbol{s}) = p, s_p) > 0\}| \leq 1 \ \forall s_p, \forall p.$$

According to this definition, *prices need not be fully revealing to traders or outsiders*, only to the policy maker who can combine $p = p(\boldsymbol{s})$ and s_p. If policy maker signals are uninformative, i.e., $|S_p| = 1$, then full revelation reduces to invertibility of the price function in $\boldsymbol{s}$, because prices alone have to reveal $\boldsymbol{s}$. In this case, the information is fully revealed to anybody who observes the price and knows the price function.

The next lemma establishes that if a price function is fully revealing, then accuracy implies $p(\boldsymbol{s}) = v(\boldsymbol{s})$ and vice versa. Hence, if prices are fully revealing, then accurate prices must fully reveal the expected asset value given all trader information $v(\boldsymbol{s})$.

Lemma 1. *A fully revealing price function $p(\boldsymbol{s})$ is accurate if and only if $p(\boldsymbol{s}) = v(\boldsymbol{s})$.*

Proof. Necessity: Accuracy implies $p(\boldsymbol{s}) = v(\boldsymbol{s})$. First, $\mathbb{E}[a(\theta, i(p, s_p))|p(\boldsymbol{s}) = p, s_j] = \mathbb{E}[a(\theta, i(\boldsymbol{s}))|p(\boldsymbol{s}) = p, s_j]$ since prices are fully revealing by assumption, and $\mathbb{E}[a(\theta, i(\boldsymbol{s}))|p(\boldsymbol{s}) = p, s_j] = \mathbb{E}[v(\boldsymbol{s})|p(\boldsymbol{s}) = p, s_j]$ by the law of iterated expectations. Plugging into the definition of accurate prices, $p(\boldsymbol{s}) = \mathbb{E}[v(\boldsymbol{s})|p(\boldsymbol{s}) = p, s_j]$ for all j (and all $\boldsymbol{s} \in \boldsymbol{S}$), and taking the conditional expectation on both sides, $\mathbb{E}[p(\boldsymbol{s})|\boldsymbol{s}] = \mathbb{E}[\mathbb{E}[v(\boldsymbol{s})|p(\boldsymbol{s}) = p, s_j]|\boldsymbol{s}]$. Again using iterated expectations yields $p(\boldsymbol{s}) = v(\boldsymbol{s})$.

Sufficiency: $p(\boldsymbol{s}) = v(\boldsymbol{s})$ implies accurate prices. This immediately follows:

$$\mathbb{E}[a(\theta, i(p, s_p))|p(\boldsymbol{s}) = p, s_j] = \mathbb{E}[a(\theta, i(p, s_p))|v(\boldsymbol{s}) = p, s_j] = v(\boldsymbol{s}) = p(\boldsymbol{s}) \text{ for all } j. \qquad \square$$

Finally, we are going to need a definition to characterize the policy maker signal structures that allow for full revelation. First, define the full inverse of $v(\boldsymbol{s})$, i.e., the set of $\boldsymbol{s} \in \boldsymbol{S}$ for which the expected asset value equals p if $\boldsymbol{s}$ is known to the policy maker as

$$v^{-1}(p) := \{\boldsymbol{s} \in \boldsymbol{S} : v(\boldsymbol{s}) = p\},$$

and the set of p for which $v^{-1}(p)$ contains more than one distinct element as

$$X := \{p \in \text{Image}(v(\boldsymbol{s})) : |v^{-1}(p)| > 1\}.$$

Thus, set X contains all prices $p = v(\boldsymbol{s})$ (assuming accurate prices) that are consistent with more than one distinct state $\boldsymbol{s} \in \boldsymbol{S}$. If X is empty, then $v(\boldsymbol{s})$ is invertible, i.e., $v^{-1}(p)$ is single-valued for all p. If prices alone do not always reveal $\boldsymbol{s}$, then X is non-empty. To achieve full revelation, the policy maker signal s_p has to discriminate between the possible states $v^{-1}(p)$ that are consistent with the observed price $p \in X$.

The following condition states that if $v(\boldsymbol{s})$ is not invertible, i.e., set $v^{-1}(p)$ contains more than one element for some price p, then the probability of receiving any signal s_p must be zero in all $\boldsymbol{s} \in v^{-1}(p)$ except for at most one.

Condition 1 (**Excluding signal structure in case of non-invertibility of $v(\boldsymbol{s})$**).

$$|X| > 0 \implies \Pr(s_p|\boldsymbol{s} = t) \cdot \Pr(s_p|\boldsymbol{s} = t') = 0 \ \forall s_p \in S_p, \ \forall t \neq t' \in v^{-1}(p), \ \forall p \in X.$$

The condition has similarity to what Cabrales et al. (2014) call 'excluding signal structure,' but it is not identical, because in their meaning it is sufficient for signals s_p to exclude one state, whereas here policy maker signals might have to exclude several.

The condition is fulfilled if $v(\boldsymbol{s})$ is invertible, so that $|X| = 0$. The following example illustrates how full revelation may still occur even if the price function $p(\boldsymbol{s}) = v(\boldsymbol{s})$ is not invertible, because the policy maker can combine the information revealed by prices with her private information to rule out all states but one.

Example 3. Suppose $\theta \in \{1, 2, 3\}$ and traders observe the state, i.e., $s_j = \theta \; \forall j$. The policy maker receives the imperfect signal $s_p = 1$ if $\theta = 1$ and $s_p = 0$ if $\theta \in \{2, 3\}$.

Now suppose that asset values (if the policy maker knew θ) are $v(\theta) = 1$ if $\theta \in \{1, 2\}$ and $v(\theta = 3) = 0$. Consequently, $v(\theta)$ is not invertible and $|X| > 0$, since $v^{-1}(1) = \{1, 2\}$. Thus, imposing accurate prices, observing merely $p = p(\theta) = v(\theta)$ does not reveal θ, since $p(\theta = 1) = p(\theta = 2) = 1$. However, condition 1 still holds, because if the policy maker observes $p = p(\theta) = v(\theta) = 1$, then she can discriminate between the two states consistent with these prices, $\theta \in \{1, 2\}$, using her private information, since $s_p = 1$ if $\theta = 1$ and $s_p = 0$ if $\theta \in \{2, 3\}$. Both prices and signal s_p are necessary for full revelation; neither one on its own is sufficient to reveal θ. □

Bond et al. (2010) provide another example that fulfills condition 1, where $\theta \in \mathbb{R}$, $v(\theta)$ is non-invertible (see Figure 1.2 below), and the policy maker signal rules out all states but one consistent with market prices, since s_p is distributed on a bounded support around θ.

1.2.2 The possibility of information revelation via prices

This section asks if a price function $p(\boldsymbol{s})$ exists which allows for both full revelation and accurate prices. There is no microfoundation for this price function yet, i.e., the section does not explain how the price function arises in some specified trading game or equilibrium concept. This foundation will be provided in subsequent sections. The analysis is separated in this way to highlight that the impossibility of fully revealing and accurate prices does not depend on this microfoundation. Instead, under some conditions it is mathematically impossible to find a price function that is both fully revealing and accurate.

When is it possible for a price function to reveal $\boldsymbol{s}$ to the policy maker (definition 2) *and* price accurately (definition 1) at the same time? Without full revelation, the policy maker has inferior information and may implement suboptimal policies, and without accurate prices, traders might lose money, hence might be better off not trading. Given correct policy maker beliefs about the price function $p(\boldsymbol{s})$, Theorem 2 shows that this is possible if and only if condition 1 is satisfied.

Theorem 2 (**Possibility of full revelation and accurate prices**). *Suppose the policy maker knows function $p(\boldsymbol{s})$ and maximizes expected utility. Then a fully revealing and accurate price function exists if and only if condition 1 holds.*

Proof. See Appendix. □

Theorem 2 characterizes the existence of fully revealing and accurate price functions in terms of policy maker preferences u, policy maker information structures $\Pr(s_p|\theta)$, asset payoff functions $a(\theta, i)$, and trader information structures $f(s_j|\theta)$. If there exists no price function that is both accurate and fully revealing, then such a price function cannot arise in equilibrium no matter the equilibrium concept.

Corollary 3. *Suppose the policy maker knows function $p(\boldsymbol{s})$ and maximizes expected utility. Then a fully revealing and accurate price function exists if $v(\boldsymbol{s})$ is invertible.*

Full revelation and accurate prices are possible either if $v(\boldsymbol{s})$ is invertible (for any policy maker signal structure), or if $v(\boldsymbol{s})$ is not invertible but the policy maker signal structure is excluding in the sense of condition 1. The requirement on the signal structure is rather strong, as it requires that any $\boldsymbol{s}$ which is not ruled out by the "price signal" is ruled out by the private signal of the policy maker s_p. For example, if traders have perfect information ($s_j = \theta\ \forall j$), and if $\theta \in \mathbb{R}$ and $s_p \sim \mathcal{N}(\mu, \sigma^2)$ with $\sigma^2 > 0$ as in many finance models, then $\Pr(s_p|\theta = t) = \phi\left(\frac{s_p - t}{\sigma}\right) > 0$ for all $t \in \mathbb{R}$ and $s_p \in S_p$, i.e., condition 1 is not fulfilled. Hence, full revelation with normally distributed policy maker signals is possible if and only if $v(\theta)$ is invertible, because policy maker signals never rule out any state.

Moreover, in the special case where the policy maker does not receive an informative signal, invertibility of $v(\boldsymbol{s})$ is necessary and sufficient for full revelation and accurate prices, again because policy maker signals never rule out any state.

Corollary 4. *Suppose the policy maker knows function $p(\boldsymbol{s})$, does not receive informative signals ($|S_p| = 1$), and maximizes expected utility. Then a fully revealing and accurate price function exists if and only if $v(\boldsymbol{s})$ is invertible.*

This simple condition—non-invertibility of the 'expected asset value given optimal policy' function $v(\boldsymbol{s})$—explains the general difficulty of finding fully revealing equilibria for most policy maker signal structures. Prices cannot reveal $\boldsymbol{s}$ and at the same time be accurate. Thus, traders have incentives to make the forecasts implicit in their trades less revealing (rather than wrong) or not to trade at all. Mathematically, the result obtains because a price function cannot at the same time be invertible (as required by full revelation) and non-invertible (as required by accuracy, which implies $p(\boldsymbol{s}) = v(\boldsymbol{s})$, if $v(\boldsymbol{s})$ is non-invertible).

Theorem 2 implies that accurate prices and full revelation are possible if condition 1 holds. The price function equal to the expected asset value function given optimal policy is an accurate forecast of the asset value, and at the same time reveals all information to the policy maker that is need to implement this asset value. Hence, condition 1 is also the condition that determines whether the forecasting problem is self-defeating or self-fulfilling.

1.2.3 Examples

In several papers with policy maker-trader interaction from the literature section, traders have perfect information about the state variable θ. It can easily be verified that invertibility

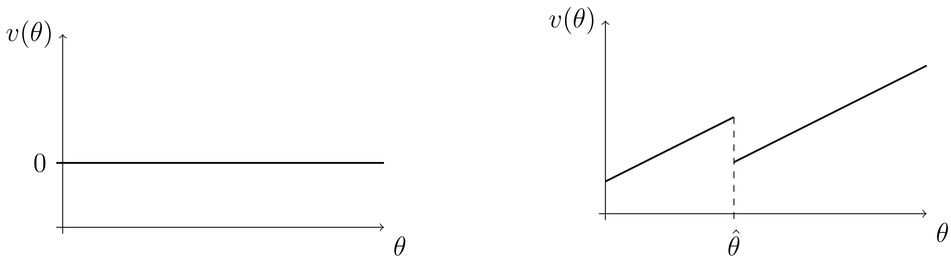

(a) Bernanke and Woodford (1997): $i(\theta) = \arg\min_i \operatorname{Var}(\theta + i)$

(b) Bond et al. (2010), Prescott (2012): intervention below threshold $\hat{\theta}$ increases the asset value

Figure 1.2: Examples for non-invertible asset values at the optimal policy $v(\theta)$ under full information

of $v(\theta)$ is not fulfilled in these papers, see Figure 1.2. For example (adapting their notation), in Bernanke and Woodford (1997)'s static model, the central bank wants to cancel out all variance due to inflation pressures θ, so that the asset value given the optimal policy using trader information is $v(\theta) = \theta + i = c$ for some constant c. In Bond et al. (2010) (similar in Prescott, 2012), the policy variable is binary, and an intervention is value increasing: $a(\theta, i = 1) > a(\theta, i = 0)$. The optimal policy calls for $i = 1$ if and only if $\theta \leq \hat{\theta}$ for some threshold $\hat{\theta}$, hence $v(\theta)$ has a discontinuous downward jump at $\hat{\theta}$, which makes it non-invertible (see Figure 1.2).

Hence, the underlying problem—non-invertibility of $v(\theta)$ or more generally failure of condition 1—is the same in these papers. Despite the same problem, preferences of the policy maker differ, which shows the problem of full revelation with self-defeating prophecies is not due to specific policy maker goals. In Bernanke and Woodford (1997), a central bank wants to minimize the variance of inflation, and in Bond et al. (2010) a board of directors wants to maximize firm value minus intervention cost. In Prescott (2012), the policy is determined by a capital conversion rule.

Despite non-invertibility of $v(\theta)$ in Bond et al. (2010), they show that full revelation may be possible under some conditions, exactly because the informative signal of the policy maker—which is uniformly distributed on a bounded support—excludes all states that are not ruled out by the price (condition 1).

Example 4. One of the most well known examples of an impressive "market forecast" was when the Challenger space shuttle exploded during take-off in 1986. While the investigators took about 4 months to officially announce the defective part and the responsible manufacturing company, the stock price of the responsible company decreased about 12% on the day of the disaster (see Maloney and Mulherin, 2003 for more details). The stock prices of other companies that supplied parts to the space shuttle dropped only 2-3%. Thus, one can argue that the financial market revealed the responsible party almost instantaneously.

To formalize the situation in a very simple model, suppose the stock value of a supplier company is $a(\theta, i) = 1 - i(\theta)$, where $i = 1$ if and only if $\theta = 1$ and $i = 0$ if and only if $\theta = 0$, i.e., $i(\theta) = \mathbf{1}\{\theta = 1\}$. Variable i is the punishment a responsible company will incur, for example due to lawsuits, lost business etc., and $\theta \in \{0, 1\}$ indicates whether the company is responsible for the crash (i.e., its parts malfunctioned).

There are insiders in the market that know θ, i.e., whether the company is responsible. If θ was revealed via prices, then the asset value would be $v(\theta) = a(\theta, i(\theta)) = 1 - i(\theta) = 0$ if the company is responsible ($\theta = 1$), and $v(\theta) = 1$ if the company is not responsible. Thus, $v(\theta)$ is invertible, condition 1 holds, and the forecasting problem of finding the responsible party was not a self-defeating prophecy. Consequently, the anticipated reaction by NASA and authorities did not hamper incentives for traders to reveal their information, which may contribute to explain why the market forecast worked well in this situation.[4] □

1.2.4 Non-existence of fully revealing rational expectations equilibria

This section investigates whether accurate prices and full revelation can occur in a rational expectations equilibrium. That is, under which conditions can the financial market aggregate and reveal private information to the policy maker in equilibrium?

A rational expectations equilibrium in this setting is defined as follows (for a similar definition in the perfect information case, see Bond et al., 2010), where both the $p(\boldsymbol{s})$ and $i(p, s_p)$ function are known in equilibrium.

Definition 3. *A rational expectations equilibrium (REE) consists of*

i. A price function $p(\boldsymbol{s}) = \mathbb{E}_{\theta, s_p}[a(\theta, i(p, s_p))|p(\boldsymbol{s}) = p, s_j]$ *for almost all* j*, and*

ii. an optimal policy function $i(p, s_p)$ *given knowledge of* $p(\boldsymbol{s})$*, i.e.,*

$$i(p, s_p) = \arg\max_i \mathbb{E}_\theta[u(\theta, i)|p(\boldsymbol{s}) = p, s_p].$$

Condition (*i.*) of definition 3 requires that the equilibrium price equals the expected value of the asset given trader information for all traders. This price function is the market clearing outcome of an unmodeled noiseless competitive market with risk neutral traders and rational expectations. See Appendix A.1 for an alternative definition explicitly modeling traders and endowments, yielding the same non-existence result, and DeMarzo and Skiadas (1998, 1999) for more details on risk neutral REE. Condition (*ii.*) requires that the policy maker acts optimally given her information s_p and the information contained in prices.

It turns out that the same condition which determines the existence of a fully revealing and accurate price function is also necessary and sufficient for the existence of a fully re-

[4] While the financial market information is certainly not evidence that holds up in court, the information may steer investigators in the right direction to find admissible evidence, thus affecting real decisions.

vealing REE. This is because Theorem 2 assumed that the policy maker knows $p(\boldsymbol{s})$ and acts optimally given her information, which is now an equilibrium requirement, and full revelation and accurate prices are also required in a fully revealing REE by definition.

Corollary 5. *A fully revealing (definition 2) rational expectations equilibrium exists if and only if condition 1 holds.*

Proof. Condition 1 of definition 3 implies that the REE price function has to be accurate (cf. definition 1). Thus, the result follows from Theorem 2. □

Thus, according to the REE concept, financial markets can both aggregate and reveal all trader information $\boldsymbol{s}$, but only if condition 1 holds. Hence, there are situations where markets cannot be strong-form informationally efficient, i.e., prices do not reflect all trader information. If prices fully revealed trader information, then in at least one state there would be mispricing, which introduces incentives to exploit the mispricing, and consequently traders do not support a fully revealing price function in equilibrium. Even in the most extreme case, where all traders perfectly know the state of the world ($s_j = \theta \ \forall j$) and perfectly know policy maker behavior, prices cannot reflect trader information if condition 1 fails to hold. The problem is not an informational one—traders know everything. Instead, accurate prices and full revelation are mutually exclusive, because the policy maker de facto "prefers to falsify trader forecasts." This result is in strong contrast to the standard models without policy maker, where asset values are exogeneous and existence of fully revealing REE is generic (see for example Radner, 1979 or Allen, 1981).

Appendix A.2 shows that the same condition is necessary and sufficient for full revelation using Perfect Bayesian Nash equilibrium in a continuum economy if traders receive perfectly correlated signals, so the results are relevant beyond the REE concept.

While Corollary 5 explains the non-existence of *fully revealing* equilibria in various applications, the conditions in Corollary 5 and the following Proposition 6 jointly explain the REE non-existence results (fully revealing or otherwise) for example in Bernanke and Woodford (1997) or Prescott (2012).

Proposition 6. *Suppose the policy maker does not receive any signals* ($|S| = 1$) *and traders are perfectly informed* ($s_j = \theta \ \forall j$)*. If* $a(t,i) \neq a(t',i) \ \forall t \neq t' \in \Theta, \ \forall i \in I$*, then there exists no 'not fully revealing' rational expectations equilibrium.*

Proof. Suppose there exists a not fully revealing REE. This implies there exist $t \neq t' \in \Theta$, such that $p(t) = p(t')$. Since the policy maker cannot distinguish the two states, she will take the same action i for $\theta = t$ and $\theta = t'$. But then for at least one of these two states the price must be inaccurate, since $p(t) = p(t')$ and $a(t,i) \neq a(t',i)$ implies $p(\theta) \neq \mathbb{E}[a(\theta,i)|s_j = \theta] = a(\theta,i)$ for at least one $\theta \in \{t,t'\}$, which contradicts this being an REE (condition 1 of definition 3). □

Intuitively, if the asset value changes monotonically in the state θ, and prices are not fully revealing, then at least two states θ and θ' have to be 'bunched together' by the price function, $p(\theta) = p(\theta')$, so the policy maker cannot distinguish θ and θ'. But if traders know θ, then they know that the asset is mispriced for at least one of these two states, hence a not-fully revealing price function cannot be part of an equilibrium.

The last two results imply that the fully revealing REE is the unique REE if condition 1 holds, the policy maker is uninformed, and $a(\theta, i)$ invertible in θ.

1.2.5 Asset design and selection

The results of the previous sections state that accurate and fully revealing price functions and fully revealing equilibria exist with an invertible $v(\boldsymbol{s})$ function. From an asset design or asset selection point of view, we can ask what kind of asset supports full revelation. Formally, differences in assets are captured by the asset payoff function $a(.,.)$ that maps (θ, i) into an asset value.

In this section, I shall consider a class of assets whose value depends on some outcome o, described by a function $o : \Theta \times I \to \mathbb{R}$ that maps the realization of state θ and policy i into a real number. Consequently, the asset value indirectly depends on state θ and policy i to the degree that it influences the outcome $o(\theta, i)$. Thus, I consider the class of assets that can be written as $a(\theta, i) = A(o(\theta, i))$ for some function $A : \mathbb{R} \to \mathbb{R}$.

The A-function may be invertible, or non-invertible such as an Arrow security:

$$A(o(\theta, i)) = \begin{cases} 1 & \text{if } o(\theta, i) \geq T, \\ 0 & \text{if } o(\theta, i) < T, \end{cases}$$

for some threshold value $T \in \mathbb{R}$. Such asset payoff functions can be observed for example with credit default swaps, which have positive value if and only if the debtor solvency $o(\theta, i)$ is below a certain threshold that does not allow him to repay his debt. Bonds arguably also have a similar structure; they have positive value if and only if the issuing entity can repay. Moreover, prediction markets (e.g., Wolfers and Zitzewitz, 2004, Siemroth, 2014) often use this asset payoff function with so called 'winner take all' securities.

The following proposition shows that invertible asset payoff functions A are never worse, but may be better suited to promote fully revealing prices than non-invertible ones. Thus, we may say that invertible asset payoff functions are 'weakly preferable' in terms of information revelation. This is because a non-invertible asset payoff function may 'bunch' or 'pool' several states in a single asset value, thereby making it impossible to infer the state from the asset value, i.e., precluding invertibility of $v(\boldsymbol{s})$. Denote the set of all invertible functions $A : \mathbb{R} \to \mathbb{R}$ by $\mathcal{A}$ and the set of non-invertible functions by $\mathcal{A}'$.

Proposition 7 (Asset design and full revelation). *Consider the class of assets that* *can be written as* $a(\theta, i) = A(o(\theta, i))$ *for any* $o : \Theta \times I \to \mathbb{R}$*. If a non-invertible function* $A' \in \mathcal{A}'$ *allows full revelation and accurate prices, then so does any invertible* $A \in \mathcal{A}$*, but the converse does not hold.*

Proof. See Appendix. □

Consequently, in this setting, we would expect assets with invertible asset payoff function $A \in \mathcal{A}$ to be more informative, all else equal, compared to assets with non-invertible asset payoff function. Proposition 7 thus provides an empirical implication of the theory, and a recommendation for anyone designing assets/markets/policies for information revelation.

Besides changing the asset value function, the proper choice of the underlying of the asset can help support information revelation. The following example investigates the project deadline assets that are used in corporate prediction markets (Cowgill and Zitzewitz, 2015).

Example 5. A project in a company may miss the deadline ($\theta = 1$) or it may meet the deadline ($\theta = 0$) with the currently available resources. Insiders—e.g., the employees that work on the project—know the state ($s_j = \theta$), while the manager (policy maker) does not. The deadline asset in the corporate prediction market pays \$1 if and only if the deadline is missed, and \$0 otherwise. The manager can react to information about whether the deadline will be met: $i = 0$ means the project does not receive additional resources (more manpower, funds, etc.), and $i = 1$ means the project receives additional resources that definitely ensure completion on time. The manager does not want to commit additional resources unless it is necessary. Consequently, the asset value is $a(\theta, i) = \mathbf{1}\{i = 0 \wedge \theta = 1\}$, i.e., the project misses the deadline if and only if the manager does not commit additional resources ($i = 0$) and the project misses the deadline without additional resources ($\theta = 1$).

To determine whether full revelation is possible in equilibrium, calculate $v(\theta)$, i.e., the asset value if θ (the trader information) were revealed to the manager. Clearly, $v(\theta = 1) = 0$, since the optimal policy is $i(\theta = 1) = 1$. Moreover, $v(\theta = 0) = 0$ with $i(\theta = 0) = 0$. Thus, the $v(\theta)$-function is not invertible in θ, and a self-defeating prophecy prevents revelation of trader information: Traders anticipate that revelation of $\theta = 1$ triggers a policy reaction that prevents the deadline being missed, hence they would lose money by trading on their information.

How can corporations solve this problem? Since they are free to design other assets in their markets, a simple adjustment fixes the problem. Consider another asset with value $a(\theta, i) = i$, i.e., the asset pays \$1 if and only if the company commits additional resources. The $v(\theta)$-function is $v(\theta = 1) = i(\theta = 1) = 1$ and $v(\theta = 0) = i(\theta = 0) = 0$, i.e., it is invertible. Thus, instead of designing an asset that predicts the outcome (deadline missed), which the company might seek to manipulate depending on state and information, another asset simply predicts the intervention decision. In equilibrium, traders have incentives to forecast the optimal policy for the policy maker, and this forecast is a self-*fulfilling* prophecy, since the policy maker wants to follow the "recommendation." In conclusion, if one is free to design assets, it will typically be possible to find outcomes and asset payoff functions such that the $v(\boldsymbol{s})$-function is invertible, which guarantees existence of a revealing equilibrium (Corollary 5), i.e., proper incentives to reveal the information. $\square$

A natural question is whether it is always possible to find an asset payoff function $a : \Theta \times I \to \mathbb{R}$ that allows for fully revealing equilibria. This question is equivalent to asking whether an invertible function $a(\theta, i(\theta))$ always exists, but it is hard to answer at this level of generality. In theory, $a(\theta, i(\theta)) = \theta$ is always invertible in θ, yet θ may not be verifiable or contractible in all applications. For example, if θ is the quality of a company CEO, then the possibility of θ-revelation depends on whether observable measures of θ exist on which asset values might condition. If $a(\theta, i(\theta)) = \theta$ is not possible, then $a(\theta, i(\theta)) = i(\theta)$ works if $i(\theta)$ is invertible and verifiable. Consequently, this question has to be answered for specific applications.

1.3 The model with noise affecting market prices

1.3.1 Set-up

This section presents a financial market model with noise where a policy maker reacts to information contained in prices and thereby changes asset values. Besides making the model more realistic, noise solves some undesirable features of noiseless models such as no trade, and allows us to investigate whether the previous results are an artifact of noiseless markets. I extend a standard constant absolute risk aversion (CARA)-normal model by introducing a policy maker. The equilibrium market clearing price p is affected by the realization of a random noise variable u, which is independent of the state θ. In the common interpretation, u is the aggregate net demand of noise traders, whose trading activity (due to exogenous reasons such as liquidity shocks) is independent of the price/state.

For rational traders and the policy maker, the noise shocks introduce a difficulty in extracting information from the price: A high asset price may indicate favorable information about the fundamental θ, or it may indicate a lot of noise trader purchases u which are unrelated to fundamentals. Consequently, traders and the policy maker will only be able to make stochastic inferences about the realization of θ from the market price, and cannot perfectly infer θ from the market price as in the previous section without noise.

More specifically, I solve for a noisy REE, where[5]

[5] I heavily borrow notation from Vives (2010), who provides a detailed derivation of the standard linear noisy REE without policy maker.

- a riskless asset (return normalized to zero) and a risky asset is traded in the financial market,

- the risky asset value is the sum of fundamental and policy,

$$a(\theta, i) = \theta + i,^{6}$$

- fundamental $\theta \sim \mathcal{N}(\bar{\theta}, 1/\tau_\theta)$, private trader signals $s_j = \theta + \varepsilon_j,\ \varepsilon_j \sim \mathcal{N}(0, 1/\tau_\varepsilon)$, private policy maker signal $s_p = \theta + \varepsilon_p,\ \varepsilon_p \sim \mathcal{N}(0, 1/\tau_\varepsilon)$, and aggregate noise trader net demand $u \sim \mathcal{N}(0, 1/\tau_u)$ are normally distributed, with $0 < \tau_\theta, \tau_\varepsilon, \tau_u < \infty$, and $\varepsilon_p, \varepsilon_j, u$ are independent of θ,

- there is a continuum of traders $j \in [0, 1]$, and all j have a CARA utility function defined over investment returns π_j, $U_j(\pi_j) = -\exp(-\rho_j \pi_j)$, with $\pi_j = (\theta + i - p)x_j$ and net demand x_j for the risky asset,

- a share $\mu \in (0, 1]$ of traders are informed, receive i.i.d. signals s_j about the realization of θ, have a coefficient of absolute risk aversion $\rho_I > 0$, and use demand strategies $X_I(p, s_j) : \mathbb{R}^2 \to \mathbb{R}$ to be specified later,

- a share $(1 - \mu) \in [0, 1)$ of traders are uninformed, have a coefficient of absolute risk aversion $\rho_U > 0$, and use demand strategies $X_U(p) : \mathbb{R} \to \mathbb{R}$ to be specified later,

- the policy maker has a utility function defined over state θ and policy i, $U(\theta, i)$, for which $i(z) = \arg\max_i \mathbb{E}[U(\theta, i)|z]$ exists for any normally distributed signal z,

- the environment (i.e., all of the above except for the realization of random variables $(\{s_j\}_{j=0}^{\mu}, s_p, u, \theta)$) is common knowledge.

CARA-utility functions exhibit no wealth effects, hence I normalize wealth to zero without loss of generality. As is standard, there are no budget constraints in this model; demands are bounded by the degree of risk aversion.

The timing is still as depicted in Figure 2.1: First, nature draws θ, then traders trade leading to a market clearing price p, the policy maker sets i, and finally asset values and payoffs realize. Since the policy maker can condition on the price, but traders cannot condition on the policy, the model yields the same equilibria if we assume simultaneous decisions of traders and policy maker.

This extended model nests the standard CARA-normal REE, which is the special case $i = 0$, i.e., where the policy maker does not affect asset values. I choose the CARA-normal

[6] In the case with an uninformed policy maker, the results can be generalized any asset payoff function with an additively separable form $a(\theta, i) = \theta + g(i)$ with arbitrary $g : I \to \mathbb{R}$. In the case of an informed policy maker, however, this linear asset payoff function is critical. See also the discussion on quasi-linear equilibria below.

parametrization to facilitate comparison with the existing literature, which heavily relies on this framework since Grossman and Stiglitz (1980) and Hellwig (1980). Indeed, Vives (2010) calls it the workhorse model in the study of financial markets with asymmetric information.

If the policy maker is informed, it is crucial that demand and policy reaction functions are linear in price p and state θ to characterize the equilibrium at least implicitly. In these cases, I will use the following utility function,

$$U(\theta, i) = \psi_1 + \psi_2 i - \psi_3 i^2/2 + \psi_4 i \cdot \theta + \psi_5 \theta, \tag{1.1}$$

with $\psi_3 > 0$ to ensure a unique policy reaction and $\psi_4 \neq 0$ for non-triviality, which leads to a linear policy function $i(p, s_p) = \beta_1 + \beta_2 p + \beta_3 s_p$. This policy maker utility function is one flexible parameterization that yields a linear reaction function, although it is not the only one. One advantage of this particular parametrization is that it yields a simple and intuitive measure of the welfare impact of the additional information provided by the financial market, as is shown below. The reaction function of an informed policy maker is

$$\frac{\partial \mathbb{E}[U(\theta, i)|p, s_p]}{\partial i} = \psi_2 - \psi_3 i + \psi_4 \mathbb{E}[\theta|p, s_p] \stackrel{!}{=} 0 \iff i(p, s_p) = \frac{\psi_2 + \psi_4 \mathbb{E}[\theta|p, s_p]}{\psi_3}.$$

A noisy rational expectations equilibrium with an informed policy maker is defined as follows (the definition for an uninformed policy maker is easily adapted).

Definition 4. *A noisy rational expectations equilibrium is a set of trading strategies contingent on the available information, $X_I(p, s_j)$ for all $j \in [0, \mu]$ and $X_U(p)$ for all $j \in (\mu, 1]$, an optimal policy function $i(p, s_p)$, and a measurable price functional $P(\theta, u)$ such that*

1. *the market for the risky asset clears:*

$$\int_0^\mu X_I(p = P(\theta, u), s_j) dj + \int_\mu^1 X_U(p = P(\theta, u)) dj + u = 0 \text{ a.s.},$$

2. *all traders j use optimal demand strategies given the available information,*

$$X_I(p, s_j) \in \arg\max_x \ \mathbb{E}_{\theta, s_p}[U_j((\theta + i(p, s_p) - p)x)|p = P(\theta, u), s_j] \ \forall j \in [0, \mu],$$
$$X_U(p) \in \arg\max_x \ \mathbb{E}_{\theta, s_p}[U_j((\theta + i(p, s_p) - p)x)|p = P(\theta, u)] \ \forall j \in (\mu, 1],$$

3. *the policy maker sets an optimal policy given the available information,*

$$i(p, s_p) \in \arg\max_i \ \mathbb{E}_\theta[U(\theta, i)|p = P(\theta, u), s_p].$$

Any set of information structures and policy maker preferences that fulfills the parametric assumptions in this noisy REE setup can be translated to the model without noise, but not vice versa, since the model without noise allows for more general signal structures and policy maker preferences.

1.3.2 Measures of price informativeness and policy maker welfare

A new measure of price informativeness

Because the realization of the noise variable u introduces price movements unrelated to fundamentals or information, the price cannot perfectly reveal trader information or the state. Thus, a different concept than full revelation (definition 2) as in the model without noise is required to assess the informational content of market prices. The new measure is how much the additional information revealed by the market price reduces the policy maker θ-forecast error on average.

To derive this measure, first note that price informativeness is affected by how aggressively informed traders use their information to trade, which is captured by the coefficient a in the informed traders' net demand strategies,

$$X_I(p, s_j) = as_j - g_I(p).$$

Given $a > 0$, and holding noise variance $1/\tau_u > 0$ constant, a larger coefficient a means informed traders increase their net demand for the asset the more favorable their information s_j. Thus, the relative weight of information compared to noise increases, and prices become more informative. To see this more clearly, consider the market clearing condition if net demand strategies take the forms $X_I(p, s_j) = as_j - g_I(p)$ and $X_U(p) = -g_U(p)$ with equilibrium price function $P(\theta, u)$, which requires aggregate excess demand to be zero,

$$\begin{aligned} &\int_0^\mu (as_j - g_I(P(\theta, u)))dj - (1-\mu)g_U(P(\theta, u)) + u = 0 \\ \iff &\frac{1}{\mu a}\left[\mu g_I(P(\theta, u)) + (1-\mu)g_U(P(\theta, u))\right] = \theta + u/(\mu a) \end{aligned}$$

where I use a law of large numbers to evaluate the integral over i.i.d. trader signals s_j (Sun, 2006). Since the left hand side depends on (θ, u) only via $P(\theta, u)$, any price function that clears the market must reveal the term on the right hand side, $z := \theta + u/(\mu a)$. Hence, the 'price signal' is an additive statistic of the fundamental θ and a normally distributed noise term $u/(\mu a)$. This information contained in the price is independent of the endogenous demand strategy parts unrelated to private information (g_I and g_U), and only depends on how aggressively the informed trade on information (captured by strategy coefficient a).

Now, $\mathrm{Var}(u) = 1/\tau_u$, $\mathrm{Var}(u/(\mu a)) = 1/((\mu a)^2\tau_u)$, thus its precision is $\mathrm{Var}(u/(\mu a))^{-1} = (\mu a)^2\tau_u$. Using the standard Bayesian updating rule for normal distributions, the information z revealed by the price leads to the estimate

$$\mathbb{E}[\theta|p] = \mathbb{E}[\theta|z] = \frac{\tau_\theta \bar{\theta} + (\mu a)^2 \tau_u z}{\tau_\theta + (\mu a)^2 \tau_u}, \tag{1.2}$$

which is the precision weighted sum of prior mean and price signal. To define a measure of price informativeness, consider the mean squared error of the policy maker "estimator" of θ given p,

$$\mathrm{MSE}(\mathbb{E}[\theta|p]) := \mathbb{E}\left[(\theta - \mathbb{E}[\theta|p])^2|p\right],$$

which is a well known measure of the deviation or forecast error of an estimator from the variable to be estimated. To measure price informativeness, take the difference of the mean squared error of the policy maker estimate *without* the information contained in the market price and the mean squared error of the estimate *with* the information contained in the market price. Thus, the measure directly captures the differences in "forecast errors" of θ due to access to the price and the information contained therein: The larger the measure, the more the price information helps to estimate θ.[7] For an uninformed policy maker, the price informativeness measure is

$$\begin{aligned}\mathrm{PI}_{\text{uninformed}} :=& \mathrm{MSE}(\mathbb{E}[\theta]) - \mathrm{MSE}(\mathbb{E}[\theta|p]) = \mathbb{E}\left[(\theta - \mathbb{E}[\theta])^2\right] - \mathbb{E}\left[(\theta - \mathbb{E}[\theta|p])^2|p\right] \\ =& \mathrm{Var}(\theta) - \mathrm{Var}(\theta|p) = \frac{1}{\tau_\theta} - \frac{1}{\tau_\theta + (\mu a)^2 \tau_u} = \frac{1}{\tau_\theta^2/((\mu a)^2 \tau_u) + \tau_\theta},\end{aligned} \tag{1.3}$$

which is decreasing in prior precision τ_θ and increasing in $\mu, |a|, \tau_u$. Intuitively, if there is less uncertainty in the θ-realization (τ_θ increases), then more information (contained in p) does not help as much in predicting θ, thus the price informativeness measure decreases. $\mathrm{PI}_{\text{uninformed}}$ is not just price informativeness for the (uninformed) policy maker, but all uninformed traders or outsiders.

Similarly, we can define price informativeness if the policy maker is informed, which is the difference in MSE given her private information s_p and the MSE given her private information and the market price,

$$\begin{aligned}\mathrm{PI}_{\text{informed}} :=& \mathrm{MSE}(\mathbb{E}[\theta|s_p]) - \mathrm{MSE}(\mathbb{E}[\theta|p, s_p]) = \mathrm{Var}(\theta|s_p) - \mathrm{Var}(\theta|p, s_p) \\ =& \frac{1}{(\tau_\theta + \tau_\varepsilon)^2/((\mu a)^2 \tau_u) + \tau_\theta + \tau_\varepsilon},\end{aligned} \tag{1.4}$$

[7]An alternative measure of price informativeness is $|\mu a|$, i.e., how aggressively the informed trade on information weighted by the share of informed traders (e.g., Bond and Goldstein, 2015). While this measure is appealing, it does not exactly capture how informative prices are for the policy maker. For example, consider the limit case of perfect information for the policy maker, i.e., the variance of the error term in the policy maker signal approaches zero. Then prices do not reveal any new information to the policy maker, but traders will still gain by trading on information ($a \neq 0$). Similarly, $|\mu a|$ does not capture changes in the prior distribution, while the MSE measure does.

which is decreasing in τ_θ and τ_ε, and increasing in $\mu, |a|, \tau_u$.

A measure of welfare gains

A central question of this paper is to what degree information from financial markets can help policy makers to improve real decisions. To make this assessment, I need to impose more structure on the policy maker utility function, and I will use the quadratic policy maker utility function (1.1) for this purpose. The welfare measure of interest is how much the financial market information improves the utility of the policy maker by improving decisions. We can think of the policy maker utility function as a welfare measure for all the (unmodeled) individuals that are affected by the policy maker decision.[8] For example, if the financial market can reveal information about future inflation shocks to the central bank, then the central bank will be better able to hit the target inflation rate or maintain price stability, thus improving welfare.

Since the optimal policy depends on the realization of θ, a utility loss can result if the policy maker sets a wrong policy due to imperfect information about θ. Define the expected utility loss of a policy maker with information set $\mathcal{I}$ due to imperfect information, i.e., the expected utility difference of a perfectly informed policy maker observing the realization θ and a policy maker observing only $\mathcal{I}$, as

$$\begin{aligned}
L(\mathcal{I}) &:= \mathbb{E}[U(\theta, i(\theta)) - U(\theta, i(\mathcal{I}))|\mathcal{I}] \\
&= \mathbb{E}[\psi_2(i(\theta) - i(\mathcal{I})) - (i(\theta)^2 - i(\mathcal{I})^2)\psi_3/2 + \psi_4\theta(i(\theta) - i(\mathcal{I}))|\mathcal{I}] \\
&= \frac{\psi_2\psi_4(\mathbb{E}[\theta|\mathcal{I}] - \mathbb{E}[\mathbb{E}[\theta|\mathcal{I}]|\mathcal{I}])}{\psi_3} - \frac{\mathbb{E}[(\psi_2 + \psi_4\theta)^2 - (\psi_2 + \psi_4\mathbb{E}[\theta|\mathcal{I}])^2|\mathcal{I}]}{2\psi_3} \\
&\quad + \frac{\psi_4^2(\mathbb{E}[\theta^2|\mathcal{I}] - \mathbb{E}[\theta\mathbb{E}[\theta|\mathcal{I}]|\mathcal{I}])}{\psi_3} \\
&= -\frac{\psi_4^2(\mathbb{E}[\theta^2|\mathcal{I}] - \mathbb{E}[\theta|\mathcal{I}]^2)}{2\psi_3} + \frac{\psi_4^2(\mathbb{E}[\theta^2|\mathcal{I}] - \mathbb{E}[\theta|\mathcal{I}]^2)}{\psi_3} \\
&= \frac{\psi_4^2}{2\psi_3}(\mathbb{E}[\theta^2|\mathcal{I}] - \mathbb{E}[\theta|\mathcal{I}]^2) = \frac{\psi_4^2}{2\psi_3}\mathrm{Var}(\theta|\mathcal{I}).
\end{aligned}$$

Thus, the utility loss due to imperfect policy maker information is proportional to the variance of θ given policy maker information, multiplied by a factor depending on how strongly the bliss point of the utility function reacts to a change in θ, i.e., how important θ is in making policy. Now we can define a welfare measure that directly captures the welfare impact of financial market information on policy maker utility. The measure is the difference in utility loss if the policy maker does not have access to the information provided by the financial market and if she does. For an uninformed policy maker, this leads to a welfare

[8] Clearly, using policy maker utility as welfare measure assumes a benevolent policy maker.

improvement due to financial market information

$$
\begin{aligned}
\Delta W_{\text{uninformed}} :=& L(\varnothing) - L(p) = \frac{\psi_4^2}{2\psi_3}(\mathrm{Var}(\theta) - \mathrm{Var}(\theta|p)) \\
=& \frac{\psi_4^2}{2\psi_3}\mathrm{PI}_{\text{uninformed}} = \frac{\psi_4^2}{2\psi_3}\frac{1}{\tau_\theta^2/((\mu a)^2\tau_u) + \tau_\theta},
\end{aligned}
\tag{1.5}
$$

which is linear and increasing in price informativeness. Intuitively, the more informative the prices, the more they improve decisions and therefore welfare. The similarity of the expressions for price informativeness and welfare impact is a consequence of the quadratic utility function of the policy maker. The welfare impact of financial market information on an informed policy maker is

$$
\begin{aligned}
\Delta W_{\text{informed}} :=& L(s_p) - L(p, s_p) = \frac{\psi_4^2}{2\psi_3}(\mathrm{Var}(\theta|s_p) - \mathrm{Var}(\theta|p, s_p)) \\
=& \frac{\psi_4^2}{2\psi_3}\mathrm{PI}_{\text{informed}} = \frac{\psi_4^2}{2\psi_3}\frac{1}{(\tau_\theta + \tau_\varepsilon)^2/((\mu a)^2\tau_u) + \tau_\theta + \tau_\varepsilon}.
\end{aligned}
\tag{1.6}
$$

1.3.3 Results with an uninformed policy maker

I begin the analysis with a policy maker who does not receive an independent signal on the fundamental θ, hence the only information available to her (apart from the prior) is contained in market prices. Usually, the CARA-normal setup is used for its tractability and simple closed-form solutions. However, the typical guess-and-verify approach guessing a linear equilibrium price function can only analyze cases with linear policy reaction functions. For nonlinear reaction functions, the approach requires knowledge of the functional form of the equilibrium price function. Thus, the approach is unsuitable to derive a condition for equilibrium existence depending on policy maker preferences. Breon-Drish (2015) recently demonstrated a novel solution approach in a model without policy maker where the typical CARA-normal assumptions can be relaxed. Adapting this new approach in an extended model with policy maker, I can solve for equilibria with nonlinear policy reaction functions if the policy maker is uninformed, and can derive a condition for equilibrium existence depending on policy maker preferences. Since the new approach does not require knowledge of the form of the equilibrium price function, it is simpler and more useful to analyze the model with arbitrary policy maker preferences. The only other noisy model of policy maker-trader interaction that I am aware of (Bond and Goldstein, 2015) only considers linear policy functions. Since the equilibria can be solved in closed-form, the model is useful for many applications. The prime application is monetary policy: There is considerable evidence that central bank policy reacts nonlinearly or asymmetrically to changes in price levels and stock prices (e.g., Weise, 1999; Kim et al., 2005; Surico, 2007; Ravn, 2012).

Equilibrium

With an uninformed policy maker, I will consider the class of equilibria where demand functions of the informed traders are possibly nonlinear in the price but additively separable from signal s_j,

$$X_I(s_j, p) = as_j - g_I(p),\ X_U(p) = -g_U(p),$$

and where the equilibrium price function $P(\theta, u)$ is continuous.[9] This class includes and generalizes the linear equilibria solved for in the standard guess-and-verify approach. Equilibrium existence or non-existence is understood within this equilibrium class.

Definition 5 (**Quasi-linear equilibrium**). *The class of equilibria where the net demand function of the informed traders takes the form $X_I(s_j, p) = as_j - g_I(p)$ for constant a and function $g_I : \mathbb{R} \to \mathbb{R}$, with continuous price function $p = P(\theta, u)$, is called quasi-linear.*

Recall that the equilibrium price function $P(\theta, u)$ maps all possible realizations (θ, u) into a price $p \in \mathbb{R}$. The market clearing condition given the demand functions is

$$\begin{aligned} &\int_0^\mu as_j - g_I(P(\theta, u))dj - (1-\mu)g_U(P(\theta, u)) + u = 0 \\ \iff\ &\frac{1}{\mu a}(\mu g_I(P(\theta, u)) + (1-\mu)g_U(P(\theta, u))) = \theta + u/(\mu a). \end{aligned} \tag{1.7}$$

Only the right hand side directly depends on (θ, u). Since the market clearing condition has to hold for all realizations (θ, u), it implies that the left hand side has to react to any change in $\theta + u/(\mu a)$. Since the left hand side depends on θ and u only via the equilibrium price function $P(\theta, u)$, it further implies that any equilibrium price function must reveal the linear statistic $\theta + u/(\mu a)$, i.e., a noisy signal of the state θ. Consequently, the price function must be invertible in $\theta + u/(\mu a)$, and defining $z := \theta + u/(\mu a)$, I shall write $P(z)$ instead of $P(\theta, u)$ in the following. Thus, if there is an equilibrium, the policy maker can infer at least the realization z from the price. Continuity of the price function ensures that prices reveal z and no more, as later shown in the proof. The contribution of Breon-Drish (2015) is to recognize that this approach allows us to pin down the information set of the uninformed (which in this setting includes the policy maker) without knowing the functional form of the equilibrium price function $P(\theta, u)$. The standard guess-and-verify approach, on the other hand, heavily relies on guessing the correct functional form of the equilibrium price function.

The utility function $U(\theta, i)$ of the uninformed policy maker is an arbitrary function of θ and i so that $i(p) \in \arg\max_i \mathbb{E}_\theta[U(\theta, i)|p = P(z)]$ always exists for normally distributed signals p and is continuous (to ensure that the price function is continuous). Moreover, for

[9] While equilibria in the class of continuous equilibria are unique (see Breon-Drish, 2015 and Proposition 8 below), Pálvölgyi and Venter (2015) show for the classical Grossman and Stiglitz (1980)-model that a continuum of discontinuous equilibria with varying price informativeness exist besides the standard linear equilibrium. Thus, we may view the focus on continuous price functions as a reasonable equilibrium refinement.

equilibrium uniqueness I require $i(p)$ to be unique. $i(p)$ is the policy maker reaction function to prices. The asset value $\theta + i(p)$ conditional on the price is normally distributed, even if the policy reaction function $i(p)$ is nonlinear, since it only depends on p. For a normally distributed asset value, the optimal demand of the informed traders with CARA utility is

$$
\begin{aligned}
X_I(s_j, p) &= \frac{\mathbb{E}[\theta + i(p) - p | s_j, p]}{\rho_I \mathrm{Var}(\theta + i(p) - p | s_j, p)} = \frac{\mathbb{E}[\theta | s_j, p] + i(p) - p}{\rho_I \mathrm{Var}(\theta | s_j, p)} \\
&= \frac{\tau_\theta + \tau_\varepsilon + (\mu a)^2 \tau_u}{\rho_I} \left(\frac{\tau_\theta \bar{\theta} + \tau_\varepsilon s_j + (\mu a)^2 \tau_u P^{-1}(p)}{\tau_\theta + \tau_\varepsilon + (\mu a)^2 \tau_u} + i(p) - p \right) \\
&= \left[\tau_\theta \bar{\theta} + \tau_\varepsilon s_j + (\mu a)^2 \tau_u P^{-1}(p) + (i(p) - p)(\tau_\theta + \tau_\varepsilon + (\mu a)^2 \tau_u) \right] / \rho_I,
\end{aligned}
\tag{1.8}
$$

which is of the form $X_I(s_j, p) = a s_j - g_I(p)$ as assumed, with

$$
g_I(p) = - \left[\tau_\theta \bar{\theta} + (\mu a)^2 \tau_u P^{-1}(p) + (i(p) - p)(\tau_\theta + \tau_\varepsilon + (\mu a)^2 \tau_u) \right] / \rho_I, \tag{1.9}
$$

and, matching the coefficient, $a = \tau_\varepsilon / \rho_I$. The demand of the uninformed traders is similarly

$$
X_U(p) = \frac{\mathbb{E}[\theta | p] + i(p) - p}{\rho_U \mathrm{Var}(\theta | p)} = \left[\tau_\theta \bar{\theta} + (\mu a)^2 \tau_u P^{-1}(p) + (i(p) - p)(\tau_\theta + (\mu a)^2 \tau_u) \right] / \rho_U, \tag{1.10}
$$

with

$$
g_U(p) = - \left[\tau_\theta \bar{\theta} + (\mu a)^2 \tau_u P^{-1}(p) + (i(p) - p)(\tau_\theta + (\mu a)^2 \tau_u) \right] / \rho_U. \tag{1.11}
$$

These demand functions are uniquely determined given that prices reveal z. This fact together with the requirement that $P(z)$ be invertible implies equilibrium uniqueness in the class of quasi-linear equilibria. Substituting $g_I(p), g_U(p)$ into the market clearing condition and setting $p = P(z)$,

$$
\begin{aligned}
&-\frac{1}{\mu a} \big(\mu \left[\tau_\theta \bar{\theta} + (\mu a)^2 \tau_u P^{-1}(P(z)) + (i(P(z)) - P(z))(\tau_\theta + \tau_\varepsilon + (\mu a)^2 \tau_u) \right] / \rho_I \\
&+ (1 - \mu) \left[\tau_\theta \bar{\theta} + (\mu a)^2 \tau_u P^{-1}(P(z)) + (i(P(z)) - P(z))(\tau_\theta + (\mu a)^2 \tau_u) \right] / \rho_U \big) = z.
\end{aligned}
\tag{1.12}
$$

As established before, market clearing requires the equilibrium price function $P(z)$ to be invertible. Thus, we can rewrite $P^{-1}(P(z)) = z$ and, abusing notation, $i(P(z)) = i(z)$ with $i(z) = \arg\max_i \mathbb{E}[U(\theta, i) | z]$. These are the price terms that only depend on the information contained in prices and do not change if we change the price function (assuming it remains invertible and thus reveals the same information z). The market clearing condition is therefore

$$
\begin{aligned}
&-\frac{1}{\mu a}\big(\mu\left[\tau_\theta\bar{\theta}+(\mu a)^2\tau_u z+(i(z)-P(z))(\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u)\right]/\rho_I \\
&+(1-\mu)\left[\tau_\theta\bar{\theta}+(\mu a)^2\tau_u z+(i(z)-P(z))(\tau_\theta+(\mu a)^2\tau_u)\right]/\rho_U\big)=z \\
&\iff -\frac{1}{\mu a}\big(\mu\left[\tau_\theta\bar{\theta}+(i(z)-P(z))(\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u)\right]/\rho_I \\
&+(1-\mu)\left[\tau_\theta\bar{\theta}+(i(z)-P(z))(\tau_\theta+(\mu a)^2\tau_u)\right]/\rho_U\big)=z(1+\mu^2 a\tau_u/\rho_I+(1-\mu)\mu a\tau_u/\rho_U) \\
&\iff P(z)\left[\mu(\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u)/(\rho_I\mu a)+(1-\mu)(\tau_\theta+(\mu a)^2\tau_u)/(\rho_U\mu a)\right] \\
&=z(1+\mu^2 a\tau_u/\rho_I+(1-\mu)\mu a\tau_u/\rho_U)+\mu[\tau_\theta\bar{\theta}+i(z)(\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u)]/(\rho_I\mu a) \\
&+(1-\mu)[\tau_\theta\bar{\theta}+i(z)(\tau_\theta+(\mu a)^2\tau_u)]/(\rho_U\mu a) \\
&\iff P(z)=\Big\{z(1+\mu^2 a\tau_u/\rho_I+(1-\mu)\mu a\tau_u/\rho_U)+\mu[\tau_\theta\bar{\theta}+i(z)(\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u)]/(\rho_I\mu a) \\
&+(1-\mu)[\tau_\theta\bar{\theta}+i(z)(\tau_\theta+(\mu a)^2\tau_u)]/(\rho_U\mu a)\Big\} \\
&/\Big\{\mu(\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u)/(\rho_I\mu a)+(1-\mu)(\tau_\theta+(\mu a)^2\tau_u)/(\rho_U\mu a)\Big\},
\end{aligned}
\tag{1.13}
$$

which is an explicit expression for the price function given optimal demands if the price function reveals z to all traders and policy maker, and if traders anticipate the policy reaction $i(z)$. For equilibrium, we still have to confirm that the price function is in fact invertible, as required for market clearing. Only the numerator depends on z, so to determine invertibility we can drop the denominator, which is a positive constant. We can also drop additive terms from the numerator. Thus, $P(z)$ is invertible if and only if

$$
\begin{aligned}
z(1+\mu^2 a\tau_u/\rho_I+(1-\mu)\mu a\tau_u/\rho_U)+i(z)[\mu(\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u)/(\rho_I\mu a) \\
+(1-\mu)(\tau_\theta+(\mu a)^2\tau_u)/(\rho_U\mu a)]
\end{aligned}
\tag{1.14}
$$

is invertible in z. These terms capture how aggregate net demand changes if z changes. But (1.14) might not be invertible if $i(z)$ is at least locally decreasing, since the terms in brackets are positive. Thus, the one and only possible cause of non-invertibility and equilibrium non-existence is the policy reaction function $i(z)$ to information z: In the standard case without policy maker, $i(z)=0$ for all z, so the equilibrium price function is linearly increasing and an equilibrium always exists. Notice the similarity to the invertibility condition in the model without noise, which requires that the expected asset value given optimal policy and trader information $v(\boldsymbol{s})$ is invertible. If we let the noise variance go to zero here ($\tau_u\to\infty$), then (1.14) reduces to invertibility of $z+i(z)$, and since $z=\theta+u/(\mu a)=\theta$ a.s. for $\tau_u\to\infty$, $\theta+i(\theta)$ is the asset value given optimal policy and trader information.

The approach taken in this section is similar to the approach taken in the section without noise: I first assume the price function is invertible in $z=\theta+u/(\mu a)$, a noisy signal of θ, and therefore fully reveals the noisy signal. Second, I obtain the reaction function $i(z)$ of the policy maker to that information, which determines how the expected asset value is affected by policy. Finally, I check whether the price function given the optimal demands of the traders is in fact invertible. If it is not, then traders are not willing to clear the market for any price function that fully reveals z and triggers policy reaction $i(z)$, so that no equilibrium in the class of quasi-linear equilibria exists. As in the model without noise, it is a self-defeating prophecy that precludes revealing financial market prices: If market prices were revealing, they would lead to policy reactions for which traders would not want to clear the market.

The following proposition gives the formal result.

Proposition 8. *An equilibrium in the class of quasi-linear equilibria exists if and only if there exists a continuous reaction function $i(z) \in \arg\max_i \mathbb{E}[U(\theta, i)|z]$ such that, for $a = \tau_\varepsilon/\rho_I$,*

$$z(1 + \mu^2 a\tau_u/\rho_I + (1-\mu)\mu a\tau_u/\rho_U) + i(z)[\mu(\tau_\theta + \tau_\varepsilon + (\mu a)^2\tau_u)/(\rho_I \mu a) \\ + (1-\mu)(\tau_\theta + (\mu a)^2\tau_u)/(\rho_U \mu a)]$$

is invertible in z. If it exists, then the equilibrium is unique in the class of quasi-linear equilibria and market prices fully reveal the noisy signal $z = \theta + u/(\mu a)$.

The trader equilibrium strategies are $X_I(s_j, p) = as_j - g_I(p)$ and $X_U(p) = -g_U(p)$ with $a = \tau_\varepsilon/\rho_I$, $g_I(p)$ given in (1.9), *and $g_U(p)$ given in* (1.11). *The equilibrium policy reaction is $i(P^{-1}(p)) = i(z)$. The equilibrium price function $P(\theta, u) = P(z)$ is given in* (1.13).

Proof. See Appendix. □

Interestingly, market prices are informative, but the trading aggressiveness a of the informed traders is constant and unaffected by policy maker preferences. The explanation is given below in example 6.

The next result makes the requirement on the reaction function $i(z)$ more salient: Assuming differentiability, its derivative has to be either large enough or small enough for all z, but it may not have both positive and large negative slopes on $z \in \mathbb{R}$. Consequently, quadratic or quartic reaction functions will produce self-defeating prophecies, whereas almost all linear functions $i(z)$ and all weakly increasing functions $i(z)$ allow for revealing equilibria. Moreover, "weak" non-invertibilities in $i(z)$ such that $\sup_z i'(z) > 0$ and $0 > \inf_z i'(z) > \xi$, where $\xi < 0$ is defined in (1.15), can be supported in equilibrium.

Corollary 9. *If the policy reaction function $i(z)$ is continuous and differentiable, then there exists an equilibrium in the class of quasi-linear equilibria if and only if $i'(z) > \xi$ almost everywhere or $i'(z) < \xi$ almost everywhere, with $a = \tau_\varepsilon/\rho_I$ and*

$$\xi := -\frac{1 + \mu^2 a\tau_u/\rho_I + (1-\mu)\mu a\tau_u/\rho_U}{\mu(\tau_\theta + \tau_\varepsilon + (\mu a)^2\tau_u)/(\rho_I\mu a) + (1-\mu)(\tau_\theta + (\mu a)^2\tau_u)/(\rho_U\mu a)} < 0. \tag{1.15}$$

Proof. A continuous function on a convex set is invertible if and only if it is strictly monotone. Thus, taking the derivative of (1.14) with respect to z and requiring it to be strictly greater or smaller than zero for almost all z (allowing for saddle points on a measure zero set) yields the result. $\square$

In the noiseless limit as $\text{Var}(u) = 1/\tau_u \to 0$, $\xi \to -1$, i.e., the policy reaction $i(z)$ may not exactly offset any increase in $z = \theta + u/(\mu a) = \theta$ (a.s.) for an informative equilibrium to exist. An example of $i'(z) = -1$ in the noiseless case is Bernanke and Woodford (1997), as discussed below.

To summarize the results so far, the new approach of solving for noisy rational expectations equilibria allows for nonlinear policy maker reaction functions, which covers a larger set of policy maker preferences and settings with self-defeating prophecies from the no-noise model. This section is therefore a large step towards modeling and analyzing the policy maker-trader interactions in the presence of noise, providing the tools for many applications beyond linear policy rules. While the results here focused on continuous reaction functions $i(z)$, there are no functional form restrictions beyond that. Proposition 8 shows that noise cannot generally solve the problem of self-defeating prophecies, thus the problem of self-defeating prophecies also occurs in more realistic models with noise.

Still, situations with policy maker objectives that implied a self-defeating prophecy without noise may not be a self-defeating prophecy with noise, hence noise may support rather than hamper information revelation via prices in some instances. To understand why, consider an equilibrium where the policy maker has the quadratic utility function (1.1), which yields a linear reaction function, with $\psi_4/\psi_3 = -1$. $\psi_4/\psi_3 = -1$ means that the policy maker sets $i(p)$ such that $\mathbb{E}[\theta + i(p)|p] = c$ for some $c \in \mathbb{R}$, and consequently the asset value $\theta + i(p)$ would be constant for every realization of θ if prices were fully revealing (i.e., condition 1 does not hold, so there is no revealing equilibrium without noise). Bernanke and Woodford (1997) is one example from the literature which investigates exactly this case. Consider an example in the noisy REE setup, where an equilibrium with the following strategies exists.

Example 6. The parameters are $\psi_4/\psi_3 = -1$, $\psi_2 = 0$, $\bar{\theta} = 3$, $\tau_\varepsilon = 2$, $\tau_\theta = 1$, $\tau_u = 1$, $\mu = 0.5$, $\rho_I = 2$, $\rho_U = 2$. The resulting noisy REE strategies are

$$X_I(p, s_j) = s_j - 2.1875p - 3,\ X_U(p) = -0.625p,$$
$$i(p) = -3 - 0.5625p.$$

Figure 3.2 plots the equilibrium price function at the mean/median value of the noise variable ($u = 0$), which is also the average/median price function due to linearity in u. The price function is increasing in θ, even though the policy maker wants to implement a constant asset value. The figure further displays the expected value of θ given the information contained in price p, which is also increasing, but at a relatively flat slope. This line is

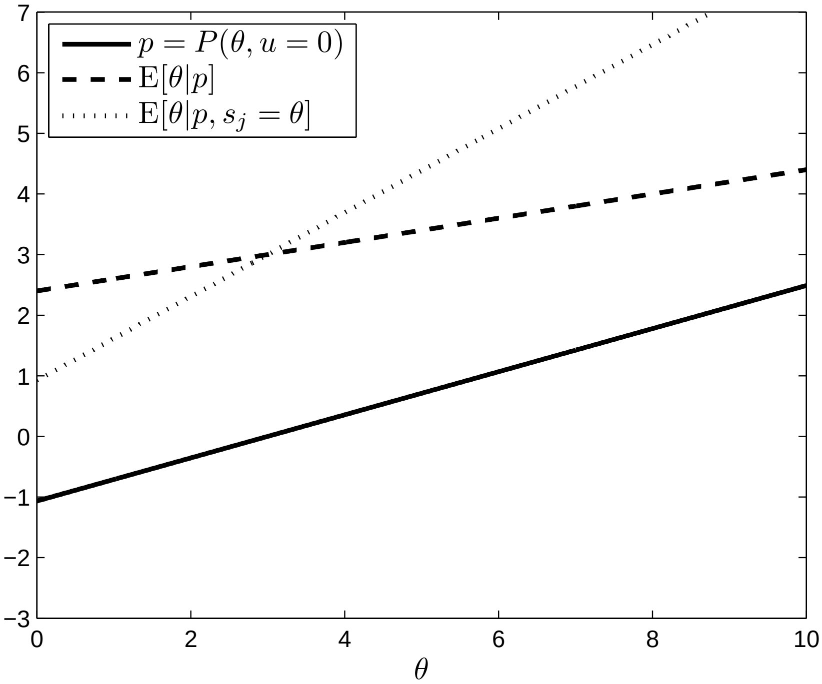

Figure 1.3: Plot of the average equilibrium price function $P(\theta, u = 0)$ (solid line), the expected value of the state variable θ for the uninformed (dashed), and the expected value of the state variable for the average informed trader (dotted).

the inference that the policy maker and uninformed traders draw from the price. If prices were fully revealing, then the expected value of θ given the price would be the 45° line, but the noise prevents full revelation of the state. The line of the expected value for the average informed trader is steeper than the line of the uninformed, reflecting their superior information about θ, but still smaller than one.

Figure 1.4 explains why the equilibrium price is informative, but not perfectly informative. It plots the average equilibrium price function, as well as the expected asset value given p (dashed)—this is the expected asset value for uninformed traders and the policy maker—and the expected asset value given p and signal $s_j = \theta$ (dotted), which is the expected asset value for the mean/median informed trader. First consider the expected asset value for the uninformed. It is constant, which just reflects the policy maker preferences $\psi_4/\psi_3 = -1$, where she sets $i(p)$ such that $\mathbb{E}[\theta + i(p)|p]$ is constant. Uninformed traders sell the asset for any price above 0—where the expected profit from buying is negative—and buy for any price below 0.[10]

The important line is the expected asset value for the average informed trader. His asset valuation and expected profit from buying increases the larger θ, since the expected asset value function is steeper than the price function. The reason is that the equilibrium price

[10]This example may seem to suggest that uninformed traders lose money, but this is just because the graphs consider the realization $u = 0$, which implies that noise traders do not trade. For any other realization of u (measure 1) they do trade and lose money in expectation to the uninformed (and informed) traders.

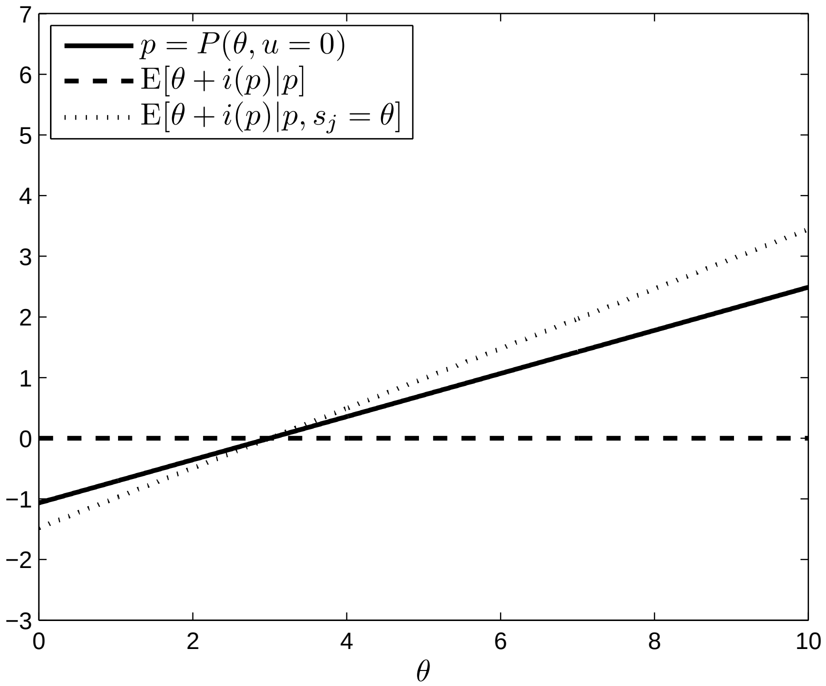

Figure 1.4: Plot of the average equilibrium price function $P(\theta, u = 0)$ (solid line), the expected value of the asset for the uninformed (dashed), and the expected value of the asset for the average informed trader (dotted).

function only imperfectly reveals information about θ, which is why the policy maker and uninformed traders tend to underestimate θ for θ-realizations above the prior mean $\bar{\theta} = 3$, and tend to overestimate for θ-realizations below the prior mean (see Figure 3.2). The informed traders have a better estimate of θ due to their signals. The expected asset value for most informed traders exceeds the price for $\theta > \bar{\theta}$, because their signals tend to be closer to the actual θ than what the price implies, so they buy more of the asset the larger their realization of s_j, which increases their estimate of θ. This is why $a > 0$ and why the equilibrium price function is informative: Informed traders have incentives to trade on their information, because noise prevents full revelation of their information, which otherwise would lead to a self-defeating prophecy. □

Thus, similar to the solution of the Grossman and Stiglitz (1980) paradox, prices in this instance are informative *because* of noise, rather than despite of noise. The explanation that the informed traders profit from information applies beyond this example, since the estimate of θ for the uninformed tends to be more tilted towards the prior mean $\bar{\theta}$, so it always pays for the informed to increase net demand for larger s_j in equilibrium.

Finally, it is also interesting to consider the limit as the variance of the noise variable u approaches zero, or equivalently, if the precision diverges to infinity. The next result shows that, in the limit, market prices are fully revealing if an equilibrium exists. Thus, as the noise variance becomes smaller and smaller, market prices become arbitrarily close to fully revealing.

Corollary 10. *If the policy maker is uninformed, then equilibrium market prices become fully revealing in the limit as the noise variance* $1/\tau_u$ *approaches zero, i.e.,*

$$\mathbb{E}[\theta|p] \xrightarrow{a.s.} \theta \text{ for } \tau_u \to \infty.$$

Proof. From Proposition 8, $a = \tau_\varepsilon/\rho_I > 0$ in equilibrium, which is independent of u or τ_u. Thus,

$$\lim_{\tau_u\to\infty} \mathbb{E}[\theta|p] = \lim_{\tau_u\to\infty} \frac{\tau_\theta\bar{\theta} + (\mu a)^2\tau_u(\theta + u/(\mu a))}{\tau_\theta + (\mu a)^2\tau_u} = \lim_{\tau_u\to\infty} \theta + u/(\mu a)) = \theta \text{ (a.s.)},$$

since $\mathbb{E}[u] = 0$ and $\mathrm{Var}(u) = 1/\tau_u \to 0$. □

Comparative statics

The last results for an uninformed policy maker determine the factors that influence price informativeness and policy maker welfare if an equilibrium exists.

Proposition 11 (Comparative statics). *In the quasi-linear noisy REE with uninformed policy maker,*

- *larger trader signal precision* τ_ε*, share of informed traders* μ*, or noise precision* τ_u*, and*
- *smaller prior distribution precision* τ_θ*, or risk aversion* ρ_I*,*

increase price informativeness (1.3)*, all else equal.*

Proof. From (1.3), price informativeness does not depend on demand strategies except for coefficient a. From Proposition 8, the equilibrium strategy coefficient is $a = \tau_\varepsilon/\rho_I$. Thus, price informativeness is increasing in $|a|, \mu$ and τ_u, and $|a|$ in turn is increasing in τ_ε and decreasing in ρ_I. Moreover, price informativeness (1.3) decreases in τ_θ. □

To assess the welfare impact of additional information from financial market prices, we need to be more explicit about policy maker preferences. The following corollary derives the comparative statics if the policy maker has quadratic utility. The results are identical to those of price informativeness, since more information about θ helps the policy maker to implement better policies, except for the effect of $|\psi_4/\psi_3|$. This coefficient in the utility function of the policy maker (1.1) captures how strongly the optimal policy depends on the realization of θ. If it is larger, then better information about θ improves welfare more.

Corollary 12 (**Comparative statics welfare**). *In the linear noisy REE with quadratic policy maker utility function* (1.1),

- *larger* $|\psi_4/\psi_3|$, *trader signal precision* τ_ε, *share of informed traders* μ, *or noise precision* τ_u, *and*

- *smaller prior distribution precision* τ_θ, *or risk aversion* ρ_I,

increase the policy maker utility gain due to financial market information (1.5), *all else equal.*

Proof. From (1.5), the welfare gain is an increasing linear function of price informativeness, hence the same comparative statics as in Proposition 11 hold. In addition, the factor ψ_4^2/ψ_3 is increasing in $|\psi_4/\psi_3|$ ($\psi_3 > 0$ by construction). □

1.3.4 Results with an informed policy maker

Equilibrium

If the policy maker is informed, i.e., receives independent information s_p, then demand functions of the informed are not of the quasi-linear form as in definition 5 unless the policy reaction function is linear. Consequently, throughout this section, I assume the policy maker has the quadratic utility function (1.1), which implies a linear reaction function of the form $i(p, s_p) = \beta_1 + \beta_2 p + \beta_3 s_p$. The ratio of two parameters from the utility function, ψ_4/ψ_3, will be of particular interest in this section, which determine how strongly and in which direction the optimal policy reacts to the realization of θ. Traders have linear demand functions of the form $X_I(p, s_j) = as_j - c_I p + b_I$ and $X_U(p) = -c_U p + b_U$.

Despite these stronger assumptions compared to the previous sections, the equilibrium cannot be explicitly solved, since the equilibrium condition (1.28) determining a is a fifth degree polynomial, which does not generally admit an analytical solution. Without solving explicitly, I first establish existence of equilibrium for almost all parameter profiles, but there need not be a unique equilibrium. Then I give conditions for equilibrium uniqueness and derive comparative statics in these unique equilibria. Finally, I explain how "policy risk" affects equilibrium trading behavior and price informativeness, and how policy risk can create a form of strategic complementarity and induce multiple equilibria.

Proposition 13 (**Equilibrium existence**). *If the policy maker is informed, then a linear noisy REE exists for almost all parameter values.*

Proof. See Appendix. □

There may be different equilibrium values of a, but all other strategy coefficients c_I, c_U, b_I, $b_U, \beta_1, \beta_2, \beta_3$ are unique given a. Thus, a—how aggressively the informed trade on information—is the only source of equilibrium multiplicity. An equilibrium may fail to exist for a negligible set of parameter values because of a self-defeating prophecy, where—if prices were informative—the policy reaction to prices $i(p, s_p)$ would exactly cancel out any favorable information about θ, so traders have no reason to change their positions for different prices, which means prices cannot clear the market. The intuition is very much as in the case of an uninformed policy maker or as in the noiseless case. Due to the restriction to linear reaction functions, self-defeating prophecies are only a knife-edge case in $\psi_4/\psi_3 \in \mathbb{R}$ space.

Without solving a explicitly, it is obvious from the equilibrium condition that $a = 0$, i.e., a completely uninformative equilibrium, can never be a solution. Thus, noisy REE are always informative if they exist.

Corollary 14. *In any equilibrium, $a \neq 0$, i.e., market prices are informative.*

Proof. Setting $a = 0$ in equilibrium condition (1.28) yields a contradiction for any $\tau_\varepsilon > 0$. □

We can again consider the limit as the noise variance approaches zero. As with an uninformed policy maker, prices come arbitrarily close to fully revealing in any equilibrium.

Corollary 15. *If the policy maker is informed, then equilibrium market prices become fully revealing in the limit as the noise variance $1/\tau_u$ approaches zero, i.e.,*

$$\mathbb{E}[\theta|p] \xrightarrow{a.s.} \theta \text{ for } \tau_u \to \infty.$$

Proof. Corollary 14 proved that $a \neq 0$ in any equilibrium. Thus,

$$\lim_{\tau_u \to \infty} \mathbb{E}[\theta|p] = \lim_{\tau_u \to \infty} \frac{\tau_\theta \bar{\theta} + (\mu a)^2 \tau_u (\theta + u/(\mu a))}{\tau_\theta + (\mu a)^2 \tau_u} = \lim_{\tau_u \to \infty} \theta + u/(\mu a)) = \theta \text{ (a.s.)},$$

since $\mathbb{E}[u] = 0$ and $\text{Var}(u) = 1/\tau_u \to 0$. □

Although existence of an equilibrium is generic (Proposition 13), uniqueness (within the class of linear equilibria) is not guaranteed. While determining a necessary and sufficient condition for uniqueness may be possible, e.g., using Sturm's theorem on the fifth degree polynomial (1.28), the resulting condition will be too complex to be of value here. Instead, I shall give simple sufficient conditions on the parameters for uniqueness. Not only are the predictions of the model sharper with a unique equilibrium, it also makes comparative statics unambiguous while they are typically not if multiple equilibria exist (see below).

The next proposition shows that sufficiently strong policy maker preferences for intervention in either direction ($\psi_4/\psi_3 >> 0$ or $\psi_4/\psi_3 << 0$) or strong risk aversion among informed

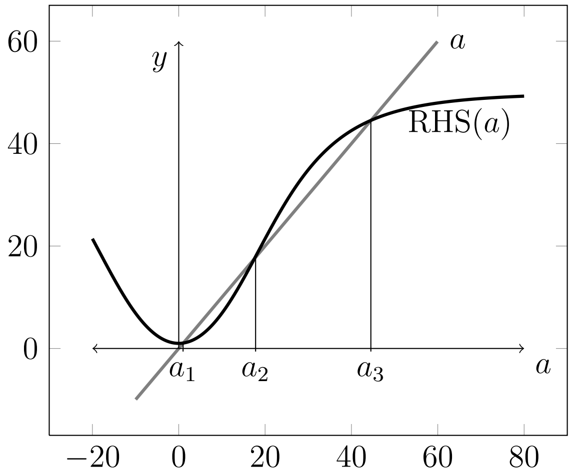

Figure 1.5: Equilibrium multiplicity. The graph plots the left hand side ($y = a$) and right hand side ($y = \text{RHS}(a)$) of equilibrium condition (1.28); the three intersections determine three different equilibrium values of a.[11]

traders guarantee a unique equilibrium in the class of unique equilibria. Technically, these conditions guarantee that the right hand side of equilibrium condition (1.28) plotted in Figure 1.5 is very flat, so that it intersects the 45° line (left hand side of the equilibrium condition) only once and the equilibrium is unique. Intuitively, these conditions guarantee that the policy maker reaction to a change in how aggressively informed traders use their information (a) changes trader payoffs only slightly, thus the complementarity between a and policy maker reaction is weak.

Strong enough strategic complementarity can induce multiple equilibria as follows. If informed traders increase $|a|$, i.e., trade more aggressively on information, then equilibrium prices become more informative. Consequently, the policy maker relies more on the price (rather than her private information) when inferring the state and making policy. This can be directly seen in the weight the policy maker places on her private information when making inferences about θ using Bayes' rule and the resulting reaction to private information in (1.27). From the perspective of the traders, a smaller weight on private information by the policy maker in determining policy (and thus asset values) typically reduces the asset variance conditional on the price (1.22), because the risk introduced by the private information of the policy maker is reduced. Thus, larger $|a|$ induces a policy maker reaction that reduces the asset variance, which makes it more attractive to trade on information, i.e., increase $|a|$. Consequently, multiple equilibria with small or large $|a|$ may exist.

Proposition 16 (Equilibrium uniqueness).

i. There exists $r^* > 0$ such that, for all $\psi_4/\psi_3 > r^*$, the linear equilibrium is unique with $a > 0$.

[11]The parameter profile used for the graph is $\psi_4/\psi_3 = -1$, $\tau_\varepsilon = 5$, $\tau_\theta = 0.1$, $\tau_u = 0.6$, $\mu = 0.1$, $\rho_I = 0.1$.

ii. There exists $r^* < 0$ such that, for all $\psi_4/\psi_3 < r^*$, the linear equilibrium is unique with $a < 0$.

iii. There exists $\rho^* > 0$ such that, for all $\rho_I > \rho^*$, the equilibrium is unique.

Proof. See Appendix. □

Comparative statics

The next proposition derives the comparative statics for the price informativeness for the cases where ψ_4/ψ_4 is either sufficiently positive or sufficiently negative, i.e., where the policy maker has strong preferences for intervention, which guarantees a unique equilibrium (Proposition 16). These are the most interesting cases, since ψ_4/ψ_3 close to zero implies that the optimal policy barely depends on θ, so the proposition considers the cases where additional information really matters for policy.

It is possible to derive comparative statics under other conditions that guarantee uniqueness, but then the comparative statics may differ for $\psi_4/\psi_3 < 0$ close to zero and $\psi_4/\psi_3 << 0$. It is also possible to derive comparative statics if there are multiple equilibria, but then comparative statics will be equilibrium specific and there will typically exist at least one other equilibrium where the statics are reversed. If there are multiple equilibria as plotted in Figure 1.5 and if $\psi_4/\psi_3 > 0$, then the comparative statics results of (i.) carry over to all equilibria where the slope of the RHS is smaller than 1, i.e., the RHS crosses from above (a_1 and a_3 in Figure 1.5), and are reversed for those where the RHS crosses from below (a_2 in Figure 1.5). Due to this possible ambiguity of the comparative statics, the proposition focuses on cases where a unique equilibrium and unambiguous comparative statics are guaranteed.

For the purpose of comparative statics, I allow the quality of information for informed traders and policy maker to vary separately. τ_ε is the precision of the trader signals and τ_ε^P is the precision of the policy maker signal.

Proposition 17 (**Comparative statics**).

i. For $\psi_4/\psi_3 > 0$ sufficiently large,

- *decreasing ψ_4/ψ_3, policy maker information precision τ_ε^P, prior distribution precision τ_θ, or informed trader risk aversion ρ_I, and*
- *increasing the share of informed traders μ, noise precision τ_u, or trader signal precision τ_ε*

increases price informativeness (1.4) *for the policy maker in the unique linear equilibrium. Moreover, all of these except for ψ_4/ψ_3 unambiguously increase the welfare gain due to financial market information* (1.6).

ii. For $\psi_4/\psi_3 < 0$ sufficiently negative,

- *decreasing policy maker information precision τ_ε^P, prior distribution precision τ_θ, or informed trader risk aversion ρ_I, and*
- *increasing ψ_4/ψ_3, the share of informed traders μ, noise precision τ_u, or trader signal precision τ_ε*

increases price informativeness in the unique linear equilibrium. Moreover, all of these except for ψ_4/ψ_3 unambiguously increase the welfare gain due to financial market information.

Proof. See Appendix. □

The intuition for these results is as follows. If $\psi_4/\psi_3 > 0$, then the policy maker has a preference for implementing policies that increase the asset value more for larger θ. Thus, the policy maker intervention amplifies the conditional asset value variance $\mathrm{Var}(\theta + i(p, s_p)|p, s_j)$ from the perspective of the traders beyond the usual uncertainty over the fundamental value, because of the "policy reaction risk." The variance amplification occurs because the policy maker has private information, which influences policy and asset value, and traders do not know the policy maker reaction from observing the market clearing price p (unlike the case of an uninformed policy maker, where p completely determines policy). This additional risk in the policy reaction leads the informed—who are risk averse—to trade less aggressively on their private information, i.e., $a > 0$ decreases and price informativeness suffers. If $\psi_4/\psi_3 << 0$, then reducing it further similarly amplifies the return variance and traders react by increasing $a < 0$, i.e., also trade less aggressively and make prices less informative. Bond and Goldstein (2015) also document policy risk (what they call endogenous risk) in their model.

If the share of informed traders μ increases, then the informed trade more aggressively on information. This is because the price becomes more informative with more informed traders, and the policy maker reacts by placing a larger weight on the price rather than her private information when inferring the state θ and making her policy decision. This reduces the return variance from the perspective of the traders, because observing p they can better infer the policy maker reaction. Similarly, less noise affecting the price (larger τ_u) and better information for traders (larger τ_ε) leads to the informed trading more aggressively, again because the policy maker uses her private information less in making policy, which reduces the variance from the perspective of the traders, i.e., reduces "policy risk." Finally, a smaller risk aversion ρ_I leads to more aggressive trading and more informative prices, which is a standard result and not directly related to the policy maker reaction.

The effect of an increased precision of the policy maker information τ_ε^P on price informativeness is potentially ambiguous. On the one hand, more precise policy maker information means the additional information from the financial market is less important, in the sense

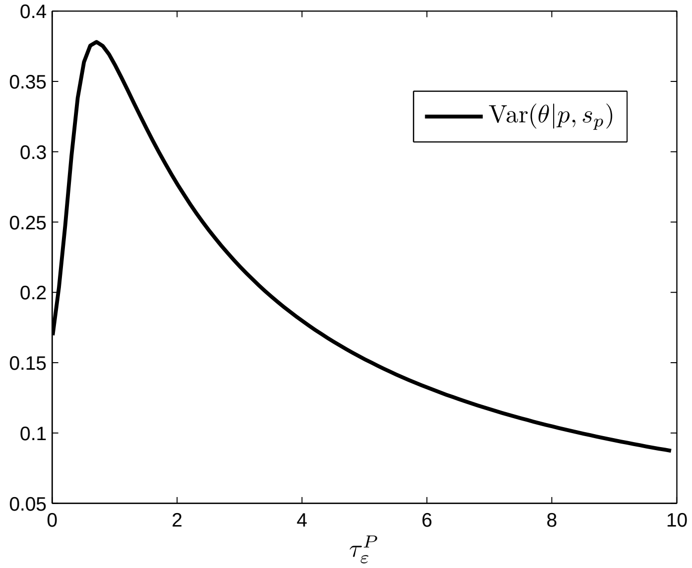

Figure 1.6: Plot of the variance of the policy maker estimate of the state θ, depending on policy maker information precision τ_ε^P and taking into account equilibrium effects on price informativeness. The parameter values of the example are $\mu = 1$, $\rho_I = 1$, $\tau_\theta = 1$, $\tau_\varepsilon = 1$, $\tau_u = 5$, $\psi_4/\psi_3 = 2$.

that it does not reduce the posterior variance over θ for the policy maker as much (direct effect). On the other hand, it can potentially increase $|a|$, i.e., how aggressively the informed trade on information, which in turn increases price informativeness (indirect effect). The proof shows that the direct effect dominates for large enough $|\psi_4/\psi_3|$. The same intuition applies to the prior distribution precision τ_θ.

Interestingly, considering the variance of the policy maker estimate $\mathrm{Var}(\theta|p, s_p)$ rather than price informativeness, more policy maker information is not uniformly better. Taking equilibrium effects into account, i.e., how informed traders change $|a|$ and hence the informational content of market prices, better information for the policy maker (as measured by a lower policy maker signal variance $1/\tau_\varepsilon^P$) can increase the variance of the policy maker estimate. Figure 1.6 plots one example, where the conditional variance is increasing in τ_ε^P for small values of τ_ε^P, because the informed trade less aggressively on information ($|a|$ decreases) due to increased policy risk. However, the figure also shows that τ_ε^P decreases the conditional variance for large enough τ_ε^P, since the gain in private information precision outweighs the loss in financial market information precision eventually.

Policy maker transparency

The comparative statics results showed that price informativeness can be strongly affected by policy risk, i.e., the risk due to the policy maker reaction, since the informed traders are risk averse. To reduce the policy risk, which exists solely because traders do not know the independent information of the policy maker s_p, the policy maker could make her information s_p public. In this case, the expected asset value for informed traders is $\mathbb{E}[\theta + i(p, s_p)|p, s_j, s_p] = \mathbb{E}[\theta|p, s_j, s_p] + i(p, s_p)$, where private information s_j is still useful in predicting part of the asset value. Moreover, the conditional asset variance becomes $\mathrm{Var}(\theta + i(p, s_p)|p, s_j, s_p) = \mathrm{Var}(\theta|p, s_j, s_p) < \mathrm{Var}(\theta + i(p, s_p)|p, s_j)$, so there is no more policy risk, since traders know all factors that determine policy (p, s_p). Therefore, if s_p is publicly disclosed, then trading aggressiveness is $a = \tau_\varepsilon/\rho_I$, as in the standard model without policy maker or with an uninformed policy maker (see above). And exactly the same reasons as in the section with uninformed policy maker explain why the informed have higher net demand for the asset with larger signal s_j.

If the policy maker does not disclose her information as assumed in all other sections so far, then a is implicitly defined as (1.28):

$$a = \frac{\tau_\varepsilon}{\rho_I \left[1 + \frac{\psi_4}{\psi_3} \frac{\tau_\varepsilon}{\tau_\varepsilon + \tau_\theta + (\mu a)^2 \tau_u} + \frac{\left(\frac{\psi_4}{\psi_3}\right)^2 \frac{\tau_\varepsilon}{\tau_\varepsilon + \tau_\theta + (\mu a)^2 \tau_u}}{1 + \frac{\psi_4}{\psi_3} \frac{\tau_\varepsilon}{\tau_\varepsilon + \tau_\theta + (\mu a)^2 \tau_u}}\right]}.$$

Clearly, for large $|\psi_4/\psi_3|$, the equilibrium a becomes very small (see also the proof of Proposition 16), because policy risk severely diminishes the incentives to trade on information for the risk averse traders. Indeed, $a \to 0$ as $|\psi_4/\psi_3| \to \infty$. Hence, a policy maker with preferences for large interventions will be better off by disclosing her information, since it removes policy risk for traders and thereby makes prices more informative.

Bond and Goldstein (2015) find the exact opposite. In their model, the policy maker should never disclose her information s_p about θ, for then prices become completely uninformative. This is because the asset value in their model depends on θ only indirectly via policy, so once the policy is known, traders have no reason to use their information about θ to trade.

1.4 Conclusion

This paper analyzes a general problem where traders have information about a state variable θ that a policy maker needs in order to implement her optimal policy i. A financial market trades an asset whose value depends both the state θ and the policy i. Under which conditions can the equilibrium prices of these assets be revealing, in the sense that trader information (or a noisy signal thereof) is revealed to the policy maker, if traders correctly anticipate the policy maker reaction?

In a model without noise, I show that there are situations where it is impossible to find a price function that both fulfills a no-mispricing condition and reveals trader information to the policy maker. Consequently, since competitive markets require no mispricing, there is no fully revealing equilibrium in these situations. The main result of this section shows that a condition on asset values, policy maker preferences, and trader/policy maker information structures is necessary and sufficient for the existence of fully revealing equilibria. Hence, the main result identifies all situations where we can expect financial markets and prediction markets to work as tools guiding policy makers in making real decisions, and where we cannot. If the condition is not fulfilled, then market prices can be fully revealing only as long as the policy maker does not react to the information contained therein. The main result can be used to design policies or assets that are more supportive of information revelation by traders.

To investigate how these results depend on the assumption of a noiseless financial market, I develop a model where prices are affected by noise shocks. Adapting a new solution approach, I can explicitly solve for equilibria with uninformed policy maker even if reaction and price functions are nonlinear. In a class of equilibria with continuous price functions, I derive a necessary and sufficient condition for the existence of revealing equilibria. Thus, the section provides useful tools for applications that allow for a broad set of policy maker preferences. In specific cases, noise can solve the problem of self-defeating prophecies, because prices do not fully reveal trader information. Hence, traders retain incentives to trade on information, making market prices at least partially revealing. Similar to Grossman and Stiglitz (1980), market prices in these cases can be informative *because* of noise, rather than despite of noise, only the reason why strong-form informational efficiency fails is self-defeating prophecies and not costly information.

A question for future research is how the necessary and sufficient conditions change if there are large strategic traders who can move prices, which may introduce price manipulation motives to trigger different policy reactions and achieve larger trading profits. Moreover, there are many important welfare aspects left to explore. For example, if the ex post optimal policy reaction function leads to a self-defeating prophecy, when can commitment to another reaction function be superior, because it supports information revelation? This question is difficult to answer in cases where a self-defeating prophecy implies equilibrium non-existence, where welfare cannot be evaluated.

Appendix: Proofs

Theorem 2 (Possibility of full revelation and accurate prices). *Suppose the policy maker knows function $p(\boldsymbol{s})$ and maximizes expected utility. Then a fully revealing and accurate price function exists if and only if condition 1 holds.*

Proof. Necessity: Existence of a fully revealing and accurate price function implies that condition 1 holds. First note that full revelation implies that the expected asset value given trader information equals $v(\boldsymbol{s}) = \mathbb{E}[a(\theta, i(\boldsymbol{s}))|\boldsymbol{s}]\ \forall \boldsymbol{s} \in \boldsymbol{S}$. As shown in Lemma 1, if prices are fully revealing, then accuracy implies $p(\boldsymbol{s}) = v(\boldsymbol{s})$.

Full revelation with accurate prices by definition implies

$$|\{t \in \boldsymbol{S} : \Pr(\boldsymbol{s} = t|p(\boldsymbol{s}) = p, s_p) > 0\}| \leq 1 \; \forall s_p, \forall p. \tag{1.16}$$

By Bayes' rule and the law of iterated expectations,

$$\begin{aligned}
\Pr(\boldsymbol{s} = t|p(\boldsymbol{s}) = p, s_p) &= \frac{\Pr(\boldsymbol{s} = t) \cdot \Pr(p(\boldsymbol{s}) = p, s_p|\boldsymbol{s} = t)}{\Pr(p(\boldsymbol{s}) = p, s_p)} \\
&= \frac{\Pr(\boldsymbol{s} = t) \cdot \mathbb{E}[\Pr(p(\boldsymbol{s}) = p, s_p)|\boldsymbol{s} = t]}{\Pr(p(\boldsymbol{s}) = p, s_p)} \\
&= \frac{\Pr(\boldsymbol{s} = t) \cdot \mathbb{E}[\Pr(s_p|p(\boldsymbol{s}) = p) \cdot \Pr(p(\boldsymbol{s}) = p)|\boldsymbol{s} = t]}{\Pr(p(\boldsymbol{s}) = p, s_p)} \\
&= \frac{\Pr(\boldsymbol{s} = t) \cdot \mathbf{1}\{p(t) = p\} \cdot \Pr(s_p|\boldsymbol{s} = t)}{\Pr(p(\boldsymbol{s}) = p, s_p)}
\end{aligned} \tag{1.17}$$

First, if $\Pr(p(\boldsymbol{s}) = p|\boldsymbol{s} = t) = \mathbf{1}\{p(t) = p\} > 0$ for exactly one $t \in \boldsymbol{S}$ for all $p \in \text{Image}(v(\boldsymbol{s}))$, then $v(\boldsymbol{s})$ is invertible and the price alone fully reveals $\boldsymbol{s}$. Hence $\nexists p \in \text{Image}(v(\boldsymbol{s})) : |v^{-1}(p)| > 1$, i.e., the antecedent of condition 1 is false and the condition holds.

Second, if there exists $p \in \text{Image}(v(\boldsymbol{s}))$ such that $\Pr(p(\boldsymbol{s}) = p|\boldsymbol{s} = t) > 0$ for more than one $t \in \boldsymbol{S}$, then $|X| > 0$, i.e., the antecedent of condition 1 is true. Thus, for condition 1 to hold, the consequent must be true as well. Since there is full revelation, i.e., (1.16) holds, we must have $\Pr(s_p|\boldsymbol{s} = t) \cdot \Pr(s_p|\boldsymbol{s} = t') = 0 \; \forall s_p \in S_p, \forall t \neq t' \in v^{-1}(p), \forall p \in X$, so that $\Pr(\boldsymbol{s} = t|p(\boldsymbol{s}) = p, s_p) > 0$ for at most one $t \in v^{-1}(p)$ for each $p \in \text{Image}(v(\boldsymbol{s}))$. Thus, the consequent is true and condition 1 holds.

Sufficiency: Condition 1 implies the existence of a fully revealing and accurate price function. Condition 1 is true either if the antecedent is false, or if both antecedent and consequent are true. First, if the antecedent of condition 1 is false, then $v(\boldsymbol{s})$ is invertible. In this case, the price function $p(\boldsymbol{s}) = v(\boldsymbol{s})$ is both fully revealing and accurate.

Second, consider the case where the antecedent and consequent of condition 1 is true. The consequent is $\Pr(s_p|\boldsymbol{s} = t) \cdot \Pr(s_p|\boldsymbol{s} = t') = 0 \; \forall s_p \in S_p, \; \forall t \neq t' \in v^{-1}(p), \; \forall p \in X$, which implies that $\Pr(\boldsymbol{s} = t|p(\boldsymbol{s}) = p, s_p)$ in (1.17) is positive for exactly one $t \in \boldsymbol{S}$ for every $p \in \text{Image}(v(\boldsymbol{s}))$, i.e., the price function $p(\boldsymbol{s}) = v(\boldsymbol{s})$ is both fully revealing and accurate. □

Proposition 7 (Asset design and full revelation). *Consider the class of assets that can be written as $a(\theta, i) = A(o(\theta, i))$ for any $o : \Theta \times I \to \mathbb{R}$. If a non-invertible function $A' \in \mathcal{A}'$ allows full revelation and accurate prices, then so does any invertible $A \in \mathcal{A}$, but the converse does not hold.*

Proof. Define $v(\boldsymbol{s}, A) := \mathbb{E}_\theta[A(o(\theta, i(\boldsymbol{s}, s_p)))|\boldsymbol{s}]$. We need to show that there is no pair $\boldsymbol{s} \neq \boldsymbol{s}'$ such that $v(\boldsymbol{s}, A') \neq v(\boldsymbol{s}', A')$ but $v(\boldsymbol{s}, A) = v(\boldsymbol{s}', A)$, i.e., whenever a pair of states is revealed solely via prices under $A' \in \mathcal{A}'$, then it must also be revealed solely by prices under all $A \in \mathcal{A}$. Formally,

$$v(\boldsymbol{s}, A') \neq v(\boldsymbol{s}', A') \implies v(\boldsymbol{s}, A) \neq v(\boldsymbol{s}', A) \ \forall A \in \mathcal{A}, \ \forall A' \in \mathcal{A}'. \tag{1.18}$$

If this condition holds, then if condition 1 holds under A', it will also hold under any A, since a price function $p(\boldsymbol{s}) = v(\boldsymbol{s}, A)$ of asset A must always rule out at least as many states $\boldsymbol{s} \in \boldsymbol{S}$ as a price function $p'(\boldsymbol{s}) = v(\boldsymbol{s}, A')$ of asset A'.

To show (1.18) holds, note that for any $A' \in \mathcal{A}'$,

$$\begin{aligned} &\mathbb{E}_\theta[A'(o(\theta, i(\boldsymbol{s}, s_p)))|\boldsymbol{s}] \neq \mathbb{E}_\theta[A'(o(\theta, i(\boldsymbol{s}', s_p)))|\boldsymbol{s}'] \\ &\implies \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}, s_p))|\boldsymbol{s}] \neq \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}', s_p))|\boldsymbol{s}']. \end{aligned} \tag{1.19}$$

Moreover, because $\mathcal{A}$ is the set of invertible functions,

$$\begin{aligned} &\mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}, s_p))|\boldsymbol{s}] \neq \mathbb{E}_\theta[o(\theta, i(\boldsymbol{s}', s_p))|\boldsymbol{s}'] \\ &\implies \mathbb{E}_\theta[A(o(\theta, i(\boldsymbol{s}, s_p)))|\boldsymbol{s}] \neq \mathbb{E}_\theta[A(o(\theta, i(\boldsymbol{s}', s_p)))|\boldsymbol{s}'] \ \forall A \in \mathcal{A}. \end{aligned} \tag{1.20}$$

Combining (1.19) and (1.20) results in (1.18), which as argued above implies that condition 1 holds for any $A \in \mathcal{A}$ whenever it holds for some $A' \in \mathcal{A}'$.

We still need to show that $A \in \mathcal{A}$ can allow for full revelation and accurate prices but $A' \in \mathcal{A}'$ might not. This immediately follows from the fact that invertibility of A is necessary (but not sufficient) for the invertibility of $v(\boldsymbol{s}) = \mathbb{E}_\theta[A(o(\theta, i(\boldsymbol{s}, s_p)))|\boldsymbol{s}]$, which guarantees that condition 1 holds. And according to Theorem 2, full revelation and accurate prices are possible if and only if condition 1 holds. □

Proposition 8. *An equilibrium in the class of quasi-linear equilibria exists if and only if there exists a continuous reaction function $i(z) \in \arg\max_i \mathbb{E}[U(\theta, i)|z]$ such that, for $a = \tau_\varepsilon/\rho_I$,*

$$\begin{aligned} z(1 + \mu^2 a\tau_u/\rho_I + (1-\mu)\mu a\tau_u/\rho_U) + i(z)[\mu(\tau_\theta + \tau_\varepsilon + (\mu a)^2\tau_u)/(\rho_I\mu a) \\ +(1-\mu)(\tau_\theta + (\mu a)^2\tau_u)/(\rho_U\mu a)] \end{aligned}$$

is invertible in z. If it exists, then the equilibrium is unique in the class of quasi-linear equilibria and market prices fully reveal the noisy signal $z = \theta + u/(\mu a)$.

The trader equilibrium strategies are $X_I(s_j, p) = as_j - g_I(p)$ and $X_U(p) = -g_U(p)$ with $a = \tau_\varepsilon/\rho_I$, $g_I(p)$ given in (1.9)*, and $g_U(p)$ given in* (1.11)*. The equilibrium policy reaction is $i(P^{-1}(p)) = i(z)$. The equilibrium price function $P(\theta, u) = P(z)$ is given in* (1.13)*.*

Proof. Conjecturing demand functions of the quasi-linear form in definition 5, the market clearing condition is (1.7). The market clearing condition has to hold for any realization of (θ, u), which appears directly only on the right hand side of (1.7). Since (θ, u) appears on the left hand side only indirectly via the equilibrium price function $P(\theta, u)$, this price function must at least reveal the statistic $z = \theta + u/(\mu a)$, so that I now write $P(z)$. That is, all (θ, u) realizations with $p = P(\theta, u)$ must be on the line $z = \theta + u/(\mu a)$. Hence, any equilibrium price function reveals at least the realization of z to the uninformed, which includes the policy maker. Below, I shall also confirm that the price function reveals at most z. Invertibility of $P(z)$ also implies that there is a unique market clearing price for every realization of z.

Using the information set for the uninformed $\{z\}$ and the information set for the informed $\{s_j, z\}$, the expected asset values are normally distributed conditional on this information. Consequently, the optimal demand of the informed is uniquely determined and given by (1.8), which is of the conjectured quasi-linear form. Matching coefficient a, $a = \tau_\varepsilon/\rho_I$. Matching g_I yields (1.9). Similarly, the optimal demand for the uninformed is uniquely given by (1.10), and matching g_U yields (1.11). Consequently, given that z is revealed by prices and that $i(z)$ is unique, the equilibrium must be unique in the class of quasi-linear equilibria if it exists, since demand functions and hence the price function (via market clearing condition) is unique. Substituting $g_I(p)$ and $g_U(p)$ into the market clearing condition (1.7) yields (1.12). Given the assumption that $P(z)$ is invertible, rewrite $P^{-1}(P(z)) = z$ and, abusing notation, $i(P(z)) = i(z)$, so that the market clearing condition can be rearranged for an explicit expression of $P(z)$ given in (1.13). Dropping constant factors and terms from the right hand side of (1.13), it is immediate that $P(z)$ is invertible as required if and only if (1.14) is invertible in z. If (1.14) is non-invertible, it contradicts market clearing and no equilibrium exists.

It remains to be shown that the continuous equilibrium price function $P(\theta, u)$ does not reveal more than the realization of z, i.e., $P(\theta, u)$ depends on (θ, u) only via $\theta + u/(\mu a)$. Lemma 2 in Pálvölgyi and Venter (2015) shows this for the Grossman and Stiglitz (1980)-model, and the proof can be applied directly to the problem here. For completeness, I will translate their proof into my notation.

I already established that all (θ, u) realizations for which $p = P(\theta, u)$ are on the line $\theta + u/(\mu a)$ (if not, then $P(\theta, u)$ would not be invertible in z as required for market clearing). Now, by contradiction, suppose $P(\theta, u)$ depends on (θ, u) not only via $z = \theta + u/(\mu a)$. This implies there exist two pairs $(\theta_1, u_1) \neq (\theta_2, u_2)$ such that $z = \theta_1 + u_1/(\mu a) = \theta_2 + u_2/(\mu a)$ with $P(\theta_1, u_1) = p_1 \neq P(\theta_2, u_2) = p_2$. Now, given continuity of $P(\theta, u)$ and using the intermediate value theorem, we can find a pair (θ^*, u^*) with $\theta^* + u^*/(\mu a) = z$ and $P(\theta^*, u^*) = (p_1 + p_2)/2$ (e.g., increase θ and decrease u). Similarly, we can find a pair (θ', u') with $\theta' + u'/(\mu a) \neq z$ and $P(\theta', u') = (p_1 + p_2)/2$ (e.g., starting at $\min\{p_1, p_2\}$, by increasing θ and keeping u constant). That is, we have two points with $P(\theta^*, u^*) = P(\theta', u') = (p_1 + p_2)/2$ and $\theta^* + u^*/(\mu a) \neq \theta' + u'/(\mu a)$. Yet this contradicts the previously established fact that all (θ, u) realizations for which $p = P(\theta, u)$ are on the same line $\theta + u/(\mu a)$. Consequently, a continuous price function $P(\theta, u)$ reveals exactly $z = \theta + u/(\mu a)$ in equilibrium. $\square$

Proposition 13 (Equilibrium existence). *If the policy maker is informed, then a linear noisy REE exists for almost all parameter values.*

Proof. The solution approach for the noisy REE is the typical "guess and verify" approach: First, a conjecture about the shape of demand functions is made. In this case, we conjecture the demand functions to be linear in signal s_j and price p, which—after imposing market clearing—gives a linear price function $P(\theta, u)$ with undetermined coefficients. Second, according to this price conjecture, the price function $P(\theta, u)$ gives the relationship between state θ and price, which is used to update traders' beliefs about θ via Bayes' rule. Third, demand functions given the information sets are computed. Fourth, the undetermined coefficients are identified, which gives the actual relationship between θ and prices.

I am going to derive a symmetric linear noisy rational expectations equilibrium, where the conjecture is that traders use strategies

$$X_I(p, s_j) = as_j - c_I p + b_I \text{ and } X_U(p) = -c_U p + b_U,$$

which yields the market clearing condition

$$\begin{aligned} \int_0^\mu (as_j - c_I p + b_I)dj + \int_\mu^1 (-c_U p + b_U)dj + u = 0 \\ \iff p = \frac{\mu a\theta + \mu b_I + (1-\mu)b_U + u}{\mu c_I + (1-\mu)c_U}, \end{aligned} \tag{1.21}$$

because an appropriate law of large numbers for i.i.d. random variables (Sun, 2006) yields $\int_0^\mu s_j dj = \mu\theta$. Define $\lambda \coloneqq (\mu c_I + (1-\mu)c_U)^{-1}$ and $\tilde{b} \coloneqq \mu b_I + (1-\mu)b_U$ to simplify notation, so that the market clearing condition is

$$p = \lambda(\mu a\theta + \tilde{b} + u).$$

Rearranging, the following linear statistic of the market clearing price p, which is informationally equivalent to p, is the state θ plus a normally distributed noise term $u/(\mu a)$,

$$\frac{p - \lambda\tilde{b}}{\lambda\mu a} = \theta + u/(\mu a).$$

The variance of the net asset value for informed traders conditional on their information is

$$\begin{aligned} \mathrm{Var}(\theta + i(p, s_p) - p|p, s_j) = \mathrm{Var}(\theta + i(p, s_p)|p, s_j) = \mathrm{Var}(\theta + \beta_1 + \beta_2 p + \beta_3 s_p|p, s_j) \\ = \mathrm{Var}(\theta(1+\beta_3) + \beta_3\varepsilon|p, s_j) = \frac{(1+\beta_3)^2}{\tau_\varepsilon + \tau_\theta + (\mu a)^2\tau_u} + \frac{\beta_3^2}{\tau_\varepsilon}, \end{aligned} \tag{1.22}$$

since $\mathrm{Var}(\theta|p, s_j) = 1/(\tau_\varepsilon + \tau_\theta + (\mu a)^2\tau_u)$, as in the standard model without policy maker, and θ and ε are independent. Similarly, for the uninformed, the variance of the asset value given the price is

$$\mathrm{Var}(\theta + i(p, s_p) - p|p) = \frac{(1+\beta_3)^2}{\tau_\theta + (\mu a)^2\tau_u} + \frac{\beta_3^2}{\tau_\varepsilon}.$$

The Bayesian updating rule for the mean of normal distributions is a precision weighted sum of prior mean and signals, where the precision of the price signal is the inverse of its conditional variance, hence the conditional expectation of the asset value is

$$\begin{aligned}\mathbb{E}[\theta + i(p, s_p) - p)|p, s_j] &= \mathbb{E}[\theta(1+\beta_3) + \beta_3\varepsilon)|p, s_j] + \beta_1 + \beta_2 p - p \\ &= (1+\beta_3)\frac{\tau_\theta\bar{\theta} + \tau_\varepsilon s_j + (\mu a)^2\tau_u\left(\frac{p-\lambda\tilde{b}}{\lambda\mu a}\right)}{\tau_\theta + \tau_\varepsilon + (\mu a)^2\tau_u} + \beta_1 + \beta_2 p - p.\end{aligned}$$

Similarly, the expectation of the asset value for uninformed traders is

$$\mathbb{E}[\theta + i(p, s_p) - p)|p] = (1+\beta_3)\frac{\tau_\theta\bar{\theta} + (\mu a)^2\tau_u\left(\frac{p-\lambda\tilde{b}}{\lambda\mu a}\right)}{\tau_\theta + (\mu a)^2\tau_u} + \beta_1 + \beta_2 p - p.$$

The well-known CARA demand functions derived from the first order conditions are given by

$$\begin{aligned}X_I(p, s_j) &= \frac{\mathbb{E}[\theta + i(p, s_p) - p)|p, s_j]}{\rho_I \mathrm{Var}(\theta + i(p, s_p) - p|p, s_j)}, \\ X_U(p) &= \frac{\mathbb{E}[\theta + i(p, s_p) - p)|p]}{\rho_U \mathrm{Var}(\theta + i(p, s_p) - p|p)}.\end{aligned}$$

Note that the trader objective is concave even if $i(p)$ is highly convex, because a single trader does not affect p, hence $i(p)$ is a constant in the trader maximization problem. Plugging in for conditional expectations and variances,

$$X_U(p) = \frac{(1+\beta_3)\frac{\tau_\theta\bar{\theta} + (\mu a)^2\tau_u\left(\frac{p-\lambda\tilde{b}}{\lambda\mu a}\right)}{\tau_\theta + (\mu a)^2\tau_u} + \beta_1 + \beta_2 p - p}{\rho_U\left[\frac{(1+\beta_3)^2}{\tau_\theta + (\mu a)^2\tau_u} + \frac{\beta_3^2}{\tau_\varepsilon}\right]},$$

which is linear in p as conjectured, hence identifying coefficients using $X_U(p) = -c_U p + b_U$ we obtain

$$c_U = \frac{1 - \beta_2 - \frac{(1+\beta_3)\mu a\tau_u/\lambda}{\tau_\theta + (\mu a)^2\tau_u}}{\rho_U\left[\frac{(1+\beta_3)^2}{\tau_\theta + (\mu a)^2\tau_u} + \frac{\beta_3^2}{\tau_\varepsilon}\right]},\ b_U = \frac{\beta_1 + \frac{(1+\beta_3)(\tau_\theta\bar{\theta} - \mu a\tau_u\tilde{b})}{\tau_\theta + (\mu a)^2\tau_u}}{\rho_U\left[\frac{(1+\beta_3)^2}{\tau_\theta + (\mu a)^2\tau_u} + \frac{\beta_3^2}{\tau_\varepsilon}\right]}. \tag{1.23}$$

For informed traders, plugging in for conditional expectations and variances,

$$X_I(p, s_j) = \frac{(1+\beta_3)\frac{\tau_\theta\bar{\theta} + \tau_\varepsilon s_j + (\mu a)^2\tau_u\left(\frac{p-\lambda\tilde{b}}{\lambda\mu a}\right)}{\tau_\theta + \tau_\varepsilon + (\mu a)^2\tau_u} + \beta_1 + \beta_2 p - p}{\rho_U\left[\frac{(1+\beta_3)^2}{\tau_\varepsilon + \tau_\theta + (\mu a)^2\tau_u} + \frac{\beta_3^2}{\tau_\varepsilon}\right]},$$

which is linear in signal s_j and price p as conjectured. Matching coefficients of $X_I(p, s_j) = as_j - c_I p + b_I$,

$$a = \frac{(1+\beta_3)\tau_\varepsilon}{\rho_I(\tau_\theta + \tau_\varepsilon + (\mu a)^2\tau_u)\left[\frac{(1+\beta_3)^2}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u} + \frac{\beta_3^2}{\tau_\varepsilon}\right]},$$
$$c_I = \frac{1 - \beta_2 - \frac{(1+\beta_3)\mu a\tau_u/\lambda}{\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u}}{\rho_I\left[\frac{(1+\beta_3)^2}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u} + \frac{\beta_3^2}{\tau_\varepsilon}\right]}, \; b_U = \frac{\beta_1 + \frac{(1+\beta_3)(\tau_\theta\bar{\theta} - \mu a\tau_u\tilde{b})}{\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u}}{\rho_I\left[\frac{(1+\beta_3)^2}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u} + \frac{\beta_3^2}{\tau_\varepsilon}\right]}. \tag{1.24}$$

Now, both c_I and c_U depend on λ and vice versa. Substitute c_I and c_U into $\lambda = 1/(\mu c_I + (1-\mu)c_U)$ and solve for λ to get

$$\lambda = \frac{1 + \frac{(1+\beta_3)\mu^2 a\tau_u}{\rho_I(\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u)\left[\frac{(1+\beta_3)^2}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u}+\frac{\beta_3^2}{\tau_\varepsilon}\right]} + \frac{(1+\beta_3)\mu(1-\mu)\tau_u}{\rho_U(\tau_\theta+(\mu a)^2\tau_u)\left[\frac{(1+\beta_3)^2}{\tau_\theta+(\mu a)^2\tau_u}+\frac{\beta_3^2}{\tau_\varepsilon}\right]}}{\frac{\mu(1-\beta_2)}{\rho_I\left[\frac{(1+\beta_3)^2}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u}+\frac{\beta_3^2}{\tau_\varepsilon}\right]} + \frac{(1-\mu)(1-\beta_2)}{\rho_U\left[\frac{(1+\beta_3)^2}{\tau_\theta+(\mu a)^2\tau_u}+\frac{\beta_3^2}{\tau_\varepsilon}\right]}}. \tag{1.25}$$

Furthermore, both b_I and b_U depend on $\tilde{b}$ and vice versa. Substitute both into $\tilde{b} = \mu b_I + (1-\mu)b_U$ and solve for $\tilde{b}$ to get

$$\tilde{b} = \frac{\mu\left[\frac{\frac{(1+\beta_3)\tau_\theta\bar{\theta}}{\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u}+\beta_1}{\rho_I\left[\frac{(1+\beta_3)^2}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u}+\frac{\beta_3^2}{\tau_\varepsilon}\right]}\right] + (1-\mu)\left[\frac{\frac{(1+\beta_3)\tau_\theta\bar{\theta}}{\tau_\theta+(\mu a)^2\tau_u}+\beta_1}{\rho_U\left[\frac{(1+\beta_3)^2}{\tau_\theta+(\mu a)^2\tau_u}+\frac{\beta_3^2}{\tau_\varepsilon}\right]}\right]}{1 + \frac{\mu^2 a\tau_u(1+\beta_3)}{\rho_I(\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u)\left[\frac{(1+\beta_3)^2}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u}+\frac{\beta_3^2}{\tau_\varepsilon}\right]} + \frac{(1-\mu)\mu a\tau_u(1+\beta_3)}{\rho_U(\tau_\theta+(\mu a)^2\tau_u)\left[\frac{(1+\beta_3)^2}{\tau_\theta+(\mu a)^2\tau_u}+\frac{\beta_3^2}{\tau_\varepsilon}\right]}}. \tag{1.26}$$

Turning to the policy maker, her utility function is strictly concave by construction, so the first order condition determines the unique reaction function

$$i(p, s_p) = \frac{\psi_2 + \psi_4\mathbb{E}[\theta|p, s_p]}{\psi_3} = \frac{\psi_2 + \psi_4\frac{\tau_\theta\bar{\theta}+\tau_\varepsilon s_p+(\mu a)^2\tau_u\left(\frac{p-\lambda\tilde{b}}{\mu a\lambda}\right)}{\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u}}{\psi_3}.$$

Matching coefficients for $i(p, s_p) = \beta_1 + \beta_2 p + \beta_3 s_p$, we obtain

$$\beta_1 = \frac{\psi_2 + \psi_4\frac{\tau_\theta\bar{\theta}-\mu a\tau_u\tilde{b}}{\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u}}{\psi_3}, \; \beta_2 = \frac{\psi_4\frac{\mu a\tau_u/\lambda}{\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u}}{\psi_3}, \; \beta_3 = \frac{\psi_4\frac{\tau_\varepsilon}{\tau_\theta+\tau_\varepsilon+(\mu a)^2\tau_u}}{\psi_3}. \tag{1.27}$$

Substitute the term for β_3 in (1.27) into coefficient a (1.24), so that the equilibrium condition for a is

$$a = \frac{\tau_\varepsilon}{\rho_I\left[1 + \frac{\psi_4}{\psi_3}\frac{\tau_\varepsilon}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u} + \frac{\left(\frac{\psi_4}{\psi_3}\right)^2\frac{\tau_\varepsilon}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u}}{1+\frac{\psi_4}{\psi_3}\frac{\tau_\varepsilon}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u}}\right]}, \tag{1.28}$$

which depends only on one endogenous strategy variable (a). The condition can be rearranged as a fifth degree polynomial in a. Fifth degree polynomials are guaranteed to have at least one solution and at most five. Given a, β_3 in (1.27) is uniquely determined. Given β_3 and a, substituting β_2 into λ in (1.25) yields a linear condition with a unique solution for λ, which in turn determines β_2 in (1.27) and c_U, c_I in (1.23), (1.24) uniquely. Substituting β_1 into $\tilde{b}$ in (1.26) yields a linear condition with a unique solution, which in turn determines β_1 uniquely. Finally, market clearing requires $\mu c_I + (1-\mu)c_U \neq 0$ for the coefficients c_I, c_U just obtained. Thus, an equilibrium exists for almost all parameter profiles, and a is the only source of equilibrium multiplicity (i.e., all other endogenous variables are uniquely determined given a). □

Proposition 16 (**Equilibrium uniqueness**).

i. There exists $r^* > 0$ such that, for all $\psi_4/\psi_3 > r^*$, the linear equilibrium is unique with $a > 0$.

ii. There exists $r^* < 0$ such that, for all $\psi_4/\psi_3 < r^*$, the linear equilibrium is unique with $a < 0$.

iii. There exists $\rho^* > 0$ such that, for all $\rho_I > \rho^*$, the equilibrium is unique.

Proof. *i.* Using the shorthand $r := \psi_4/\psi_3$, recall equilibrium condition (1.28):

$$a = \frac{\tau_\varepsilon}{\rho_I\left[1 + r\frac{\tau_\varepsilon}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u} + \frac{r^2\frac{\tau_\varepsilon}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u}}{1+r\frac{\tau_\varepsilon}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u}}\right]}. \tag{1.28}$$

Every a that fulfills condition (1.28) forms an equilibrium; uniqueness requires a single intersection of a and the right hand side (RHS) of the condition. Note that the RHS of (1.28) is symmetric about $a = 0$, since a enters the condition only quadratically. Note also that the RHS of (1.28) is positive for all $a \geq 0$ at all $r > 0$. Using the shorthand $\tau := \tau_\varepsilon + \tau_\theta + (\mu a)^2\tau_u$, the slope of the RHS of (1.28) in a is

$$\frac{\partial \mathrm{RHS}(a)}{\partial a} = \tau_\varepsilon\rho_I^{-1}[2r\tau_\varepsilon\tau_u\mu^2 a\tau^{-2} + 2r^2\tau_\varepsilon\tau_u\mu^2 a(\tau + r\tau_\varepsilon)^{-2}] / \left[1 + r\frac{\tau_\varepsilon}{\tau} + \frac{r^2\tau_\varepsilon}{\tau + r\tau_\varepsilon}\right]^2. \tag{1.29}$$

This slope is positive for $r > 0$ and approaches zero as $r \to \infty$ at all $a > 0$, since the denominator grows at a quadratic rate in r, while the numerator only grows at a linear rate. Now, at $a = 0$, $\mathrm{RHS}(a) > a = 0$. Since a large r makes the slope of $\mathrm{RHS}(a)$ arbitrarily small, there exists some $r^* > 0$ such that the slope of $\mathrm{RHS}(a)$ is less than 1 for all $r \geq r^*$, which guarantees that $\mathrm{RHS}(a)$ crosses a exactly once in $a \geq 0$, i.e., (1.28) has only one non-negative solution. It remains to be shown that there is no further solution $a = \mathrm{RHS}(a)$ in $a < 0$. This immediately follows from the fact that $\mathrm{RHS}(a) > 0$ for all $a \in \mathbb{R}$ if $r > 0$.

ii. We want to show that there exists $r^* < 0$ such that $r < r^*$ guarantees a unique solution to (1.28), and that this solution is $a < 0$. The RHS of the equilibrium condition (1.28) is negative for $a < 0$ close to zero if and only if $1 + r\frac{\tau_\varepsilon}{\tau_\varepsilon+\tau_\theta+(\mu a)^2\tau_u} < 0$, which holds with $r < -\frac{\tau_\varepsilon+\tau_\theta}{\tau_\varepsilon}$. Thus, for negative enough $r < 0$, $\text{RHS}(a) < 0$ close to $a = 0$.

The slope of the RHS (1.29) approaches zero as $r \to -\infty$ at all $a \in \mathbb{R}$, since the denominator grows at a quadratic rate whereas the numerator changes only at a linear rate in r. Thus, there exists some $r^* < 0$ such that the slope of $\text{RHS}(a)$ is smaller than 1 and $\text{RHS}(a = 0) < 0$ for all $r \leq r^*$, which guarantees that the RHS intersects a exactly once for $a \leq 0$ and does not intersect for $a > 0$.

iii. The slope of the RHS in a given in (1.29) approaches zero as $\rho_I \to \infty$. The same arguments as in (*i.*) and (*ii.*) imply that we can find a threshold $\rho^* > 0$ such that, for all $\rho_I > \rho^*$, the slope of the RHS is less than 1 everywhere, which guarantees a unique equilibrium. □

Proposition 17 **(Comparative statics).**

i. *For $\psi_4/\psi_3 > 0$ sufficiently large,*

- *decreasing ψ_4/ψ_3, policy maker information precision τ_ε^P, prior distribution precision τ_θ, or informed trader risk aversion ρ_I, and*
- *increasing the share of informed traders μ, noise precision τ_u, or trader signal precision τ_ε*

increases price informativeness (1.4) *for the policy maker in the unique linear equilibrium. Moreover, all of these except for ψ_4/ψ_3 unambiguously increase the welfare gain due to financial market information* (1.6).

ii. *For $\psi_4/\psi_3 < 0$ sufficiently negative,*

- *decreasing policy maker information precision τ_ε^P, prior distribution precision τ_θ, or informed trader risk aversion ρ_I, and*
- *increasing ψ_4/ψ_3, the share of informed traders μ, noise precision τ_u, or trader signal precision τ_ε*

increases price informativeness in the unique linear equilibrium. Moreover, all of these except for ψ_4/ψ_3 unambiguously increase the welfare gain due to financial market information.

Proof. We can rewrite the equilibrium condition (1.28), which implicitly defines a, as $\text{RHS}(a, t) - a = 0$, where $\text{RHS}(a, t)$ is the right hand side of (1.28) as a function of a

and parameter t. By the implicit function theorem, we can determine how a changes in equilibrium in response to a small change in parameter t:

$$\frac{\partial a}{\partial t} = -\frac{\mathrm{RHS}_t}{\mathrm{RHS}_a - 1},$$

where RHS_t is the partial derivative with respect to t. The uniqueness proofs of Proposition 16 show that RHS_a is less than 1 for $|\psi_4/\psi_3|$ large enough, hence it follows that a increases if and only if $\mathrm{RHS}_t > 0$ and decreases if and only if $\mathrm{RHS}_t < 0$.

i. Using the shorthands $r := \psi_4/\psi_3$ and $\tau := \tau_\varepsilon + \tau_\theta + (\mu a)^2\tau_u$,

$$\mathrm{RHS}_r = \underbrace{-\tau_\varepsilon\rho_I^{-1}}_{<0}\underbrace{\left[1 + r\tau_\varepsilon/\tau + \frac{r^2\tau_\varepsilon/\tau}{1 + r\tau_\varepsilon/\tau}\right]^{-2}}_{>0}\left[\tau_\varepsilon/\tau + \frac{2r\tau_\varepsilon/\tau + r^2(\tau_\varepsilon/\tau)^2}{(1 + r\tau_\varepsilon/\tau)^2}\right], \tag{1.30}$$

which is negative if and only if the last bracket is positive. Requiring the term to be positive is equivalent to the condition

$$r^2(\tau_\varepsilon/\tau + (\tau_\varepsilon/\tau)^2) + 2r(1 + \tau_\varepsilon/\tau) + 1 > 0, \tag{1.31}$$

which holds for all $r \geq 0$. Thus, $\mathrm{RHS}_r < 0$ and a decreases with larger r.

Next, RHS_μ, which is positive for $r > 0$:

$$\mathrm{RHS}_\mu = \underbrace{-\tau_\varepsilon\rho_I^{-1}}_{<0}\underbrace{\left[1 + r\tau_\varepsilon/\tau + \frac{r^2\tau_\varepsilon/\tau}{1 + r\tau_\varepsilon/\tau}\right]^{-2}}_{>0}\underbrace{\left[-2r\tau_\varepsilon a^2\mu\tau_u/\tau^2 - 2r^2\tau_\varepsilon a^2\mu\tau_u/(\tau + r\tau_\varepsilon)^2\right]}_{<0 \text{ if } r>0}. \tag{1.32}$$

Moreover, $\mathrm{RHS}_{\tau_\theta} > 0$:

$$\mathrm{RHS}_{\tau_\theta} = \underbrace{-\tau_\varepsilon\rho_I^{-1}}_{<0}\underbrace{\left[1 + r\tau_\varepsilon/\tau + \frac{r^2\tau_\varepsilon/\tau}{1 + r\tau_\varepsilon/\tau}\right]^{-2}}_{>0}\underbrace{\left[-r\tau_\varepsilon/\tau^2 - r^2\tau_\varepsilon/(\tau + r\tau_\varepsilon)^2\right]}_{<0 \text{ if } r>0}. \tag{1.33}$$

It can similarly be verified that $\mathrm{RHS}_{\rho_I} < 0$ and $\mathrm{RHS}_{\tau_u} > 0$.

To investigate the effects of τ_ε and τ_ε^P separately, we have to derive the equilibrium condition for a again, which is more complicated in the more general case:

$$a = \frac{\left(1 + r\frac{\tau_\varepsilon^P}{\tau_\theta + \tau_\varepsilon^P + (\mu a)^2\tau_u}\right)\tau_\varepsilon}{\rho_I\left[\left(1 + r\frac{\tau_\varepsilon^P}{\tau_\theta + \tau_\varepsilon^P + (\mu a)^2\tau_u}\right)^2 + (\tau_\theta + \tau_\varepsilon + (\mu a)^2\tau_u)\left(r\frac{\tau_\varepsilon^P}{\tau_\theta + \tau_\varepsilon^P + (\mu a)^2\tau_u}\right)^2/\tau_\varepsilon^P\right]}. \tag{1.34}$$

This is also a fifth degree polynomial in a, guaranteeing a solution. A tedious but straightforward computation using the quotient rule shows that the slope of the RHS in

a gets arbitrarily close to zero for $|r| \to \infty$, since the denominator of RHS_a grows at a quartic rate while numerator grows only quadratically, thus we also have uniqueness for $|r|$ large enough in this more general case. It remains to determine the sign of the slope of the RHS in τ_ε and τ_ε^P for the comparative statics. Abbreviating $\beta_3 = r\frac{\tau_\varepsilon^P}{\tau_\theta+\tau_\varepsilon^P+(\mu a)^2\tau_u}$,

$$\begin{aligned}\frac{\partial \text{RHS}(1.34)}{\partial \tau_\varepsilon} &= \frac{(1+\beta_3)\rho_I[(1+\beta_3)^2 + \tau\beta_3^2/\tau_\varepsilon^P)] - (1+\beta_3)\tau_\varepsilon\rho_I\beta_3^2/\tau_\varepsilon^P}{[\rho_I((1+\beta_3)^2+\tau\beta_3^2/\tau_\varepsilon^P)]^2} \\ &= \frac{(1+\beta_3)\rho_I[(1+\beta_3)^2 + (\tau_\theta + (\mu a)^2\tau_u)\beta_3^2/\tau_\varepsilon^P)]}{[\rho_I((1+\beta_3)^2+\tau\beta_3^2/\tau_\varepsilon^P)]^2} > 0,\end{aligned} \tag{1.35}$$

since $r > 0$ implies $\beta_3 > 0$, thus $a > 0$ in equilibrium increases with an increase of τ_ε. Finally, denoting $\beta_3' = \partial\beta_3/\partial\tau_\varepsilon^P = r\frac{\tau_\theta+(\mu a)^2\tau_u}{(\tau_\varepsilon^P+\tau_\theta+(\mu a)^2\tau_u)^2} > 0$,

$$\begin{aligned}\frac{\partial \text{RHS}(1.34)}{\partial \tau_\varepsilon^P} > 0 \iff & \beta_3'\tau_\varepsilon\rho_I[(1+\beta_3)^2 + \tau\beta_3^2/\tau_\varepsilon^P] \\ & -(1+\beta_3)\tau_\varepsilon\rho_I[2(1+\beta_3)\beta_3' + \tau(2\beta_3\beta_3'\tau_\varepsilon^P - \beta_3^2)/\tau_\varepsilon^{P2}] \\ = -(\tau_\varepsilon\rho_I(1+\beta_3)^2\beta_3' + \tau_\varepsilon\rho_I\tau 2\beta_3\beta_3'/\tau_\varepsilon^P + \tau_\varepsilon\rho_I\tau\beta_3^2\beta_3'/\tau_\varepsilon^P) &+ \tau\rho_I\tau_\varepsilon(1+\beta_3)\beta_3^2/\tau_\varepsilon^{P2} > 0 \\ \iff -(1+\beta_3)^2\beta_3' - \tau 2\beta_3\beta_3'/\tau_\varepsilon^P - \tau\beta_3^2\beta_3'/\tau_\varepsilon^P &+ \tau(1+\beta_3)\beta_3^2/\tau_\varepsilon^{P2} > 0 \\ \iff \frac{r(-(1+r)(\tau_\theta+(\mu a)^2\tau_u)(\tau_\varepsilon^P+\tau_\theta+(\mu a)^2\tau_u)^2 + r\tau_\varepsilon(\tau_\varepsilon^{P2}(1+r) - (t+(\mu a)^2\tau_u)^2))}{(\tau_\varepsilon^P+\tau_\theta+(\mu a)^2\tau_u)^4} &> 0,\end{aligned}$$

which holds for $r \to \infty$, since the positive term grows at a cubic rate while the negative terms grow at most at a quadratic rate. Thus, for $r > 0$ large enough, $\text{RHS}_{\tau_\varepsilon^P} > 0$, hence $\partial a/\partial\tau_\varepsilon^P > 0$.

Price informativeness (adapting (1.4) for $\tau_\varepsilon^P \neq \tau_\varepsilon$)

$$\text{PI}_{\text{informed}} = \frac{1}{(\tau_\theta+\tau_\varepsilon^P)^2/((\mu a)^2\tau_u) + \tau_\theta + \tau_\varepsilon^P}$$

increases in $|a|, \mu, \tau_u$, decreases in $\tau_\theta, \tau_\varepsilon^P$, and does not change with $\psi_4/\psi_3, \tau_\varepsilon, \rho_I$. The total effect of parameter t on price informativeness is

$$\frac{d\text{PI}}{dt} = \frac{\partial \text{PI}}{\partial t} + \frac{\partial \text{PI}}{\partial |a|}\cdot\frac{\partial |a|}{\partial t}.$$

Thus, all of the above comparative statics also apply to price informativeness except for τ_θ and τ_ε^P, which decrease price informativeness directly ($\partial\text{PI}/\partial\tau_\theta < 0$) but increase $|a|$ ($\partial\text{PI}/\partial|a| \cdot \partial|a|/\partial\tau_\theta > 0$), making the effect potentially ambiguous. The effects of $\psi_4/\psi_3, \tau_\varepsilon, \rho_I$ are indirect via a ($\partial\text{PI}/\partial t = 0$).

The indirect effect is very small for $|r| \to \infty$, which leads to $\text{RHS}_{\tau_\theta} \to 0$ and $\text{RHS}_a \to 0$, hence $\partial|a|/\partial\tau_\theta = \text{RHS}_{\tau_\theta}/(\text{RHS}_a - 1) \to 0$. The same applies to τ_ε^P. Consequently, for

τ_θ and τ_ε^P, the direct effect on price informativeness dominates for large $|r|$, and price informativeness decreases with larger τ_θ and τ_ε^P.

Since the welfare gain in (1.6) is price informativeness times $\psi_4^2/(2\psi_3)$, all of the above comparative statics except for ψ_4/ψ_3 carry over to the welfare measure.

ii. $\text{RHS}_r < 0$ for $r < 0$ sufficiently negative. This follows from the fact that $\text{RHS}_r < 0$ (see (1.30)) if and only if (1.31) holds, which it does for $r << 0$ sufficiently negative. Thus, a decreases in response to a less extreme $\psi_4/\psi_3 < 0$.

$\text{RHS}_\mu < 0$, since the last term in (1.32) is positive for $r \to -\infty$. $\text{RHS}_{\rho_I} > 0$, since $\text{RHS}(a) < 0$ in equilibrium (Proposition 16). $\text{RHS}_{\tau_\theta} < 0$, since the last term in (1.33) is positive for $r \to -\infty$. $\text{RHS}_{\tau_\varepsilon} < 0$, since $1 + \beta_3 < 0$ for $r << 0$, which makes the slope in (1.35) negative. It can similarly be verified that $\text{RHS}_{\tau_u} < 0$ and $\text{RHS}_{\tau_\varepsilon^P} < 0$ if $r << 0$.

Since $r << 0$ implies $a < 0$ in equilibrium (Proposition 16), a decrease in a increases $|a|$. Consequently, the same arguments as in (*i.*) show that the same comparative statics results apply to price informativeness. As in (*i.*) all of these results carry over to the welfare measure except for ψ_4/ψ_3. $\square$

Chapter 2

Cutting Out the Middleman: Crowdinvesting, Efficiency, and Inequality[1]

2.1 Introduction

In most industrialized economies, financial wealth is distributed far more unequally than income. Late last century, 60 percent of American households possessed almost no financial wealth (one percent of total financial wealth), while the top five percent of households held more than two thirds of financial wealth (Wolff, 2002). Income inequality, in comparison, was much lower. The poorest 60 percent of US households received about 22 percent of total income. More recently, Saez and Zucman (2014) found that the bottom 90 percent of American households owned about 23 percent of wealth, but received 60 percent of income in 2012. A similar disparity of income and wealth distributions can be observed in many other countries (Davies et al., 2011). Moreover, according to Piketty (2014), the inequality of the wealth distribution increased over time in several industrialized nations.

In this paper, we show that decentralized investment processes which rely exclusively on the "wisdom of the crowd" can efficiently aggregate information about the potential success of new consumption goods, and channel funds to projects that need them the most. Efficient information aggregation requires that potential consumers of these new products have enough wealth to invest on the capital market. A major mismatch between the income and wealth distribution of consumers instead leads to inefficient investment choices that cannot be fully corrected by financial intermediaries.

Our results are derived from a Bayesian investment game with dispersed and correlated information about the future demand of new products. In our model, a firm invents a novel consumption good and tries to raise capital for production. The more money the firm

[1]This chapter is joint work with Hans Peter Grüner.

attracts, the more it can produce for later sale. Consumers look for investment opportunities on the capital market to increase their income for consumption, and may as one option invest in the new firm. We assume there is aggregate demand uncertainty, but tastes among consumers within the same class are correlated. Consequently, consumers can use their own preference for the new product as a signal about future aggregate demand and therefore profitability of the firm. We show that consumers who like the new product are more optimistic about its demand and invest, while consumers who dislike the new product invest elsewhere. Thus, consumption driven investment directs capital towards firms that are likely to find many customers.

We model the investment process as a form of crowdfunding, where shares of the firm are directly sold to many small consumers ("the crowd"), and the proceeds are used to increase production capacity.[2] Our main result is that crowdinvestments can efficiently allocate capital to firms if all consumers are wealthy enough to invest. If, however, some groups of consumers cannot invest but will receive income for consumption, then crowdinvestments reflect the preferences of the wealthy, but not necessarily those of all consumers. Thus, capital is misallocated and production does not always scale up with future demand, because consumers who are unable to invest still consume. In a way, the capital market works as an information aggregation device similar to a vote, except consumers without wealth to invest do not get a vote. The same efficiency results obtain if we consider debt instead of equity crowdfunding.

In our baseline model only consumers hold information about preferences for the innovative product in the population. In an extension to our baseline model, we introduce financial intermediaries such as investment funds, who can acquire information about consumer preferences and compete with crowdinvestors on the capital market. We show that these professional investors cannot completely rectify the capital misallocation that arises when some groups of consumers are unable to invest. Moreover, if all consumers can invest a sufficient amount, then financial intermediaries are driven out of the market. This is because the decentralized information about future aggregate demand among consumers is costlessly aggregated on the capital market. Financial intermediaries, on the other hand, have to acquire information at a cost. Thus, the crowd in the aggregate acts like an insider whose superior information drives intermediaries out. Professional investors may be active in the market only if the crowd is not perfectly informed, i.e., cannot aggregate all consumer preferences because not all consumers are able to invest.

Crowdfunding, i.e., financing forms that draw on the masses (e.g., consumers, the general

[2]This kind of crowdinvesting differs from traditional forms of financing such as IPOs, where typically a predetermined share of the firm is sold to larger institutional investors, with assistance of an underwriter. Unlike IPOs, our crowdinvestment process does not determine a share price; rather, crowdinvestors invest an amount of capital and are entitled to a share of firm earnings in proportion to their investment. Otherwise both forms are similar in that equity shares are sold, and indeed a specific form of IPOs, *direct public offerings*, are very close to the crowdinvestment process we consider.

public) rather than a few professional financial intermediaries, is becoming very popular. It encompasses different funding models with equity contracts (as we consider here), debt contracts, material rewards in case of success, or donations, and has grown considerably in recent years and may soon rival traditional funding forms. "In all, crowdfunding platforms have raised some \$2.7 billion and successfully funded more than a million campaigns in 2012, [...] with an 81% increase to \$5.1 billion expected for 2013. By 2025, the global crowdfunding market could reach between \$90 billion and \$96 billion—roughly 1.8 times the size of the global venture capital industry today" (Fortune, 2014). Unlike other sorts of crowdsourcing (e.g., decision making by crowds, software development by crowds), crowdfunding is not primarily designed to rely on the "wisdom of crowds" (Surowiecki, 2005). It was meant to replace conventional financial intermediation for cases when banks, venture capital firms and others were unwilling to provide funding. Yet our main result shows that crowdfunding can achieve both: raise capital, while at the same time aggregating information.

The structure of the paper is as follows. After reviewing the related literature, section 2.2 presents the basic model and the main result with crowdinvestors. Section 2.3 extends the model by adding a sector of what we call investment funds, who compete with crowdinvestors on the capital market. In section 2.4, we discuss further extensions of our basic analytical framework, including a model with sequential investments that permits investors to learn from other players' investment decisions during a crowdfunding campaign and a model with pre-order crowdfunding. The last section concludes. In the appendix, we analyze the extension to nonlinear production technologies.

2.1.1 Related literature

Our paper is related to several distinct strands of literature, which we shall discuss only briefly. First, our contribution is related to a large literature that studies the effects of wealth inequality on allocative efficiency and in particular on the functioning of capital markets. A non-exclusive list is Aghion and Bolton (1997), Banerjee and Newman (1993), de Mesa and Webb (1992), Galor and Zeira (1993), Grüner (2003), Grüner and Schils (2007), and Piketty (1997). All these papers emphasize the link between credit market imperfections and agents' investments into private production possibilities. Investors with little wealth either do not get credit for their individual investment projects, or they only get credit at a higher interest rate. This is why the distribution of wealth has macroeconomic implications. In the present paper, we instead consider the link between inequality of the wealth distribution and the investment in firms/technologies seeking funding. Moreover, a major difference of our model compared to existing incomplete markets models is that agents hold private information about consumption needs, which is also a signal about the realization of relevant aggregate uncertainty.

Second, our paper is part of an emerging literature on crowdfunding that already spans multiple disciplines including economics, business, and law. Several empirical studies investigate the determinants of fundraising success (e.g., Agrawal et al., 2011; Ahlers et al., 2012; Mollick, 2014; Li and Martin, 2014). Agrawal et al. (2013) and Belleflamme et al. (2015) review the first empirical findings and discuss economic concepts relating to crowdfunding. According to Agrawal et al. (2013), early results suggest that crowdinvesting can replace traditional sources of financing, just as we find in our model. They also remark that, when investments are linked to and motivated by an earlier access to the product, financing by the crowd may be able to provide information about demand to the entrepreneur that would not be available from venture capitalists. Our argument, instead, applies beyond campaigns with presale and is based on the information that the individual consumption preference reveals about others. In one of the first theoretical treatments, Belleflamme et al. (2014) investigate whether an entrepreneur should rather use pre-order or equity-based crowdfunding. In their model, crowdinvestors are motivated by "community benefits" (utility from contributing) rather than investment return considerations as in our case.

Third, our paper contributes to the financial intermediation literature (for reviews, see Bhattacharya and Thakor, 1993 or Gorton and Winton, 2003). Several reasons for the emergence of financial intermediation have been put forth. One of the most prominent reasons given is that financial intermediaries can reduce aggregate transaction costs (e.g., monitoring and screening costs) or benefit from economies of scale (e.g., Diamond, 1984; Williamson, 1986; Boot and Thakor, 1997). Relatedly, is has been argued that financial intermediaries can solve information asymmetry problems, like the problem of information credibility in markets for information (e.g., Leland and Pyle, 1977; Allen, 1990): A buyer of information (e.g., about asset values, entrepreneur project quality) typically cannot verify its quality at the time of purchase. An intermediary who uses the information to invest and sell shares of his portfolio can credibly signal that the information is good, unlike the original seller. Moreover, in the presence of heterogeneous priors about project success among agents, it has been argued that intermediaries can emerge to channel funds from pessimists, who are unwilling to invest without screening for good projects, to optimistic entrepreneurs, who want to go forward with projects without screening (Coval and Thakor, 2005). Intermediaries are neither optimistic nor pessimistic, and can therefore credibly commit to screening, unlike the optimistic entrepreneurs.

Our model shows that direct financing of firms can be superior in terms of costs and capital allocation, because it utilizes the decentralized information of the crowd. We show that in certain situations crowdinvestors have a cost advantage over financial intermediaries, because they can use information they already have and intermediaries first have to purchase. Suppose crowdinvestors *could not* use their own preference as free information on the profitability of the firm as they do in our model. Then financial intermediaries would emerge due to a standard cost advantage argument and would be able to drive individual traders out of the market, because they could purchase advantageous information which individual investors cannot afford. Therefore, existing decentralized information among consumers is crucial in our theory.

Our results contribute to the literature comparing market-based and bank-based financial systems (for a review, see Allen and Gale, 2001). Like us, Allen and Gale (1999) consider the problem of financing new technologies, but provide an alternative explanation why market finance (the analogue to our crowdinvestment) might emerge instead of financial intermediation. They show that sufficiently strong diversity of opinion (heterogeneous prior beliefs) among traders will favor market finance over intermediation, because an intermediary is more likely to make a suboptimal decision from the perspective of investors. In contrast to their model, our analysis focuses on the effect of the investor information and wealth distribution rather than diversity of opinion on the efficiency of capital allocation and the extent of financial intermediation, and investors in our model have a common prior.

It has also been noted that initial public offerings can aggregate useful information about the future success of projects (e.g., Benveniste and Spindt, 1989). The present paper studies the case where this information concerns the attractiveness of a firm's products for a population of consumers who also act as investors.

Our paper is closely related to Subrahmanyam and Titman (1999), who investigate the firm choice between public (market) financing and private financing (intermediation). In their model, two random variables influence the growth opportunities of an enterprise. Investors in the financial market can either decide to acquire costly information on one of these variables, or some subset of investors receives free ("serendipitous") information during their day-to-day activities on the other variable by chance. For example, a retail store employee might come across information about the future demand for a certain product while at work, which he can use to evaluate investment alternatives. The entrepreneur in their model decides whether to go public or use private finance, and he anticipates whether he receives more information about business growth opportunities from the market via stock prices or from the private financier. The entrepreneur uses the information he receives to decide how to invest the proceeds from selling shares of his company in growth opportunities.

While the analysis of Subrahmanyam and Titman (1999) explains firm choices between public and private finance based on informational benefits to the firm, we instead focus on the impact of wealth distribution and information among investors on financial market structure. To examine the effect of the wealth distribution, our analysis requires a more specific microfoundation of the "free information" which investors hold: Consumer preferences are correlated, hence by virtue of having preferences, consumers also possess some information about the aggregate demand for products. The consumer information in our model can therefore be viewed as systematic rather than serendipitous.

The second major difference is that the entrepreneur in Subrahmanyam and Titman (1999) chooses the financing form himself, whereas we assume that the entrepreneur offers equity in a crowdinvestment campaign, and competition between many small investors and professional investors determines who holds equity in this financial market. One of Subrahmanyam and Titman's main findings is that market financing is favored if costless information is more widespread, and we show in our context exactly when more of the costless information is incorporated in the capital market as a consequence of the wealth distribution of investors.

2.2 Crowdinvestment: The baseline model

2.2.1 Consumers and endowments

Consider an economy which is populated by a continuum of consumer-investors indexed by $i \in [0, 1]$, who we will also call 'crowdinvestors.' Each consumer has an initial endowment of wealth w_i in period 1 and receives an exogenous income y_i in period 2. Income and wealth are measured in monetary units. Individuals consume in period 2 and use the capital market to increase their income in period 2. They can invest any positive amount of money at the riskless rate R, i.e., one unit invested in period 1 turns into R units in period 2. The riskless rate R is exogenously given. In period 2, two consumption goods are available: consumption c (at a normalized price of 1) and the novel consumption good x. Consumers have private information about their preference for the novel good. Preferences are represented by the following utility function:

$$u(c_i, x_i, \theta_i) = c_i + \theta_i x_i^{\alpha}, \tag{2.1}$$

with $0 < \alpha < 1$. The parameter θ_i is private information of consumer i, with $\theta_i \in \{0, 1\}$, i.e., consumers either derive utility from consuming good x, or they do not.

There is a spot market for goods c and x in period 2. But there is neither a credit market on which consumers may borrow against future income y_i nor a forward market for the innovative good x. A credit market friction is key to our results because, on a perfect credit market, all consumers could borrow against their future income in order to finance the efficient investment in their preferred technology. Still, the assumption of no credit markets is stricter than necessary and only made for simplicity here.[3] Nonexistence of a forward market is an appropriate assumption if the innovative good x has important features that are not contractible at the funding stage, which is the case for many of the investment projects financed by crowdinvestors. Without a forward market, companies cannot finance their investments drawing on the current sales revenues and must rely on external funding.

[3]As will become clear later, it is sufficient to assume a wedge between borrowing and saving rates due to credit market frictions (e.g., Galor and Zeira, 1993), because borrowing requires an excess return from investing in equilibrium, which is incompatible with efficient capital allocation. Thus, allowing borrowing in imperfect credit markets does not change our efficiency results.

In section 2.4.2 we show that the pre-order crowdfunding and forward markets are similarly affected by wealth constraints as a equity crowdfunding market.

2.2.2 The Bayesian investment game

There is aggregate risk regarding the share of consumers who would like to consume good x in period 2. The share of consumers s who would like to consume this good is distributed according to

$$s := \int_0^1 \theta_i di = \begin{cases} \beta > 1/2 & \text{with probability } 1/2, \\ 1-\beta & \text{with probability } 1/2. \end{cases}$$

Observing his private signal $\theta_i = 1$, a consumer updates his beliefs that state $s = \beta$ has realized. The corresponding posterior probability is

$$\Pr(s = \beta | \theta_i = 1) = \frac{\frac{1}{2}\beta}{\frac{1}{2}\beta + \frac{1}{2}(1-\beta)} = \beta.$$

There are $m > 1$ firms which have access to a technology for the production of good x. Each firm produces according to the linear technology:

$$x_{sup}(X) = X,$$

where x_{sup} denotes the produced amount (supply) of the novel good and X the aggregate investment made in period 1.[4] Consumers may invest any amount $\hat{x}_i$ in these firms, and the total size of all firms is determined by the investments of all consumers

$$X = \int_0^1 \hat{x}_i di.$$

All firms act as price takers in period 2 and distribute profits to all shareholders according to their relative investment shares.

In period 2, consumers receive their exogenous income y_i and the return on their riskless or risky investments. Let $\tilde{y}_i$ be the total budget available to consumer i in period 2. An equilibrium is defined as follows.

Definition 6. *An equilibrium of the model consists of*

i. a consumption plan $x_i(p)$ for each consumer,

ii. an investment plan $\hat{x}_i(\theta_i)$ for each consumer, and

iii. a relative price function $p(X, s)$ for good x,

such that

[4] The robustness of the main result for nonlinear technologies is discussed in the appendix.

i. the consumption plan maximizes utility (2.1) *subject to the consumer's period 2 budget constraint,*

ii. the investment plans constitute a Bayesian Nash equilibrium of the investment game subject to the wealth constraints, taking into account the consumption plans and the relative price $p(X,s)$*, and*

iii. at price $p(X,s)$ *the aggregate demand for good* x *equals supply* x_{sup}.

Note that we save on notation by not including wealth w_i in the investment plan $\hat{x}_i(\theta_i)$, since consumer i's wealth is already associated with the index i.

2.2.3 Equilibrium on the goods market

In period 2, at a given price of the novel good p, a consumer maximizes (2.1) subject to the budget constraint

$$\tilde{y}_i \geq c_i + px_i.$$

Solving the maximization problem yields the individual demand for good x,

$$x_i(p) = \left(\frac{\alpha\theta_i}{p}\right)^{\frac{1}{1-\alpha}} = \theta_i \left(\frac{\alpha}{p}\right)^{\frac{1}{1-\alpha}}, \ \theta_i \in \{0,1\}. \tag{2.2}$$

Aggregate demand is therefore

$$x(p) = \int_0^1 \theta_i di \left(\frac{\alpha}{p}\right)^{\frac{1}{1-\alpha}} = s\left(\frac{\alpha}{p}\right)^{\frac{1}{1-\alpha}},$$

leading to inverted aggregate demand

$$p = \alpha\left(\frac{s}{x(p)}\right)^{1-\alpha}.$$

Since producing firms act as price takers on the product market in period 2, aggregate investment $X = x_{sup}$ determines the good's price in equilibrium according to

$$p = p\left(X,s\right) = \alpha\left(\frac{s}{X}\right)^{1-\alpha}.$$

The equilibrium return on investment in the production of good x simply equals the good's price,

$$r = \frac{p\left(X,s\right)X}{X} = p\left(X,s\right).$$

2.2.4 Equilibrium investment

Consider now a possible symmetric Bayesian Nash equilibrium of the investment game, where each consumer with preference/signal $\theta_i = 1$ invests the same amount $\hat{x}$, whereas $\theta_i = 0$ types do not invest. Consumers who invest a positive amount $\hat{x}$ less than w_i must be indifferent between an investment in the innovation and an investment at the risk-free rate R. This follows from condition *ii.* of definition 11. Therefore, the equilibrium investment $\hat{x}$ of consumers who care about the good can be determined as follows.

$$\begin{aligned} R &= \mathbb{E}_s[p\,(s \cdot \hat{x}, s)\,|\theta_i = 1] \\ &= \beta\alpha\left(\frac{\beta}{\beta\hat{x}}\right)^{1-\alpha} + (1-\beta)\,\alpha\left(\frac{1-\beta}{(1-\beta)\,\hat{x}}\right)^{1-\alpha} = \alpha\hat{x}^{\alpha-1} \\ \iff \hat{x} &= \left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}}. \end{aligned}$$

Hence, the equilibrium aggregate investment in state $s = \beta$ is

$$X = \beta\left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}}$$

and it is

$$X = (1-\beta)\left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}}$$

in state $s = 1-\beta$. The model leads to a result contradicting a standard economic intuition. The return on investment in the novel good is not higher in the "good" state of the world with high demand compared to the state with low demand for the novel good:

$$p = \alpha\left(\frac{s}{s\hat{x}}\right)^{1-\alpha} = R.$$

The reason is that equilibrium investment is proportional to the share of consumers demanding the good in the future, which makes the equilibrium good's price and therefore investment return state-independent.

We now compare this market equilibrium outcome to a planner's solution, assuming the planner knows the realization of s. An investment in x has an opportunity cost of R units of the consumption good c in period 2. Hence, social welfare is maximized when all individuals consume

$$x_i = \left(\frac{\alpha\theta_i}{R}\right)^{\frac{1}{1-\alpha}}.$$

This is the quantity demanded at a relative price of R. Any deviation of equilibrium prices from this level reduces social welfare. We state this result formally for later use.

Lemma 18. *With a linear production technology, the capital allocation is Pareto-efficient if and only if aggregate investment is $X = s\left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}}$. This outcome is realized in a market* *equilibrium if and only if the good's market clearing price is $p = R$ independent of the state s.*

Proof. See Appendix. □

Therefore, the symmetric equilibrium derived above maximizes social welfare.

Proposition 19. *When all consumers hold wealth $w_i \geq \left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}}$, there is a symmetric investment equilibrium in which all consumers invest an amount $\hat{x} = \left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}}$ in the capacity for the production of good x, if and only if they would like to consume good x themselves ($\theta_i = 1$). This symmetric equilibrium is Pareto-optimal and maximizes utilitarian welfare.*

According to Proposition 19, crowdinvestments efficiently replace a missing forward market for good x. The more consumers are interested in the good, the more are willing to invest and finance production. Thus, firms do not have to convince third parties that their business idea is worth investing in; instead, the source of funding is consumers who already find the product attractive.

Clearly, this efficiency result rests on the assumed linearity of the production function. In the presence of nonlinearities, the resulting equilibrium would generally not be efficient anymore. However, equilibrium investment still increases with the share of interested consumers (see the appendix for the model with nonlinear production technology).

2.2.5 Two wealth classes and equality

In the remainder of this paper, we consider an economy which is composed of two groups of consumers of equal size. The fraction of consumers who care about good x may differ across groups. The shares s_1 and s_2 of consumers who care about good x are independently distributed according to

$$s_1 = 2 \cdot \int_0^{0.5} \theta_i di = \begin{cases} \beta > 1/2 & \text{with probability } \frac{1}{2} \\ 1-\beta & \text{with probability } \frac{1}{2} \end{cases}$$

$$s_2 = 2 \cdot \int_{0.5}^{1} \theta_i di = \begin{cases} \beta > 1/2 & \text{with probability } \frac{1}{2} \\ 1-\beta & \text{with probability } \frac{1}{2}. \end{cases}$$

When a consumer from the first group ($g = 1$) observes signal $\theta_i = 1$, he receives information regarding the aggregate preference distribution in his own group, but still relies on his prior to estimate demand in the other group. Hence, he attaches the following posterior probabilities to the vector of states (s_1, s_2):

(s_1, s_2)	(β, β)	$(1-\beta, 1-\beta)$	$(1-\beta, \beta)$	$(\beta, 1-\beta)$
$\Pr((s_1, s_2) \vert \theta_i = 1, g = 1)$	$\frac{\beta}{2}$	$\frac{1-\beta}{2}$	$\frac{1-\beta}{2}$	$\frac{\beta}{2}$

In case a consumer from group 1 receives signal $\theta_i = 0$, the posterior probabilities are

(s_1, s_2)	(β,β)	$(1-\beta,1-\beta)$	$(1-\beta,\beta)$	$(\beta,1-\beta)$
$\Pr((s_1,s_2)\|\theta_i=0,g=1)$	$\frac{1-\beta}{2}$	$\frac{\beta}{2}$	$\frac{\beta}{2}$	$\frac{1-\beta}{2}$

Again, there is an equilibrium in which all consumers invest the same positive amount if they care about good x. The equilibrium investment $\hat{x}$ of these investors fulfills

$$\begin{aligned}
R &= \mathbb{E}_s[p(s\cdot\hat{x})\,|\theta_i=1] \\
&= \frac{\beta}{2}\alpha\left(\frac{\beta}{\beta\hat{x}}\right)^{1-\alpha} + \frac{1-\beta}{2}\alpha\left(\frac{1-\beta}{(1-\beta)\,\hat{x}}\right)^{1-\alpha} + \frac{1}{2}\alpha\left(\frac{\frac{1}{2}}{\frac{1}{2}\hat{x}}\right)^{1-\alpha} \\
&= \alpha\hat{x}^{\alpha-1} \\
\Longleftrightarrow\ \hat{x} &= \left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}}.
\end{aligned}$$

Assuming all consumers can afford this investment, equilibrium aggregate investment is

$$X = \begin{cases} \beta\left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}} & \text{if } s_1 = s_2 = \beta \\ (1-\beta)\left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}} & \text{if } s_1 = s_2 = 1-\beta \\ \frac{1}{2}\left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}} & \text{if } s_1 \neq s_2. \end{cases}$$

Again, aggregate investment scales up one-to-one with demand, the good's price is independent of the state, and the equilibrium maximizes social welfare.

A model with standard debt contracts promising a certain return of R instead of equity contracts would yield the same equilibrium outcome if all consumers hold enough wealth. Just as in the efficient equity equilibrium, every unit invested would yield a return of R to the investor and the firm would make zero profits.

2.2.6 The role of wealth inequality

Suppose now that consumers of group 2 (the poor) do not hold any wealth ($w_i = 0$) to invest, but they have some income $y_i > 0$ in period 2 for consumption. The poor cannot borrow against their own future income in order to finance an investment. Therefore, they do not invest on the capital market. Consider an equilibrium in which only wealthy consumers (group 1) with $\theta_i = 1$, invest an identical amount $\hat{x} > 0$. The individual symmetric investment $\hat{x}$ of wealthy consumers is determined by equating the expected investment return with its opportunity cost R:

$$
\begin{aligned}
R &= \mathbb{E}_s[p\,(s\cdot\hat{x})\,|\theta_i=1,g=1] = \sum_{s_1,s_2}\Pr((s_1,s_2)|\theta_i=1,g=1)\cdot\alpha\left(\frac{\frac{s_1+s_2}{2}}{\frac{s_1}{2}\cdot\hat{x}}\right)^{1-\alpha}\\
&= \frac{\beta}{2}\alpha\left(\frac{\beta}{\frac{\beta}{2}\hat{x}}\right)^{1-\alpha} + \frac{1-\beta}{2}\alpha\left(\frac{1-\beta}{\frac{1-\beta}{2}\hat{x}}\right)^{1-\alpha} + \frac{\beta}{2}\alpha\left(\frac{\frac{1}{2}}{\frac{\beta}{2}\hat{x}}\right)^{1-\alpha} + \frac{1-\beta}{2}\alpha\left(\frac{\frac{1}{2}}{\frac{1-\beta}{2}\hat{x}}\right)^{1-\alpha}\\
&= \frac{1}{2}\alpha\left(\frac{2}{\hat{x}}\right)^{1-\alpha} + \frac{\beta}{2}\alpha\left(\frac{1}{\beta\hat{x}}\right)^{1-\alpha} + \frac{1-\beta}{2}\alpha\left(\frac{1}{(1-\beta)\,\hat{x}}\right)^{1-\alpha}.
\end{aligned}
$$

Consequently, depending on the state of the world, there are only two aggregate investment levels:

$$
X = \begin{cases} \beta\hat{x} & \text{if } s_1 = \beta \\ (1-\beta)\,\hat{x} & \text{if } s_1 = 1-\beta. \end{cases}
$$

Equilibrium investment and thus supply only depends on the preferences of the wealthy, but aggregate demand depends on the preferences of all consumers. The new equilibrium does not maximize social welfare, because it does not take into account the marginal social benefit from an investment in production capacity to satisfy demand of the poor consumer group. We now show that efficient investment is possible if and only if all groups of consumers have sufficient aggregate wealth to invest, so that aggregate investment depends on the preferences of consumers from all groups who later consume the new product.

Proposition 20. *Consider the case of two distinct groups of consumers, and suppose that the preference and wealth distribution within each group is independent. Then there exists an efficient equilibrium if and only if aggregate wealth in each group is sufficient for interested consumers to finance production of the group's efficient consumption, i.e.,*

$$
2\cdot\int_0^{0.5} w_i di \geq (\alpha/R)^{\frac{1}{1-\alpha}} \text{ and } 2\cdot\int_{0.5}^{1} w_i di \geq (\alpha/R)^{\frac{1}{1-\alpha}}. \tag{2.3}
$$

Proof. Sufficiency: If all consumers have $w_i \geq (\alpha/R)^{1/(1-\alpha)}$, then section 2.2.5 demonstrates that an efficient equilibrium exists. The derivation is virtually identical with heterogeneous individual investment yielding the same aggregate investment.

Necessity: To be shown: If an efficient equilibrium exists, then aggregate wealth fulfills (2.3). In an efficient equilibrium, aggregate investment scales up linearly with the share of interested consumers in each group (Lemma 18):

$$
X = \frac{s_1+s_2}{2}(\alpha/R)^{1/(1-\alpha)}. \tag{2.4}
$$

Denote the investment amount of investor i if $\theta_i = 1$ by $\hat{x}_i(\theta_i = 1)$ and if $\theta_i = 0$ by $\hat{x}_i(\theta_i = 0)$. Recall that group 1 are all consumers $i \in [0, 0.5]$ and group 2 are all consumers $i \in (0.5, 1]$.

Now we can write aggregate investment X in terms of investment strategies of all consumers,

$$
\begin{aligned}
X &= \int_0^{0.5} [s_1 \hat{x}_i(\theta_i = 1) + (1 - s_1)\hat{x}_i(\theta_i = 0)]di + \int_{0.5}^{1} [s_2 \hat{x}_i(\theta_i = 1) + (1 - s_2)\hat{x}_i(\theta_i = 0)]di \\
&= \int_0^{0.5} [s_1(\hat{x}_i(\theta_i = 1) - \hat{x}_i(\theta_i = 0)) + \hat{x}_i(\theta_i = 0)]di \\
&\quad + \int_{0.5}^{1} [s_2(\hat{x}_i(\theta_i = 1) - \hat{x}_i(\theta_i = 0)) + \hat{x}_i(\theta_i = 0)]di.
\end{aligned} \tag{2.5}
$$

Since by assumption an efficient equilibrium exists, both (2.4) and (2.5) have to hold for all realizations of (s_1, s_2). This is only possible if $\int_0^1 \hat{x}_i(\theta_i = 0)di = 0$, i.e., consumers of type $\theta_i = 0$ do not invest. Hence, simplifying (2.5) and equating aggregate investment with efficient aggregate investment (2.4), the following conditions hold in any efficient equilibrium:

$$
\begin{aligned}
2 \cdot \int_0^{0.5} \hat{x}_i(\theta_i = 1)di &= (\alpha/R)^{1/(1-\alpha)} \text{ and} \\
2 \cdot \int_{0.5}^{1} \hat{x}_i(\theta_i = 1)di &= (\alpha/R)^{1/(1-\alpha)}.
\end{aligned} \tag{2.6}
$$

The investment budget constraint requires $\hat{x}_i(\theta_i = 1) \leq w_i$ for all i. Thus, (2.6) implies

$$
\begin{aligned}
2 \cdot \int_0^{0.5} w_i di &\geq (\alpha/R)^{1/(1-\alpha)} \text{ and} \\
2 \cdot \int_{0.5}^{1} w_i di &\geq (\alpha/R)^{1/(1-\alpha)}.
\end{aligned}
\qquad \square
$$

Independence of wealth and preference distribution within the group ensures that wealthy consumers can invest on behalf of their less wealthy fellow group members, because they have the same preference distribution. A consequence of Proposition 20 is that a mismatch of the income and wealth distribution on group level may lead to an inefficient allocation of financial capital. The reason is that the mismatch between wealth and income distribution leads to a mismatch between the ability to invest and the future propensity to consume. The result is related to the limits of arbitrage literature[5], because consumers from at least one group have information that would allow them to arbitrage away excess returns, but they cannot completely act on it due to wealth constraints.

In Proposition 20, consumer groups are characterized by their correlated preferences. The following corollary considers the special case in which groups are characterized by their wealth endowment and where preferences within the wealth classes are correlated.

Corollary 21. *Consider the case of two distinct groups of consumers, and suppose w_i is constant within each group. Then there exists an efficient equilibrium if and only if*

[5] See, for example, Shleifer and Vishny (1997), and Gromb and Vayanos (2010) for a review.

consumers in each group hold enough wealth to finance production of their own efficient consumption in case of $\theta_i = 1$, $w_i \geq (\alpha/R)^{1/(1-\alpha)}$.

In general, one could imagine additional equilibria to the efficient ones we illustrated if all consumers have sufficient wealth, but the next proposition shows that there are no equilibria where the capital allocation is inefficient.

Proposition 22. *If all consumers have wealth $w_i \geq (\alpha/R)^{1/(1-\alpha)}$, then an inefficient equilibrium does not exist.*

Proof. See Appendix. □

The intuition behind this result is as follows. If there were an inefficient equilibrium, then there would be at least one state where the return on investment exceeds R. Because the preferences of crowdinvestors contain information about the likelihood of such a state, they would best respond by increasing investment. For example, if state $(\beta, 1-\beta)$ has a return exceeding R, then crowdinvestors with $\theta_i = 1$ in group 1 assign a high probability to this state, and (collectively) increase investment until they expect a return of R. If, on the other hand, state $(1-\beta, \beta)$ has an excess return, then $\theta_i = 1$ crowdinvestors from group 2 assign a high probability to this state, and increase investment. Together, the two groups of interested consumers can remove any excess return, because they have enough wealth to arbitrage away mispricing.

2.3 Financial intermediaries and market research

2.3.1 The extended model

In this section, we add a financial sector consisting of $N \in \mathbb{N}$ investment funds[6], indexed by j, with exogenous large endowment $W_j > 0$, who may acquire information about consumer preferences and maximize expected investment returns. They can either make conventional investments with return R, or they can invest in the novel consumption good with variable return. These funds may be viewed as arbitrageurs, who arbitrage away excess returns in the investment of the firm producing the novel good.

We assume that investment funds have no information[7] on the realization of consumer preferences (unlike consumers, whose preference θ_i is informative). Funds may acquire information about the realization of preferences in the consumer population to identify

[6]We call the financial market intermediaries "investment funds," but these may be replaced by any other large investing institutional entity, such as banks, venture capital firms, hedge funds, pension funds, or investment banking divisions.

[7]This assumption is made to simplify the exposition, and any *imperfect* information about the realization of s for investment funds yields the same results concerning efficient investment for any $N \in \mathbb{N}$.

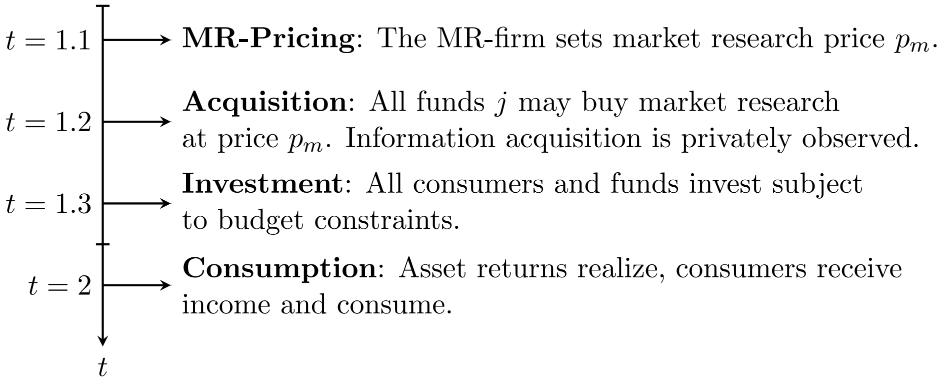

Figure 2.1: The timing of decisions.

worthwhile investment opportunitites. This can be thought of as buying market studies which evaluate the revenue potential of the new product or commissioning consumer surveys.

Formally, we represent the "market research" information by two binary and independent signals about the preference realization in the wealthy (1) and poor (2) consumer group, $m \in \{0,1\} \times \{0,1\}$. The signal quality is exogenously given by

$$\begin{aligned}\gamma := \Pr(m_1 = 1|s_1 = \beta) &= \Pr(m_1 = 0|s_1 = 1-\beta) \\ = \Pr(m_2 = 1|s_2 = \beta) &= \Pr(m_2 = 0|s_2 = 1-\beta) > 1/2.\end{aligned}$$

Market research is offered by a monopolist market research (MR) firm, which sells the same signal m to all interested buyers, i.e., signals are perfectly correlated. Neither the assumption that the MR sector is monopolistic nor that signals are perfectly correlated drives our results, as will become clear shortly. For non-triviality, we assume the MR firm can produce market research (i.e., conduct surveys, gather and analyze data) at sufficiently low cost $c > 0$, so that it can always offer market research at positive market research price p_m. If the MR firm sells market research to $0 \leq n \leq N$ funds, then its profit is given by $\pi_{MR} = np_m - \mathbf{1}\{n > 0\}c$.

In contrast to the model of section 2.2, aggregate investment is now the sum of the infinitesimally small crowdinvestments $\hat{x}_i$ and the investments of the "large" financial sector entities f_j. Thus aggregate investment in good x is

$$X = \int_0^1 \hat{x}_i di + \sum_{j=1}^{N} f_j.$$

The timing of decisions is displayed in Table 2.1. And now that we added new players to the game, we extend the equilibrium definition as follows.

Definition 7. *An equilibrium of the extended model consists of*

i. a market research price p_m set by the MR-firm at $t = 1.1$, *ii. an acquisition plan $a_j(p_m) \in \{0,1\}$ to purchase market research $m \in \{0,1\} \times \{0,1\}$ for each investment fund at $t = 1.2$,*

iii. an investment plan $\hat{x}_i(\theta_i)$ for each consumer at $t = 1.3$,

iv. an investment plan $f_j(p_m, m_j)$ for each investment fund at $t = 1.3$, where $m_j = m$ iff $a_j = 1$ and $m_j = \varnothing$ iff $a_j = 0$,

v. a consumption plan $x_i(p)$ for each consumer,

vi. a relative price function $p(X, s)$ for good x,

so that

i. the market price p_m maximizes expected profits of the market research firm at $t = 1.1$, taking into account $a_j(p_m)$ of all j,

ii. the information acquisition plans a_j and investment plans $\hat{x}_i$ and f_j constitute a Bayesian Nash equilibrium of the investment game subject to the wealth constraints, taking into account the consumption plans and the relative price $p(X, s)$,

iii. the consumption plan x_i maximizes utility subject to the consumer's future budget constraint, and

iv. at the price $p(X, s)$, the demand for good x equals production capacity X.

In our extended model, there are two possible sources of inefficiency, (i) that the creation of market research wastes cost $c > 0$ (new), and (ii) that state-contingent investment in the novel product is inefficient in the sense of Lemma 18 (as before). Since we assume that there is sufficient aggregate wealth in the economy to fund production of the efficient consumption in every state and also that utility is transferable, Pareto-efficiency from an ex-ante perspective requires that neither of the two kinds of inefficiencies occur, i.e., requires that no market research is carried out and that the capital allocation is efficient.

Definition 8. *Pareto-efficiency from an ex-ante perspective involves all agents in the economy (consumers, funds, market research firm), and requires that*

i. the market research cost $c > 0$ is not wasted, and

ii. the state-contingent capital allocation is efficient (Lemma 18).

The following analysis focuses on the possibility of efficient state-contingent investment, i.e., efficiency of the capital allocation (ii), which is necessary but not sufficient for Pareto-optimality. Our results show that Pareto-efficiency with an unequal wealth distribution fails not only because the market research cost is wasted, but because the capital allocation cannot be efficient even if market research is acquired in equilibrium.

2.3.2 Equilibrium existence

We first establish the existence of an equilibrium in the extended version of our model.

Proposition 23. *An equilibrium in which all crowdinvestors play pure strategies exists.*

Proof. See Appendix. □

The following sections analyze the equilibrium properties, especially with respect to the wealth distribution of crowdinvestors, in more detail.

2.3.3 The impossibility of efficient investment with active funds

To characterize the set of possible equilibria in more detail, we next show that efficient state dependent investment and active funds are inconsistent. The main obstacle to achieving efficient investment with active investment funds is an informational friction: Funds first have to buy the information that allows them to adjust their investment, but there are no excess returns in an efficient equilibrium that would incentivize them to buy market research. We discuss these obstacles in more detail in section 2.3.5.

Proposition 24. *There exists no equilibrium with an efficient state-dependent capital allocation in which investment funds invest.*

Proof. See Appendix. □

This result is independent of the wealth distribution of consumers. The proof proceeds in two main steps. First, suppose there is an efficient equilibrium where funds invest. Efficiency implies the investment return is R in every state (Lemma 18). But then it does not pay to buy market research for price $p_m > 0$, since return R can be realized elsewhere without this additional cost. Second, given that funds must be uninformed in an efficient equilibrium, their investment is constant over states s. Aggregate investment may still react to changes in s, since consumers may invest depending on their preferences. However, they do not invest as much as they would if investment funds were inactive, i.e., not as much as in the efficient equilibrium, since this would imply an expected return of less than R. But if consumers invest less, then the slope of aggregate investment $X(s)$ in s cannot be equal to $(\alpha/R)^{1/(1-\alpha)}$ as in the efficient equilibrium. That is, investment cannot scale up one-to-one with future aggregate demand. Consequently, there exists at least one state where aggregate investment is inefficient, which contradicts the earlier assumption that an efficient equilibrium in which funds invest exists.

2.3.4 Equilibrium if all consumers can invest

As benchmark, we again consider the case where all consumers have wealth $w_i \geq (\alpha/R)^{1/(1-\alpha)}$. In this case, the equilibrium of section 2.2 persists after adding investment funds: All consumers with type $\theta_i = 1$ invest, which is efficient and gives an investment return of R in each state (Proposition 19). Given this investment strategy by crowdinvestors, it does not pay for funds to participate; they do not buy market research and do not invest.

Proposition 25. *If all consumers have wealth $w_i \geq (\alpha/R)^{1/(1-\alpha)}$, then there exists an equilibrium where the consumer investment strategies are the same investment strategies as in Proposition 19 ($\hat{x}_i = \theta_i(\alpha/R)^{1/(1-\alpha)}$), and investment funds neither acquire information nor invest. This equilibrium is efficient.*

Proof. Suppose all consumers with $\theta_i = 1$ invest $\hat{x}_i = (\alpha/R)^{1/(1-\alpha)}$.

Investment stage: The profit of one of the N corporate investors when using investment strategy f_j with opportunity cost R and information set $\mathcal{I}_j$, given the investment strategies $\hat{x}_i$ of all consumers, is

$$\mathbb{E}_s[\pi_j(f_j, f_{-j}, \hat{x})|\mathcal{I}_j] = f_j(\mathbb{E}_s[p(f, \hat{x})|\mathcal{I}_j] - R).$$

The first order condition of Cournot competition with respect to f_j, taking investment strategies of all other players as given, is

$$0 = \mathbb{E}_s[p'(f, \hat{x})|\mathcal{I}_j]f_j + \mathbb{E}_s[p(f, \hat{x})|\mathcal{I}_j] - R \iff \mathbb{E}_s[p(f, \hat{x})|\mathcal{I}_j] = R - \mathbb{E}_s[p'(f, \hat{x})|\mathcal{I}_j]f_j, \quad (2.7)$$

hence funds aim to realize a price $p > R$, since $p' < 0$. However, the investments of the consumers are enough to realize a price $p = R$ in all states. Hence, first order condition (2.7) cannot be fulfilled with equality for any positive f_j, and the optimal choice is a corner solution $f_j = 0$ for all j.

Acquisition stage: Since investment funds do not invest, buying market research is strictly dominated for $p_m > 0$. □

2.3.5 Equilibrium if one group of consumers cannot invest

If a group of consumers is poor and cannot invest, then there may be investment opportunities for the financial sector. If the poor are interested in the novel good and the wealthy are not, then future demand for the novel good will be large but investment by the wealthy and consequently supply will be small. Hence, the price of the novel good p—which is also the per unit return of an investment in the novel good—is larger than R. In this state it would pay for the financial sector to swoop in and arbitrage away (part of) the excess return on investment, because wealthy investors underestimate future demand for the novel good.

However, as a consequence of Proposition 20 and Proposition 24, there will be some inefficiency in capital allocation whenever there is a group of consumers that does not have enough wealth to invest. Throughout this section, we assume that all consumers of group 2 (the poor) have no wealth, i.e., $\int_{0.5}^{1} w_i di = 0$.

Corollary 26. *There exists no equilibrium with an efficient state-dependent capital allocation if aggregate wealth in any group is less than* $1/2 \cdot (\alpha/R)^{1/(1-\alpha)}$.

Proof. *If investment funds do not invest, then the equilibrium cannot be efficient.* This follows from Proposition 20.

If investment funds invest, then the equilibrium cannot be efficient. This follows from Proposition 24. □

In order to see why an efficient outcome is impossible if some consumer groups cannot invest, we describe the frictions involved in more detail. One obstacle to efficiency is the market power of investment funds if $N < \infty$. Efficient investment implies that all investors make zero profits compared to the outside option at rate R, but if the fund sector is not perfectly competitive, then funds will withhold some investment to drive up prices (and therefore investment returns). This can be directly seen in the first order condition (2.7) of the fund investment problem. Thus, even if funds were perfectly informed about the state of consumer preferences s, they would not want to remove all inefficiency, as this would imply zero profits (or in fact a loss, since becoming informed is costly).

If the fund sector is competitive ($N \to \infty$), then an efficient equilibrium is still not possible. To understand why, consider the following proposition, which establishes that, if the investment fund sector is competitive, then aggregate investment will not be affected by market research in equilibrium.

Proposition 27. *Suppose the investment fund sector is competitive* ($N \to \infty$)*, so that* $X = \int \hat{x}_i di + \int f_j dj$. *Then there exists no equilibrium where a positive mass of funds buys market research for* $p_m > 0$.

Proof. Suppose there is an equilibrium with a positive mass of funds buying market research and investing in the novel good using the superior information. Because a single investment fund j is small and its investment does not influence p, j can deviate by not buying research, keep investing, and making the same investment return as before, yet saving cost $p_m > 0$. □

Proposition 27 shows that information acquisition is subject to a free-rider problem in a continuum of investment funds. As soon as aggregate investment reacts to market research information—which can only be the case if a positive probability mass of funds acquire it—then it pays to deviate for informed funds to not buying market research, and free-ride on the information incorporated in the aggregate investment by others.[8] Consequently, even

[8]The same argument would apply to crowdinvestors if they were allowed to buy market research. Hence, assuming that consumers may also buy market research would not change our results.

if there is a continuum of investment funds, no or only finitely many funds will become informed in equilibrium, but their impact on aggregate investment is negligible.[9]

Thus, with a competitive fund sector, the market for information breaks down. This result has a similar flavor as the one in Grossman and Stiglitz (1980) for financial markets, who show that there is no fully revealing equilibrium with costly information acquisition, because uninformed traders can free-ride on the information of informed traders.

Finally, even if a competitive fund sector somehow got hold of the market research signal for free, this would still not lead to efficient investment, unless market research was noiseless ($\gamma = 1$). That is, a noisy signal ($\gamma < 1$) prevents efficiency, because a wrong market research signal—which occurs with positive probability—leads to an inefficiently high or low investment.

Thus, an efficient equilibrium if not all groups of consumers can invest the efficient amount exists only if $\gamma = 1$, $N \to \infty$, and market research is costlessly available. But this is equivalent to a situation where the consumer preference realization is common knowledge, which is not realistic.

Our results show that financial intermediaries cannot fully correct the inefficiency that arises when wealth and income distribution do not match. However, they may still play a useful role in increasing social welfare in such situations. To see why this is so, consider as a simple example the case where no consumer holds any wealth. Then the addition of intermediaries is unambiguously welfare improving—even without the possibility to purchase market research.

2.4 Further extensions of the baseline model

2.4.1 Crowdinvestors and intermediaries

As we established in sections 2.2 and 2.3, efficient investment in our setting is possible if and only if aggregate wealth in both groups is large enough. This result is robust to a number of modifications and extensions of our model. It holds for any finite number of investment funds. When the number of funds approaches infinity, then there is a free-rider problem and the market for information breaks down (Proposition 27). Consequently, even a competitive fund sector is not sufficient for an efficient outcome.

The inefficiency results are also unaffected if several market research firms engage in Bertrand competition. In this case funds instead of MR firms extract profits, but the price of market research must stay positive in equilibrium, since it is costly to produce. As we

[9] Moreover, independent market research signals cannot yield efficient investment either. Although a law of large numbers guarantees that many independent market research draws m_j, $j = 1, \ldots, N$ reveal the state as $N \to \infty$ perfectly even for $\gamma < 1$, the market for information would break down, because it does not pay for funds to become informed (Proposition 27).

showed above, acquisition of market research at positive prices is incompatible with an efficient equilibrium.

Our results are robust to changes of our assumption concerning market research signal correlation. If market research signals are independent rather than perfectly correlated, then for $N < \infty$ the Cournot objective (2.7) of the funds still precludes efficient investment; for $N \to \infty$ the free-rider problem still prevents efficient aggregate investment. For the same reasons, our results hold for any market research signal quality $\gamma \in [0, 1]$.

It is also easy to see that extensions such as a larger number of states, a larger number of consumer groups, or the possibility for crowdinvestors to acquire market research do not change our conclusions. On the latter, as in the case of a competitive fund sector, an incentive to free-ride on the (costly) market research information of others prevents efficiency.

Consequently, the only situation where financial intermediaries can bring about efficient investment is if (i) perfect information about the preference realization in the population is (ii) costlessly available, and (iii) the investment fund sector is competitive.

2.4.2 Forward markets and pre-order crowdfunding

An important segment of the crowdfunding industry permits consumers to pre-order products. This segment essentially provides firms with an organized forward market that also serves as a financing device for innovations. It requires that the properties of the innovation are contractible at the date of pre-ordering. We now briefly discuss the robustness of the link between wealth/income distribution and efficiency when a pre-order crowdfunding instead of an equity crowdfunding market exists.

The analysis of a forward market requires a modified equilibrium concept. In addition to the objects listed in definition 11, an equilibrium at a given borrowing rate $B > R$ (due to credit market frictions) consists of a forward market price p_1, preference contingent forward market demand $x_i^F(\theta_i, p_1)$ (replacing all consumer investment plans) and a credit demand schedule $d_i(\theta_i, p_1)$ for each consumer. As before, we assume the innovative firm invests its entire revenue. Note that if $p_1 > 1$, then there is supply left after serving forward contract obligations in $t = 2$ and a market clearing spot market price p_2 realizes. Once more we consider the case with two wealth classes from Section 2, one wealthy and one poor. Since the forward market is the only way of raising capital for production, an unconstrained efficient outcome requires that the forward market revenue per interested consumer is $(\alpha\theta_i/R)^{1/(1-\alpha)}$. Obviously, this can only be achieved if interested consumers of *both* groups pay on average $(\alpha\theta_i/R)^{1/(1-\alpha)}$. Moreover, if there is no resale between consumers at $t = 2$, then efficiency requires that forward market payments by all interested consumers must equal $(\alpha\theta_i/R)^{1/(1-\alpha)}$ exactly. Due to $B > R$, wealthy and poor consumers face different opportunity costs in terms of $t = 2$ consumption when buying good x on the forward market. However, different opportunity costs of consumption and the fact that all interested agents consume the same amount of x is not compatible with utility maximization of all consumers. Therefore, pre-order crowdfunding can work efficiently only if the poor are sufficiently wealthy, just as we find in the case of equity crowdfunding.

2.4.3 Sequential investments

On most crowdfunding and crowdinvestment platforms, the current aggregate investment into a project is observable at any point in time for potential investors. An important question is whether our previous inefficiency results change if aggregate investment to date is observable. In that case wealthy crowdinvestors might learn something about the preferences of the poor, and consequently adjust their investment. In order to study this question, we extend the simultaneous investment game from section 2.2 to a simple sequential two stage investment game.

Consider the following modification of the baseline setup from section 2.2. In $t = 0$, all crowdinvestors may condition their investment plans $\hat{x}_i^{t=0}(\theta_i)$ only on their own private information, leading to aggregate investment $X_0 = \int_0^1 \hat{x}_i^0(\theta_i)di$. In $t = 1$, all crowdinvestors may condition their investment plans $\hat{x}_i^1(\theta_i, X_0)$ on their private information and aggregate investment from the previous investment stage. The equilibrium concept from definition 11 can be readily extended to the present setup by replacing the one stage by the two stage investment plans. In equilibrium, investors can adjust their investment to the realization of X_0, and use the information contained in X_0 about the distribution of θ_i when investing at $t = 1$. Overall investment by crowdinvestor i in the company is $\hat{x}_i^0(\theta_i) + \hat{x}_i^1(\theta_i, X_0)$, i.e., the sum of the investments in $t = 0$ and $t = 1$, with $\hat{x}_i^t \geq 0$ as before.

It is straightforward to show that all equilibria from the baseline model can be extended to equilibria in this dynamic model. Hence, the set of equilibria is weakly larger in the dynamic model.

Proposition 28. *Any equilibrium with investment strategy profile $\{\hat{x}_i(\theta_i)\}_i$ from the baseline model in section 2.2 can be extended to an outcome-identical equilibrium in the dynamic model.*

Proof. Take any equilibrium investment strategy profile $\{\hat{x}_i(\theta_i\}_i$ from the baseline model. Consider the following equilibrium candidate for the dynamic model:

$$\hat{x}_i^0(\theta_i) = 0 \ \forall i,$$
$$\hat{x}_i^1(\theta_i, X_0) = \hat{x}_i(\theta_i) \ \forall i.$$

Since nobody invests in $t = 0$, $X_0 = 0$ in all states, so aggregate investment is uninformative. Consequently, at $t = 1$, investors have the same information they have in the baseline model, so if $\hat{x}_i(\theta_i)$ is an equilibrium strategy in the baseline model, it also must be an equilibrium strategy in the last investment period of the dynamic model.

It remains to be shown that there is no profitable deviation at $t = 0$. A unilateral deviation of investing at $t = 0$ means investor i has the same information compared to the candidate strategy, and it does not change the investments by other investors, since i has no mass and does not affect X_0. Consequently, i is indifferent between investing earlier or investing according to the equilibrium candidate strategy. □

The question now is whether efficient equilibria exist in the dynamic model that do not exist in the baseline model (Proposition 20) due to the possibility of reacting to aggregate investment. First, note that an efficient equilibrium does not exist if consumers of one of the groups do not have any wealth. The intuition is quite simple: If investors of a group cannot invest at all, then nothing can be learned about their preferences from observing aggregate investment. However, efficient equilibria exist if the poor consumers cannot invest enough on their own, but enough so that aggregate investment becomes informative about their preferences.[10] More specifically, efficient equilibria exist as long as the poor consumers have some wealth and the wealthy have enough to cover the rest. The efficient equilibria are coordination equilibria in the sense that the poor consumers first invest and reveal their preference distribution to the wealthy consumers, who later invest on behalf of the poor.[11]

Proposition 28 implies that inefficient equilibria exist along with the efficient coordination equilibria. A reasonable equilibrium refinement is to require that no weakly dominated strategies are played in equilibrium. In our dynamic investment game, investing at $t = 0$ is a weakly dominated strategy: Clearly, any investment at $t = 0$ can be postponed to $t = 1$ without any drawbacks. However, if—off equilibrium—a large amount is invested at $t = 0$, so that the investment return will be below R in any state, then a player who postponed his own investment to $t = 1$ could still react by observing X_0 and not investing. A player who already committed at $t = 0$ has no such option.[12] Hence, if we restrict attention to equilibria without weakly dominated strategies, then efficient equilibria exist in the dynamic model only if they also exist in the static model of section 2.2, and our main results carry over to the dynamic model.

[10]See the online appendix for a detailed example at `http://gruener.vwl.uni-mannheim.de/fileadmin/user_upload/gruener/pdf/crowd-appendix.pdf`.

[11]There can be different efficient equilibria that overcome the wealth constraints of the poor, but all rely on the fact that the poor group reveals its preferences via investments at $t = 0$, so that others know how much more they have to invest in order to arbitrage away mispricing. This is why these efficient equilibria do not exist in the static model, or whenever aggregate investment is not observable. And clearly these efficient equilibria survive if we added more investment stages in the model or even set up a continuous time investment game.

[12]Indeed, in parimutuel betting—where as in our case the profits and losses are shared among all who invest—it is typically observed that bettors wait to place their bets until the very last moment in order to be able to react to new information (and not reveal their information to others), see, for example, (Ottaviani and Sørensen, 2009) and the references therein.

2.4.4 Nonlinear production technologies

We analyze nonlinear production technologies in detail in the appendix and only give a brief summary here. So far we only considered a linear production technology. A concave or convex technology implies that the efficient aggregate investment from a planner's perspective is nonlinearly increasing in the share of interested consumers. The aggregate investment made by crowdinvestors, on the other hand, is linearly increasing in the share of interested consumers if all consumers invest, and is not strictly increasing if some groups never invest. Thus, the market cannot achieve efficient investment as in the case of a linear production technology.

The main question is whether our previous results hold in terms of (ex ante) welfare, i.e, whether welfare is higher if all consumers can invest compared to the case where the poor consumer group cannot. In the appendix we compare two scenarios: In the equal wealth case, all consumers have sufficient wealth to make their investments. In the unequal wealth case, consumers of the poor group have no wealth whereas the wealthy have twice the wealth. The income distribution is the same in both scenarios.

We find that ex ante welfare is larger in the equal wealth scenario for concave, linear, and slightly convex production functions. Welfare is larger in the unequal wealth scenario only if there is a sufficiently large convexity in the production technology. Thus, our results generalize except for sufficiently convex production technologies in the sense that crowdinvesting yields higher welfare when all consumers have enough wealth to invest.

Given small nonlinearities, the reason why welfare is higher when wealth and income distributions match is the same as in the linear case: Since all consumers can invest, aggregate investment reacts to changes in the share of interested consumers, which is not the case with a wealth/income distribution mismatch. For a more detailed explanation of these findings we refer the reader to the appendix.

2.5 Conclusion

In most industrialized countries, wealth is far more concentrated in the population than income. We investigate the implications of this empirical fact when consumers invest in the capital market to increase their income for consumption, as can be observed in crowdinvestment campaigns. If tastes in the population are correlated, then consumers can use their own consumption preferences as signal for the profitability of firms. Consequently, consumers invest in companies whose products they like,[13] and firms can attract more capital for production if their product is well received among consumers.

[13] In the present model, the risk-neutral crowdinvestors make investment decisions based on the expected investment return. A different mechanism that may lead to similar investment patterns as we describe consists in consumers trying to hedge against price increases of products they like. As one option, consumers could hedge by investing in the company making the product, because price increases also lead to higher

We show that this pattern leads to an efficient capital allocation if all consumers who later consume also invest in the capital market. In this case, firms who need the most funding to build production capacity get the most funding. However, we also show that a wealth and income distribution mismatch may lead to an inefficient capital allocation. This is because firms with products favored by the wealthy will attract the most funding, but these are not necessarily the firms that meet the highest demand and therefore need the most funding.

We also show that financial intermediaries cannot completely fix the capital misallocation that arises with a mismatch of wealth and income distribution. The reason is that the acquisition of information about consumer preferences is costly. And even if the intermediaries were perfectly informed, they would not want to fully fund the new product, since the efficient investment implies zero profits. Unlike most of the financial intermediation literature, our setup is one where financial intermediaries are at an information cost disadvantage compared to consumers, who together hold enough information to perfectly predict future demand and therefore profitability of investments in firms.

By allowing intermediaries to compete with crowdinvestors, we endogenously determine the extent to which markets rely on intermediaries. If all consumers can invest in the capital market, then efficient investment is possible, which leaves no margin for profit and hence no room for intermediaries. The picture changes if some groups of consumers cannot invest, leading to over- or underinvestment in the new product, depending on the realization of consumer preferences. In this case it can pay for intermediaries to acquire information and reduce underinvestment.

Our analysis generates several testable predictions. First, if preferences are correlated among consumers, then consumers should tend to invest in firms whose products they like. This behavior should not (only) be driven by a sympathy for a brand name or the firm, but by the favorable information that the own preference for a product contains. Second, our model predicts that a very unequal wealth distribution relative to the income distribution limits the scope of direct financing mechanisms such as crowdinvesting (compared to intermediated finance). If wealth is concentrated only among few consumers, then the information aggregation function of crowdinvestment campaigns—a strong advantage compared to intermediated finance—is impaired. The (mis-)match of income and wealth distribution among the actual consumers of the product, and not the population, is crucial: Information aggregation of preferences for luxury products aimed at wealthy consumers may work even with very unequal wealth distributions in the population, because these consumers have sufficient wealth and income to invest and buy. But information aggregation may not work with products aimed at less fortunate consumers, who consume but cannot invest. Third, and relatedly, funding outcomes should on average be more efficient when the wealth dis-

returns on equity. An analysis investigating the impact of such a hedging motive would have to assume that investors are risk-averse.

tribution of consumers better matches the income distribution. This could either be tested across countries, or alternatively, within a country by comparing product success after different crowdinvestment campaigns that target consumers from different wealth and income groups.

Recent technological advances and the widespread use of the internet made it possible to match a large amount of investors with projects or firms seeking funding at substantially lower cost.[14] Thus, firms and projects that were previously too small to offer equity directly to the public, and therefore had to rely on financial intermediaries, now have access to the money and wisdom of crowds. Our results show that the improved access to financing from crowdinvestors increases the efficiency of capital allocation for those small firms, if the mismatch of wealth and income distribution of consumers is not too large. Hence, our paper shows that crowdinvesting may be a valuable financial innovation, which can improve social welfare.

[14]The UK crowdinvestment-platform crowdcube is one example. On this platform, 178 businesses collected on average about £281,000 from crowdinvestors, who on average invested about £415 (official statistics from 7th of January 2015, crowdcube.com).

Appendix: Proofs

Proof of Lemma 18. Concavity of the utility function (2.1) in consumption x_i implies that x_i must be equal for all $\theta_i = 1$ types in the social optimum, and equal zero for all $\theta_i = 0$ types. No waste and feasibility requires that per capita production equals per capita consumption for $\theta_i = 1$ types, i.e., $x_i = \hat{x}_i$. Thus, the social planner determines a constant per capita investment $\hat{x}_i = \hat{x}$ for all $\theta = 1$ types.

Because an investment $\hat{x}$ has opportunity cost of R units of c_i consumption, the budget constraint of the economy is

$$\int w_i - c_i - \theta_i R\hat{x}\ di = 0 \iff \int c_i\ di = \int w_i - \theta_i R\hat{x}\ di.$$

The planner's problem determines $\hat{x}$ to maximize total welfare,

$$\max_{\hat{x}} \int \theta_i x_i^\alpha + c_i\ di \text{ s.t. } \int c_i\ di = \int w_i - \theta_i R\hat{x}\ di.$$

Substituting from the budget constraint and using $x_i = \hat{x}$, this is equivalent to the unconstrained problem

$$\max_{\hat{x}} \int \theta_i \hat{x}^\alpha + w_i - \theta_i R\hat{x}\ di.$$

The first order necessary and sufficient condition of the concave objective is

$$0 = \int \alpha\theta_i\hat{x}^{\alpha-1} - \theta_i R\ di \iff \hat{x} = \left(\frac{\alpha}{R}\right)^{\frac{1}{1-\alpha}},$$

where $x_i = \hat{x} = (\alpha/R)^{\frac{1}{1-\alpha}}$ is also the socially optimal per capita consumption for $\theta_i = 1$ types. The corresponding efficient aggregate investment is $X = s\hat{x}$.

In a market equilibrium, consumption choices $x_i(p) = (\alpha/p)^{\frac{1}{1-\alpha}}$ for $\theta_i = 1$ types depend on market clearing price p, and are socially optimal if and only if aggregate investment is such that $p = R$ in every state. $\square$

Proof of Proposition 22. In a first step, we will show that consumers with type $\theta_i = 0$ do not invest in equilibrium if $w_i \geq (\alpha/R)^{1/(1-\alpha)}$. In a second step, we will show that if consumers from group 1 with $\theta_i = 1$ invest, then so do consumers from group 2 with $\theta_i = 1$, and vice versa. Finally, we show that equilibrium investment must be efficient.

First step: consumers with $\theta_i = 0$ do not invest in equilibrium. Denote the price in state $s = (\beta, \beta)$ by p_{11}, the price in state $s = (\beta, 1-\beta)$ by p_{10} and so on. Then we can write the expected returns of crowdinvestors of type $\theta_i = 1$ in group 1 and 2 as

$$\mathbb{E}_s[p|\theta_i = 1, g = 1] = \frac{\beta}{2}p_{11} + \frac{\beta}{2}p_{10} + \frac{1-\beta}{2}p_{00} + \frac{1-\beta}{2}p_{01},$$
$$\mathbb{E}_s[p|\theta_i = 1, g = 2] = \frac{\beta}{2}p_{11} + \frac{1-\beta}{2}p_{10} + \frac{1-\beta}{2}p_{00} + \frac{\beta}{2}p_{01}.$$

Similarly, the expected returns of consumers with $\theta_i = 0$ are

$$\begin{aligned}\mathbb{E}_s[p|\theta_i = 0, g = 1] &= \frac{1-\beta}{2}p_{11} + \frac{1-\beta}{2}p_{10} + \frac{\beta}{2}p_{00} + \frac{\beta}{2}p_{01},\\ \mathbb{E}_s[p|\theta_i = 0, g = 2] &= \frac{1-\beta}{2}p_{11} + \frac{\beta}{2}p_{10} + \frac{\beta}{2}p_{00} + \frac{1-\beta}{2}p_{01}.\end{aligned} \tag{2.8}$$

We want to show that consumers with $\theta_i = 0$ always expect a weakly lower investment return compared to consumers with $\theta_i = 1$. Thus, comparing $\mathbb{E}_s[p|\theta_i = 1, g = 2]$ with $\mathbb{E}_s[p|\theta_i = 1, g = 1]$,

$$\begin{aligned}&\mathbb{E}_s[p|\theta_i = 1, g = 1] \geq \mathbb{E}_s[p|\theta_i = 0, g = 2]\\ \iff &\frac{2\beta - 1}{2}p_{11} \geq \frac{2\beta - 1}{2}p_{00} \iff p_{11} \geq p_{00}.\end{aligned} \tag{2.9}$$

Denote the investment amount of investor i if $\theta_i = 1$ by $\hat{x}_i(\theta_i = 1)$ and if $\theta_i = 0$ by $\hat{x}_i(\theta_i = 0)$. Recall that group 1 are all consumers $i \in [0, 0.5]$ and group 2 are all consumers $i \in (0.5, 1]$. Now we can rewrite condition (2.9) in terms of investment strategies. After simplifying, (2.9) is equivalent to

$$\int_0^{0.5} \hat{x}_i(\theta_i = 0)di + \int_{0.5}^{1} \hat{x}_i(\theta_i = 0)di \geq 0, \tag{2.10}$$

which always holds true. Moreover, any positive aggregate investment by consumers with $\theta_i = 0$ from either group leads to $\mathbb{E}_s[p|\theta_i = 1, g = 1] > \mathbb{E}_s[p|\theta_i = 0, g = 2]$. Using the same reasoning, we also get $\mathbb{E}_s[p|\theta_i = 1, g = 2] \geq \mathbb{E}_s[p|\theta_i = 0, g = 1]$, and if (2.10) holds with strict inequality, then $\mathbb{E}_s[p|\theta_i = 1, g = 2] > \mathbb{E}_s[p|\theta_i = 0, g = 1]$.

Note that consumers only invest if their expected return is equal to or exceeds R, otherwise investing at the riskless rate R is a profitable deviation. Therefore, whenever consumers with $\theta_i = 0$ from either group invest, i.e., (2.10) holds with strict inequality, then consumers with $\theta_i = 1$ expect a return exceeding R. However, this cannot occur in equilibrium if $w_i \geq (\alpha/R)^{1/(1-\alpha)}$ for all i. Suppose (2.10) holds with strict inequality, then it is optimal for all consumers with $\theta_i = 1$ to increase their investment $\hat{x}_i(\theta_i = 1)$ until their expected return equals R. For at least one consumer this deviation must be feasible, since the wealth endowment w_i is sufficient for all consumers with $\theta_i = 1$ to invest $\hat{x}_i(\theta_i = 1) = (\alpha/R)^{1/(1-\alpha)}$, which guarantees a return of R or less. But if consumers with $\theta_i = 1$ expect a return of R, then by (2.9), consumers with $\theta_i = 0$ expect a return below R, which contradicts that (2.10) holds with strict inequality. Consequently, no consumer of type $\theta_i = 0$ invests in equilibrium.

Second step: If consumers with $\theta_i = 1$ from one group invest, then so do consumers $\theta_i = 1$ from the other group in equilibrium. Suppose that there is an equilibrium where some consumers with type $\theta_i = 1$ from group 1 (without loss of generality) invest in x, which implies $\mathbb{E}_s[p|\theta_i = 1, g = 1] \geq R$, otherwise not investing would be a profitable deviation. Now suppose to the contrary that investors from group 2 do not invest in this case, which implies they expect a weakly lower return from investing,

$$\begin{aligned} & \mathbb{E}_s[p|\theta_i = 1, g = 1] \geq \mathbb{E}_s[p|\theta_i = 1, g = 2] \\ \iff & \frac{2\beta - 1}{2} p_{10} \geq \frac{2\beta - 1}{2} p_{01} \iff p_{10} > p_{01}. \end{aligned} \tag{2.11}$$

Since the aggregate demand is the same in state $s = (\beta, 1-\beta)$ and $s = (1-\beta, \beta)$, price differences between these two states must be due to differences in aggregate investment. Rewriting condition (2.11) in terms of investment strategies gives

$$\begin{aligned} & \int_0^{0.5} [\beta \hat{x}_i(\theta_i = 1) + (1-\beta)\hat{x}_i(\theta_i = 0)] di + \int_{0.5}^1 [(1-\beta)\hat{x}_i(\theta_i = 1) + \beta \hat{x}_i(\theta_i = 0)] di \\ \leq & \int_0^{0.5} [(1-\beta)\hat{x}_i(\theta_i = 1) + \beta \hat{x}_i(\theta_i = 0)] di + \int_{0.5}^1 [\beta \hat{x}_i(\theta_i = 1) + (1-\beta)\hat{x}_i(\theta_i = 0)] di \\ \iff & \int_0^{0.5} [\hat{x}_i(\theta_i = 1) - \hat{x}_i(\theta_i = 0)] di \leq \int_{0.5}^1 [\hat{x}_i(\theta_i = 1) - \hat{x}_i(\theta_i = 0)] di \\ \iff & \int_0^{0.5} \hat{x}_i(\theta_i = 1) di \leq \int_{0.5}^1 \hat{x}_i(\theta_i = 1) di, \end{aligned} \tag{2.12}$$

where the last line follows from the fact that consumers with $\theta_i = 0$ do not invest in equilibrium (see first step). Thus, if some consumers from group 1 with $\theta_i = 1$ invest (i.e., $\int_0^{0.5} \hat{x}_i(\theta_i = 1) di > 0$), then consumers from group 2 with $\theta_i = 1$ must also invest ($\int_{0.5}^1 \hat{x}_i(\theta_i = 1) di > 0$). The same reasoning holds in the opposite direction as well: If consumers from group 2 with $\theta_i = 1$ invest, then so must those from group 1.

Condition (2.12) implies that whichever group of consumers with $\theta_i = 1$ invests less in the aggregate has a larger expected return. The argument in the next paragraph uses fact (2.12) to establish that, in equilibrium, we must have

$$\int_0^{0.5} \hat{x}_i(\theta_i = 1) di = s_1(\alpha/R)^{1/(1-\alpha)}, \int_{0.5}^1 \hat{x}_i(\theta_i = 1) di = s_2(\alpha/R)^{1/(1-\alpha)}, \tag{2.13}$$

which leads to a price of R in all states and is efficient (Lemma 18).

To show (2.13), suppose $\int_0^{0.5} \hat{x}_i(\theta_i = 1) di > s_1(\alpha/R)^{1/(1-\alpha)}$ and $\int_{0.5}^1 \hat{x}_i(\theta_i = 1) di > s_2(\alpha/R)^{1/(1-\alpha)}$, then $\mathbb{E}_s[p|\theta_i = 1, g = 1] < R$ and $\mathbb{E}_s[p|\theta_i = 1, g = 2] < R$, and not investing is a profitable deviation. Suppose $\int_0^{0.5} \hat{x}_i(\theta_i = 1) di < s_1(\alpha/R)^{1/(1-\alpha)}$ and $\int_{0.5}^1 \hat{x}_i(\theta_i = 1) di < s_2(\alpha/R)^{1/(1-\alpha)}$, then $\mathbb{E}_s[p|\theta_i = 1, g = 1] > R$ and $\mathbb{E}_s[p|\theta_i = 1, g = 2] > R$, and investing more is profitable and feasible. Suppose $\int_0^{0.5} \hat{x}_i(\theta_i = 1) di > s_1(\alpha/R)^{1/(1-\alpha)}$ and $\int_{0.5}^1 \hat{x}_i(\theta_i = 1) di < s_2(\alpha/R)^{1/(1-\alpha)}$, then either $\mathbb{E}_s[p|\theta_i = 1, g = 1] < R$ or $\mathbb{E}_s[p|\theta_i = 1, g = 2] > R$. Suppose $\int_0^{0.5} \hat{x}_i(\theta_i = 1) di < s_1(\alpha/R)^{1/(1-\alpha)}$ and $\int_{0.5}^1 \hat{x}_i(\theta_i = 1) di > s_2(\alpha/R)^{1/(1-\alpha)}$, then either $\mathbb{E}_s[p|\theta_i = 1, g = 1] > R$ or $\mathbb{E}_s[p|\theta_i = 1, g = 2] < R$. □

Proof of Proposition 23. We shall confirm that all equilibrium requirements of definition 7 can be fulfilled.

A unique market clearing price p exists for all aggregate investment levels X and all realizations of preferences (s_1, s_2). In the consumption stage, consumers use the demand function (2.2), which by construction maximizes utility.

Every price p_m set by the market research firm induces a Bayesian investment game at the acquisition and investment stage. In this investment game, all crowdinvestors i choose $\hat{x}_i \in [0, w_i]$ for each $\theta_i \in \{0, 1\}$ and all funds choose $(a_j, f_j) \in \{0, 1\} \times [0, W_j]$ for each $p_m \in \mathbb{R}_+$ and $m_j \in \{\{0, 1\}^2, \varnothing\}$, where $w_i \in [0, \infty)$ and $W_i \in [0, \infty)$.

Consider first a reduced game, where the strategy space for funds is $f_j \in [0, W_j]$ and information acquisition decisions $(a_1, a_2, \ldots, a_N)$ are exogenous. Then strategy spaces of all investors are compact and convex, and strategy f_j is concave and continuous in the expected payoff π_j for a given strategy profile $(f_{-j}, \hat{x})$ of all other investors, where

$$\mathbb{E}[\pi_j(f_j, f_{-j}, \hat{x})|\mathcal{I}_j(a_j)] = f_j(\mathbb{E}[p(f_j, f_{-j}, \hat{x})|\mathcal{I}_j(a_j)] - R),$$

and $\hat{x}_i$ is quasi-concave and continuous for crowdinvestors i. Thus, the Debreu-Glicksberg-Fan theorem (e.g., Theorem 1 in Reny, 2008) guarantees the existence of a pure strategy equilibrium for any exogenous profile $(a_1, a_2, \ldots, a_N)$.

Going back to the actual game with fund strategy space $(a_j, f_j) \in \{0, 1\} \times [0, W_j]$, which is not convex, every information acquisition profile $(a_1, a_2, \ldots, a_N)$ induces a reduced game for which we just showed a pure strategy equilibrium exists. By allowing mixed strategies in a_j, we can convexify the strategy space to $[0, 1] \times [0, W_j]$, and the expected payoffs from the mixed strategies are just linear combinations of the payoffs of the reduced game. Since a linear combination is quasi-concave and continuous, the Debreu-Glicksberg-Fan theorem guarantees existence of an equilibrium of the investment game, with possible mixing in a_j and corresponding f_j for all j and pure strategies for crowdinvestors $\hat{x}_i$.

We still have to show that, given the outcomes of the investment game for every p_m, there exists a profit maximizing price p_m for the MR firm. For a given p_m, all funds j determine the information acquisition decision by solving the problem

$$\max_{a_j \in [0,1]} \mathbb{E}[\pi_j | a_j, a_{-j}, f, \hat{x}] - a_j p_m,$$

where the set of mixed strategies $[0, 1]$ is compact, and $\mathbb{E}[\pi_j | a_j, a_{-j}, f, \hat{x}] - a_j p_m$ is continuous in p_m. Berge's maximum theorem implies that $a_j(p_m)$—the expected demand for market research by fund j—is upper hemi-continuous (uhc) in p_m. Aggregate expected demand for market research is $\sum_j a_j(p_m)$. The profit function for the market research firm is given by

$$\pi_{MR}(p_m) = p_m \sum_j a_j(p_m) - \mathbf{1}\left\{\sum_j a_j(p_m) > 0\right\} c.$$

Since summation and integration preserves upper hemi-continuity, $\sum_j a_j(p_m)$ is uhc. Moreover, the product of two non-negative uhc correspondences p_m and $\sum_j a_j(p_m)$ is uhc. The negative of the last term $\mathbf{1}\left\{\sum_j a_j(p_m) > 0\right\} c$ is lower hemi-continuous, since the indicator function $\mathbf{1}\{x \in X\}$ is lower hemi-continuous if and only if X is an open set. Consequently, $-\mathbf{1}\left\{\sum_j a_j(p_m) > 0\right\} c$ is uhc, and thus $\pi_{MR}(p_m)$ is uhc.

We can find an upper bound for a profit maximizing p_m, since no fund will buy market research if p_m is larger than the maximally possible earnings in the capital market, which are bounded. Denote such a bound by $0 < P < \infty$. Then, the market research firm chooses $p_m \in [0, P]$, which is a compact set, hence the Weierstrass extreme value theorem implies there exists a p_m which maximizes $\pi_{MR}(p_m)$. □

Proof of Proposition 24. Suppose an equilibrium with efficient investment exists in which some investment funds invest, which implies that the return on investment is R in every state (Lemma 18). In this case it does not pay for funds to buy market research at any price $p_m > 0$, as funds can by assumption obtain an investment return R by investing elsewhere without paying p_m. Consequently, investment funds must be uninformed in any efficient equilibrium, and invest a state independent amount $F := \sum_j f_j > 0$ in every state.

In any equilibrium, each consumer can condition his investment plan $\hat{x}_i$ on θ_i. Consequently, aggregate investment by consumers depending on the preference realization can be written as

$$\int_0^1 \hat{x}_i di = \int [s\hat{x}_i(\theta_i = 1) + (1-s)\hat{x}_i(\theta_i = 0)]di.$$

Efficiency requires that the price in each state equals R. In particular,

$$R = \alpha\left(\frac{\beta}{F + \int[\beta\hat{x}_i(\theta_i = 1) + (1-\beta)\hat{x}_i(\theta_i = 0)]di}\right)^{1-\alpha} \quad \text{if } s = \beta, \tag{2.14}$$

$$R = \alpha\left(\frac{1-\beta}{F + \int[(1-\beta)\hat{x}_i(\theta_i = 1) + \beta\hat{x}_i(\theta_i = 0)]di}\right)^{1-\alpha} \quad \text{if } s = 1-\beta, \tag{2.15}$$

and combining (2.14) and (2.15) implies

$$(2\beta - 1)F = (1 - 2\beta)\int \hat{x}_i(\theta_i = 0)di.$$

This condition is fulfilled with $F = \int \hat{x}_i(\theta_i = 0)di = 0$, which contradicts the assumption that investment funds invest. For $F > 0$ it implies $\int \hat{x}_i(\theta_i = 0)di < 0$, but this is impossible, thus contradicting efficiency. □

Chapter 3

An Experimental Analysis of Information Acquisition in Prediction Markets[1]

3.1 Introduction

Besides their role as allocation mechanisms, economists have long argued that markets can aggregate and convey information via prices. Increasingly, these markets are understood as tools to extract and aggregate private information in a way which can help organizations and policy makers to make decisions. Of particular interest are prediction markets (also called information markets), which are financial markets with assets whose values depend on the occurrence of events such as election outcomes, sports outcomes, or the default of a bank. For example, a typical election market asset pays \$1 if and only if the incumbent wins, and \$0 otherwise, while the complementary asset pays \$1 if and only if the challenger wins. Thus, the prices of these assets can be interpreted as a market probability forecast of the respective candidate winning the election.

In the present paper we build on this interest in the ability of information markets to convey information through prices. As pointed out by Grossman and Stiglitz (1980), there is a fundamental conflict between the efficiency of market prices and the incentives for traders to acquire information. Until now, most research on prediction markets has been about their ability to aggregate existing knowledge. We take a step back and ask how much information traders have in the first place. The main question we address in this paper is if and when prediction markets promote information acquisition.

Our design and findings are novel in the following ways. First, we vary the trading environment in terms of endowments and initial information to identify factors that induce more information acquisition and lead to more accurate price forecasts. Second, we investigate

[1] This chapter is joint work with Lionel Page.

trader characteristics that are associated with information acquisition and try to understand why these characteristics induce information acquisition. Third, beyond confirming the previous finding that acquiring information may lead to lower net profits, we also show that traders do not change the information acquisition behavior over many rounds of play. Fourth, we show that the overacquisition of information is a possible explanation why prediction markets provide accurate forecasts, because more information in the market typically leads to lower prediction errors of market prices. Finally, we suggest that the overacquisition phenomenon we observe is similar to rent overdissipation in the contest literature, and the contest nature of the market may explain the information acquisition behavior.

In our experiment, participants trade an asset whose value depends on the realization of the state of nature. Prior to trading, traders may acquire costly and imperfect signals about the state of nature. We use this experimental design to determine which factors—in terms of trading environment and individual trader characteristics—drive information acquisition on the individual level. Furthermore, we investigate whether these factors affect the market behavior, and in particular, whether increased information acquisition leads to lower prediction errors of market prices.

We find that traders invest a substantial amount of resources to collect information. In line with theoretical predictions, traders are more likely to acquire information if they have a larger endowment in cash and assets, if their existing information is inconclusive, and if they are more risk seeking. However, we find that most traders overinvest in information, so that traders obtain a negative average profit net of information costs. This finding persists even in the last rounds of the experiment, i.e., we find no evidence of reduced information acquisition over time to avoid these losses. This behavior is inconsistent with rational expectations theory and reminiscent of overdissipation in classical contest and tournament games (Anderson et al., 1998; Konrad, 2009; Sheremeta, 2010). Interestingly, while this overinvestment in information hurts net profits of traders, it improves the predictive accuracy of market prices. This phenomenon is a novel piece in the puzzle to explain why prediction markets work well in practice: Not only do they aggregate information, they also motivate a substantial amount of information acquisition. By varying the amount of initial information in the market, we also find that information acquisition is strongest if information in the market is weak. This finding helps to explain why prediction markets consistently produce good forecasts, even though initial information among traders is weaker in some markets than in others.

The remainder of the chapter is organized as follows: Section 1 reviews the relevant literature and related our findings. Section 2 briefly describes the theoretical background and a set of hypotheses regarding factors which drive information acquisition and their effect on market prices. Section 3 describes our experimental design, Section 4 presents our results and Section 5 concludes. The appendix presents additional analyses.

3.1.1 Related literature

Our study contributes to the literature investigating the effect of information acquisition in financial markets, which is a surprisingly small subset of the large experimental asset market literature.[2] Early experimental studies by Sunder (1992) and Copeland and Friedman (1992) test rational expectations equilibrium (REE) predictions about aggregate market statistics such as market prices and extent of information acquisition. Their findings are roughly consistent with predictions of REE or noisy REE. In particular, Sunder (1992) and Copeland and Friedman (1992) found that, on average, informed traders obtain the same net profits as uninformed traders. Consistent with the original Grossman and Stiglitz (1980) article, these early studies consider perfect information signals, while later studies consider more realistic information structures with imperfect information. Ackert et al. (1997) focus on pricing and information aggregation by comparing asset prices to REE predictions in markets with few imperfect signals. They find that asset prices are consistent with the rational expectations prediction in roughly half of the markets, and that auction prices of information are typically positive, but sometimes converge toward zero. Huber (2007) and Huber et al. (2011) extend the early studies to allow for different levels of information. Contrary to the early studies, they find that informed traders can receive lower net profits on average compared to uninformed traders. In particular, they find that intermediately informed traders perform the worst. We implement different degrees of information slightly differently, by allowing subjects to acquire several i.i.d. signals, but also find that informed traders on average perform worse than uninformed traders.

These previous studies focus on aggregate market statistics and the value of information. We take a more microfocused approach by investigating information acquisition on the individual level and how these individual decisions impact the market as a whole. For that purpose, our experimental design differs from earlier designs in that we vary factors suspected to influence the information acquisition decision, such as trader endowment and informativeness of initial information. We also administer a post-experimental questionnaire to gather individual trader characteristics, which allows us to investigate for instance the effect of risk aversion on information acquisition.

[2]For an earlier study of markets with a similar types of Arrow-Debreu securities, see Plott and Sunder (1988). For reviews of the experimental asset market literature, see Sunder (1995) and Noussair and Tucker (2013). Deck and Porter (2013) provide a survey on laboratory studies on prediction markets, and Palan (2013) gives a survey on experimental studies devoted to bubbles and crashes.

3.2 Theoretical background and hypotheses

3.2.1 Theoretical background

Markets play a key role in providing agents incentives to invest in the acquisition of new information. Better informed traders can make profits by taking advantage of situations where market prices deviate from asset values. Thus, informed traders act as arbitrageurs and help incorporate the public and private information into market prices. In the classical rational expectations equilibrium (REE) framework the price is fully revealing, it is a sufficient statistic for the realization of private information held by all traders. The REE framework can be seen as a formalization of the efficient market hypothesis, where markets aggregate information perfectly.

While this concept is very influential in financial economics, Grossman (1976) pointed out that in theory market prices cannot be fully efficient if information is costly, as fully revealing prices remove any incentive to acquire information. Grossman and Stiglitz (1980) attempted to solve this conflict in their noisy REE framework, which extends the standard REE with uncertainty, for example stochastic noise trader demand. Thus, market clearing prices in a noisy REE are typically not fully revealing, since a price movement may be due to favorable information on the asset value, or due to a different realization of the noise variable. Because the asset markets we consider have a zero-sum game character, any information revealed by a trader's orders in REE reduces the value of this information, and therefore reduces incentives to acquire costly information. In the extreme case without noise, the fully revealing REE predicts no information acquisition (Grossman and Stiglitz, 1980) and no trade (Milgrom and Stokey, 1982). In the noisy REE, insiders can gain an edge by acquiring superior information because uninformed traders can only partially infer their information from market prices. Thus, noisy REE with information acquisition exist (e.g., Grossman and Stiglitz, 1980, Peress, 2004).

Relative to the study of the ability of financial markets to efficiently aggregate information, there has been relatively little work exploring under which conditions these markets support information acquisition. Siemroth (2014) suggests that the information acquisition channel is one of the main reasons why prediction markets work in practice, as richer traders—who have the most weight in determining market prices—are able to acquire more information, making them better informed.

In the following subsection, we formulate several hypotheses regarding information acquisition—taken from both noisy REE and non-REE models—to be tested in the experiment.

3.2.2 Hypotheses

The experiment is designed to test the following hypotheses regarding factors which drive information acquisition (hypotheses 1-4) and how these factors affect the accuracy of market prices (hypotheses 5 and 6).

Starting with traders' characteristics, one could expect risk attitudes to play a role in information acquisition. Acquiring (imperfect) information can be seen as a risky strategy, because a potentially higher investment return comes at a certain cost. In a decision theoretic setting, Cabrales et al. (2014) demonstrated that the appeal of information acquisition decreases with risk aversion as measured by the Arrow-Pratt coefficient of absolute risk aversion. That is, all else equal, a less risk averse trader acquires information whenever a more a risk averse trader does. Similarly, in the noisy REE model of Peress (2004), his theorem 2 implies that less risk averse traders acquire more precise signals, i.e., acquire more information.

Hypothesis 1. *More risk averse traders are less likely to acquire information.*

We also hypothesize that the tendency to acquire information depends on the level of endowment of the trader. The following prediction is derived from the theoretical models of Peress (2004) and Siemroth (2014). The idea is that richer traders can trade more, and in a non-fully revealing market the value of information scales up with how many profitable trades can be made with it. Thus, richer traders should acquire more information.

Hypothesis 2. *Traders with larger endowment acquire more information.*

Furthermore, we hypothesize that the choice to acquire further information depends on the level of information traders already possess. Siemroth (2014) shows that traders with beliefs strongly favoring one investment option over another are less likely to acquire information, because it is unlikely to change the investment decision. Conversely, if a trader can barely distinguish which investment alternative is better, then he has strong incentives to acquire information.

Hypothesis 3. *Traders acquire more information if their initial information is inconclusive.*

A distinctive prediction of noisy REE or Bayesian Nash equilibrium with information acquisition in large markets is that informed traders earn higher gross returns to compensate them for the cost of information. On average, they receive the same net profits as uninformed traders with the same endowment. A rejection of this hypothesis implies either that traders may be rational but care about something other than expected net returns, or they care only about expected net returns but systematically pick actions that are suboptimal.

Hypothesis 4. *Informed traders achieve higher gross returns compared to uninformed traders, and the additional gross profit equals the information costs.*

The following two hypotheses describe how the previous individual level decisions impact the market as a whole. A higher accuracy of asset prices implies a smaller difference between the actual asset value and market prices measured by an appropriate distance metric (see the results section, where we explain our measure in detail). First, we predict that more information acquisition increases the accuracy of prices.

Hypothesis 5. *More information acquisition improves the accuracy of the market price.*

Moreover, in line with Hypothesis 2, we predict that larger endowment in the economy leads to more accurate prices, because it supports information acquisition.

Hypothesis 6. *Larger initial endowments in the economy improve the accuracy of the market price.*

Note that the first three hypotheses are inconsistent with a fully revealing rational expectations equilibrium, which predicts that traders do not acquire costly information (Grossman and Stiglitz, 1980). The noisy REE of Peress (2004) predicts all of these hypotheses except for hypothesis 3, because traders in his model have no heterogeneous beliefs at the time of information acquisition. Hence, it does not make any predictions about the effect of heterogeneous initial beliefs on information acquisition. Siemroth (2014) investigates the effect of traders' initial beliefs on their willingness to acquire new information, and the prediction from hypothesis 3 is a result from this model. We conjecture that a noisy REE with initial belief differences would yield the same prediction.

3.3 Experimental design

We recruited 9-12 participants for each of the 8 sessions of the experiment, with a total of 90 participants. The sessions were conducted at the Queensland University of Technology, and the subjects were local students predominantly in economics and business recruited via the ORSEE system (Greiner, 2004). Instructions were printed on paper for all participants to read. Participants were prohibited to talk during the experiment and only interacted via the experimental software, which was run in z-Tree (Fischbacher, 2007). Each session consisted of one practice round and 12 main rounds (markets). The practice round was neither used for payment nor used in our analysis.

States of nature and information. At the start of each market round, the computer randomly chooses a state of nature, A or B with 50-50 chances. The state of nature determines the content of a virtual urn which contains 10 balls. In state A, the urn contains 6 black balls and 4 white balls. In state B, the urn contains 4 black balls and 6 white balls. All of this is known to subjects. The realization of the urn is revealed perfectly to subjects only at the end of the round when we display the payoff information.

In all rounds, traders have the opportunity to trade several units of an asset which at the end of the round pays off either 10 ECU (experimental currency units) if the state of nature is A (the majority of balls in the urn is black) or 0 ECU if the state of nature is B (the majority of balls in the urn is white). Thus, the task of the traders is to figure out whether the majority of balls in the urn is black or white.

Before trading, participants receive or acquire information about the state of nature. First, every trader receives two independent signals in the form of two balls drawn from the virtual urn. These two draws may be observed privately or publicly, depending on the treatment (see below). The two free draws are what we call initial information of traders. Second, traders may use their initial endowment of cash to buy additional draws. They can acquire up to 10 additional draws at the cost of 3 ECU each. To do this, the participant enters the number of draws he/she wants to acquire, and these are then revealed together on the next screen. Thus, the information acquisition decisions are made prior to trading, and all traders make them simultaneously, mirroring static equilibrium models with information acquisition.

Formally, each draw is a signal $s_n \in \{A, B\}$, where $s_n = A$ represents a black draw, and all draws are made with replacement. In state A, 60% of the balls are black, hence the probability of drawing a black ball (a correct signal) is 60%, and the reverse is true for state B. In the instructions, we briefly explain to the subjects what a Bayesian posterior $\Pr(A|s_1, s_2, \ldots)$ given a profile of ball draws $(s_1, s_2, \ldots)$ is, and the program computes and displays the posterior probability for the profile of draws the subject observed.[3]

Trading. At the beginning of each round, traders see the realization of their initial information, which treatment is active in this round, and their endowment. Then they make the information acquisition decision, followed by the trading screen. Throughout the trading window, information such as their ball draw profile, the corresponding Bayesian posterior probability that the urn has a majority of black balls, and their available cash and number of assets are displayed in a frame (see screenshots in the appended instructions). Trading is organized as a continuous double auction. Traders can buy further assets with cash, either by accepting standing sell offers or by posting a buy offer and having it accepted by another trader. Or traders can sell assets for cash, by accepting buy offers or creating sell offers. We do not allow short-selling to keep the design as simple as possible, and traders cannot spend more cash than they have.

Besides standing buy and sell offers, traders also see the prices of all executed trades in a list. The trading window lasts 3 minutes. When trading closes, the state of nature is revealed, along with this round's profit and the average transaction price in this round. This summary screen ends the round. At the end of the session, every participant answers

[3]We show the posterior to help subjects make the most of the information they observe. Thus, the results we find later will not merely be due to subjects being unable to make sense of their information. For the same reason, Ackert et al. (1997) recruited only subjects who had taken a statistics class where Bayes' rule was covered.

a post-experimental questionnaire including a survey question on risk attitudes taken from Dohmen et al. (2011), which has been validated with incentivized choices in lottery games.

Treatments. We implemented a within-subject design, which allows us to exploit the within-subject variation in the trading environment to estimate the treatment effects using panel data regression models. Specifically, one of three treatments is randomly chosen at the beginning of each round. In the first (baseline) treatment, participants observe two ball draws privately as initial information. All draws are made independently, so different subjects may see different information. Traders start the round with a cash endowment of 40 ECU and 4 assets.

In addition to this baseline treatment, we implemented two different treatments to investigate how different market characteristics can influence information acquisition. First, we designed a treatment where the initial two draws are public. Some markets are characterized by more private and diverse knowledge (e.g., markets on sales targets within a firm) while some others are characterized by more homogeneous information publicly available (e.g., markets on political election outcomes). This treatment mimics markets where most initial information is publicly available. In this treatment, it is common knowledge among traders that others see the same two draws. As in the baseline treatment, they start with an endowment 40 ECU and 4 assets. Second, we designed a treatment where about half of the traders get an endowment twice as large (80 ECU and 8 assets).[4] This treatment allows us to study how information acquisition and price accuracy differ by endowment in the market and whether richer traders tend to acquire more information. The three treatments are therefore the following:

1. Initial information is privately observed, endowment of 40 ECU and 4 assets.

2. Initial information is publicly observed, endowment of 40 ECU and 4 assets.

3. Initial information is privately observed, endowment 40 ECU and 4 assets (with probability 1/2) or endowment 80 ECU and 8 assets (with probability 1/2).

Payment. At the end of the experiment, the program randomly selected one of the 12 main rounds as payment period, and converted the portfolio value in ECU (cash holdings, plus asset payouts if the state was A) to Australian dollars with 1ECU=A$0.29. This performance-dependent value plus a show up fee of $10 was the payment made to subjects. The average payment was $27.5. Sessions typically lasted between 80 and 90 minutes.

[4]We want to investigate how endowments facilitate incorporation of private information in prices and how endowments affect mispricing. Since overpricing requires sales of the asset, we have to increase not only the cash endowment but also the asset endowment to address this question. As a consequence, this treatment tests the joint effect of a higher endowment and additional assets.

3.4 Results

3.4.1 Summary statistics

Table C.1 in the appendix describes all variables used in the following analyses. Table 3.1 shows trader characteristics (elicited via the post-experimental questionnaire) and summary statistics of their choices in the experiment. On average, subjects executed between 6 and 7.2 trades per round, depending on the treatment, with most trades in the high endowment treatment.[5] Moreover, traders on average acquired between 1.7 and 2.2 draws, depending on the treatment, also with most acquisitions in the high endowment treatment. The relatively high standard deviation reflects a large heterogeneity in choices to acquire information. In a given round, on average 33% of subjects did not acquire information and 18% of subjects acquired 4 or more signals. The distribution of these choices is displayed in Figure 3.1. Interestingly, the frequency of these choices is not monotonically decreasing; acquiring two or three draws is more common than acquiring just one.

Table 3.1: Summary statistics on subject-round level by treatment

	Baseline		PublicDraw		HighEndowment	
	Mean	SD	Mean	SD	Mean	SD
Demographic variables						
Male	0.449	0.498	0.406	0.492	0.379	0.486
Age	24.94	8.203	25.38	8.123	24.56	7.630
Risk aversion	4.051	2.112	3.994	2.036	4.079	2.021
Experience	0.505	0.501	0.503	0.501	0.462	0.499
Economics/Business Major	0.632	0.483	0.637	0.481	0.600	0.491
Experimental outcomes						
Num acquired draws	1.689	1.898	1.891	1.846	2.221	2.111
Acquired information (dummy)	0.603	0.490	0.694	0.462	0.710	0.454
Num trades	6.519	6.096	6.031	5.608	7.200	5.118
Transaction LAPE	4.599	1.314	4.267	1.509	4.049	1.381
Bid-ask spread	1.298	1.538	1.146	1.515	1.150	1.484
Gross profit	0	17.89	0	15.97	0	23.65
Net profit	-5.068	18.04	-5.672	16.76	-6.662	22.82
Observations	370		320		390	

Transaction LAPE: linear prediction error from the transaction price, see (3.2).

[5]We delegated a more thorough analysis of the trading behavior to the appendix.

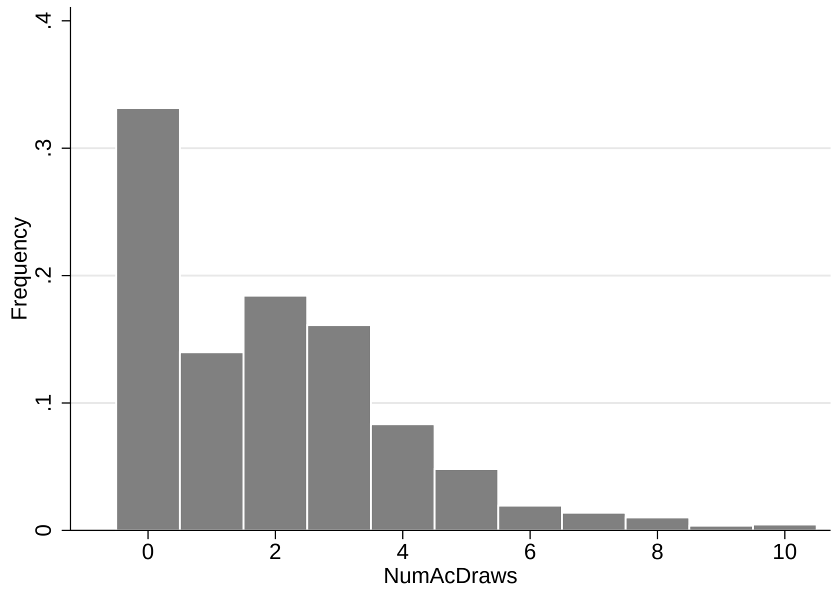

Figure 3.1: Histogram of the information acquisition choices (NumAcDraws), $N = 1080$

3.4.2 What drives information acquisition?

Table 3.2 displays the results of our analysis of the factors influencing the acquisition of information.[6]

Demographic variables. Looking at whether demographic characteristics predict information acquisition, we find that most of them, in particular gender, do not show a significant effect. We find, however, that previous experience in betting or investing in real markets is associated with a lower probability of information acquisition, as is higher risk aversion measured by the inverted Dohmen et al. (2011) 11 point risk attitude scale. A one point increase of risk aversion on the scale decreases the *probability* of acquiring at least one more draw by about 4 percentage points. Risk aversion is not significantly related to the *number* of acquired draws, however, but the sign is negative as in the probability of information acquisition regression.[7] While the theories cited above explain the risk aversion finding, the reason for the effect of experience is less obvious. Perhaps experienced traders have a better understanding of the costs and benefits of information when trading on a market, while inexperienced traders overestimate the value of information and acquire more. Age, experience, and a major in economics or business are associated with significantly lower

[6]Since NumAcDraws is a count data variable, we also ran the analyses in Table 3.2 using Poisson regression, negative binomial regression, conditional (fixed effects) Poisson regression, and hurdle models. We also ran logit regressions for the binary InfoAcquired, which is the probability of acquiring at least one more draw. In all cases, the results are very similar. Some of the nonlinear regressions are reported in Table C.2 in the appendix.

[7]All of these results are the same except for small changes in point estimates if we include age and risk aversion as third degree polynomials instead of linear trends. Moreover, since the Dohmen et al. (2011) scale is ordinal, we also ran the analysis with dummies for the upper 5 and lower 5 points of the 11 point scale, and the results are the same.

amounts of information acquisition. All of these have some dimension of experience, so these estimates may indicate that more experience in broad terms reduces the demand for information.

Result 1. *More risk averse and more experienced traders are less likely to acquire information.*

Effect of information content of initial draws. From Hypothesis 3, we expect traders to acquire more information when their initial information does not strongly favor one of the states of nature. In our experiment, traders can get three different profiles of initial information. First, they can draw one white and one black ball. In that case the Bayesian posterior probability of a black urn given this profile is 50%. Second, they can either draw two black balls (posterior probability of a black urn is 69.2%) or two white balls (posterior probability of a black urn is 30.8%). We created a dummy, ConclusiveInitDraws, which takes a value 1 when the first two draws have the same color, therefore favoring one of the state of the world ("conclusive evidence"). Including this dummy as control, we find that subjects acquire significantly less information if the initial information is conclusive both measured by the number of draws and the probability to acquire additional draws in all specifications (Table 3.2 columns 3-6). All else equal, conclusive initial draws reduce the probability to acquire additional draws by about 9 percentage points. Thus, if traders have evidence that points towards buying or selling, then they are significantly less likely to look for more information.

Result 2. *Traders acquire more information if their initial information is inconclusive.*

Effect of trader endowment. Our high endowment treatment allows us to investigate how the trader level of endowment affects information acquisition. Estimates in the first column of Table 3.2 show that more endowment in the economy (HighEndowmentTreat) induces significantly more information acquisition; about 0.53 more draws per trader on average. However, there may be two effects at play: (1) the knowledge that others might have high endowment and therefore can afford more information, and (2) the possibility that the trader himself has a high endowment, which is randomly assigned. Once we disentangle these two effects by introducing a variable indicating whether the subject started with a high endowment this period (HighEndowment), significance is retained only by the individual level variable (column 3). Consequently, it is not the knowledge that others might have more money that drives information acquisition, but the fact that the subject itself is endowed with more cash and assets to acquire information and trade. According to the point estimates in columns (3) and (4), doubling the initial endowment increases the amount of additional draws by 0.83 on average and the probability to acquire information by about 11 percentage points; both are highly significant effects. The increase in the amount of acquired draws corresponds to a wealth elasticity of information acquisition of about 0.48.

Table 3.2: What drives information acquisition?

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) OLS
Dependent variable	NumAcDraws	NumAcDraws	NumAcDraws	InfoAcquired	NumAcDraws	InfoAcquired
PubDrawsTreat	0.201	0.276***	0.234***	0.095***	0.167*	0.080***
	(0.150)	(0.081)	(0.083)	(0.028)	(0.096)	(0.030)
HighEndowmentTreat	0.531***	0.482***	0.082	0.027	0.137	0.068*
	(0.142)	(0.111)	(0.123)	(0.028)	(0.181)	(0.039)
HighEndowment			0.828***	0.111***	0.688***	0.066
			(0.161)	(0.035)	(0.214)	(0.042)
ConclusiveInitDraws			-0.387***	-0.087***	-0.477***	-0.130***
			(0.113)	(0.031)	(0.143)	(0.034)
RiskAversion					-0.118	-0.038**
					(0.103)	(0.017)
Experience					-0.732**	-0.181**
					(0.300)	(0.078)
Age					-0.034**	0.002
					(0.014)	(0.003)
Male					0.387	0.103
					(0.287)	(0.070)
EconBusMajor					-0.723**	-0.095
					(0.341)	(0.072)
Constant	1.689***	1.685***	1.881***	0.653***	3.930***	0.871***
	(0.102)	(0.056)	(0.083)	(0.022)	(0.896)	(0.142)
Subject Fixed Effects	no	yes	yes	yes	no	no
Cluster SE	no	yes	yes	yes	yes	yes
R^2	0.01	0.02	0.09	0.05	0.14	0.10
Observations	1080	1080	1080	1080	1068	1068
Clusters		90	90	90	89	89

Note: NumAcDraws is the number of acquired draws on top of the initial two free draws. InfoAcquired is a dummy variable, indicating that at least one additional draw was acquired. The baseline treatment is private draws and homogeneous endowment. If applicable, standard errors are clustered at the subject level. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Result 3. *Traders with larger endowment acquire more information, but knowledge that other traders have more endowment does not affect information acquisition.*

Effect of public draws. In our public treatment we replaced private individual draws by two publicly observable draws. We find that traders are more likely to acquire information in this setting. If the initial two draws are publicly rather than privately observed, then all fixed effects regressions indicate a positive effect on the propensity for traders to acquire information. According to columns (3) and (4), public draws significantly increase the number of additional draws acquired by about 0.23, and the probability to acquire at least one draw by about 9.5 percentage points. However, the reason behind this result is ambiguous. In addition to making initial draws publicly observable, the entire market has less information in the public draws treatment (2 draws), compared to the entire market in the baseline treatment (24 draws, 12 subjects with 2 draws each). Thus, our result suggest two possible interpretations. Either traders acquire more information when the initial information available in the market is weak, because the value of additional information is larger. Or it might indicate that traders acquire more information if most of the existing information is publicly available. Such a "level playing field" (symmetric initial information) may induce traders to acquire more information in order to gain an advantage.

In addition to aggregating private information, prediction markets provide incentives for traders to acquire additional information, and our results show that these are stronger when the available information is thin or most of it is publicly available. While the capacity of prediction markets to aggregate information is less important in these situations, the incentives they provide for traders to acquire additional information are a crucial feature that may help them outperform other forecasting institutions such as polls.

Result 4. *Traders acquire more information if the initial information is public and weak.*

3.4.3 How does information acquisition affect trader profits?

We now test whether Hypothesis 4 holds with more informed traders making higher gross profits and whether these gains equal the information costs.[8] To answer this question, we first define how trader profits are computed. Let the cash endowment at the beginning of a round (before any decision is made) be Cash^0 and the initial asset endowment be Assets^0. The endowments at the end of the round, i.e., after trading closes, are denoted by Cash^1 and Assets^1, respectively. Moreover, let the state of nature be represented by $\theta_{sr} \in \{0, 1\}$, with $\theta_{sr} = 1$ if and only if the state is A. We compute the net profit in ECU for each subject i in each round r of session s as

$$\text{NetProfit}_{sri} = (\text{Cash}^1_{sri} - \text{Cash}^0_{sri}) + 10 \cdot \theta_{sr}(\text{Assets}^1_{sri} - \text{Assets}^0_{sri}), \tag{3.1}$$

[8]Additional results on the links between information acquisition and trading activity are included in appendix.

Table 3.3: How does information acquisition affect profits?

	(1) OLS	(2) OLS	(3) OLS
Dependent variable	NetProfit	NetProfit	NetProfit
PubDrawsTreat	-0.604	-0.340	-0.241
	(1.494)	(1.483)	(1.389)
HighEndowmentTreat	-1.594	-0.921	-1.284
	(1.420)	(1.708)	(1.570)
HighEndowment		0.050	1.760
		(1.976)	(2.467)
NumAcDraws		-1.312***	-2.133***
		(0.303)	(0.552)
Constant	-5.068***	-2.851**	-1.458
	(1.017)	(1.132)	(1.266)
Subject Fixed Effects	no	no	yes
Cluster SE	no	no	yes
R^2	0.00	0.02	0.02
Observations	1080	1080	1080
Clusters			90

Note: NetProfit is the difference of portfolio values at the end and beginning of a round for each trader. The baseline treatment is private draws and homogeneous endowment. If applicable, standard errors are clustered at the subject level. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

where $\text{Cash}^1_{sri} - \text{Cash}^0_{sri}$ includes information acquisition costs, revenues from asset sales or spending for asset purchases. In words, profits are the difference of the values of the trader portfolios at the end and at the beginning of the round, where assets are evaluated at their fundamental value $10 \cdot \theta_{sr}$. Put differently, net profits are the initial portfolio value plus gross investment return less information acquisition costs. Consequently, if a subject neither acquires information nor trades, then he has a net profit of zero for certain.

Table 3.3 displays regressions with net profits as dependent variable.[9] We do not display regressions with demographic variables, as none of them showed a significant effect. Columns (2) and (3) of Table 3.3 include the number of acquired draws as an explanatory variable. In specifications with and without fixed effects, more information acquisition is associated with significantly lower net profits. According to the more conservative estimate, acquiring an additional draw reduces profits by about 1.3 ECU, which is almost half of the acquisition cost (3 ECU). Thus, although informed traders earn significantly higher gross profits by

[9]Results with the net rate of return rather than net profits in absolute terms are very similar. Moreover, including NumAcDraws as third degree polynomial rather than linearly yields virtually the same average effects (not displayed). Note that because the asset market absent information acquisition is a zero sum game, average net profits by treatment are affected only by more or less information acquisition.

trading,[10] on average these are not sufficient to recoup their information costs. This is strong evidence against the rational expectations equilibrium prediction that informed traders earn the same net profits as uninformed traders (Hypothesis 4). This result is not driven by self-selection of low ability traders (i.e., less able traders systematically being the ones that "overacquire" information) as it is significant and if anything stronger in the model with fixed effects (3).[11]

Result 5. *Better informed traders obtain higher gross profits, but they make net losses after subtracting information costs.*

Learning. We investigated whether this overacquisition of information persists over time or whether traders learn that their early information acquisition choices were excessive. We find no evidence of such learning. Comparing information acquisition in the first one to three rounds and comparing with information acquisition in the last one to three rounds yielded no significant differences (not displayed). Similarly, comparing the first few rounds to all remaining rounds or just including the round number as regressor yielded no significant difference. In fact, we did not find a single specification which would indicate that information acquisition decreases over time. Thus, even after more than 10 rounds of play, subjects did not learn that extensive information acquisition is associated with lower profits.[12] This finding is robust to including subject fixed effects, so it is not single traders that are driving this result.

Similarly, we investigated whether traders learn to make better use of the information, by allowing for a non linear learning pattern. We split the rounds in three sub-periods, a beginning, a middle and an end period. We then estimated whether the effect of information on traders' net profit differs across these periods. Specifically we ran regressions allowing the effect of information acquisition to differ for early, intermediate, and late rounds (recall that r indexes the round):

$$\begin{aligned}\text{NetProfit}_{sri} = \alpha + \beta\text{NumAcDraws}_{sri} + \gamma\text{NumAcDraws}_{sri} \cdot \mathbf{1}\{r \geq 10\}\\ +\delta\text{NumAcDraws}_{sri} \cdot \mathbf{1}\{r \leq 3\} + \eta'\mathbf{X}_{sri} + \varepsilon_{sri},\end{aligned}$$

where $\mathbf{1}\{x\}$ is the indicator function and $\mathbf{X}_{sri}$ includes for example the treatment dummies as before. We varied the upper threshold of 10, the lower one of 3, and whether fixed effects

[10] A regression with gross profits (excluding information costs) rather than net profits increases the point estimate by 3, all else equal, making it a significant positive effect.

[11] If the result from column (2) was due to selection, then the NumAcDraws coefficient in the fixed effects regression in column (3) should not be negative, but instead it is even more negative compared to (2).

[12] As pointed out by a reviewer, learning might require traders to experiment with different levels of information acquisition and observe the outcomes. In our sample, 40% of subjects had a standard deviation of 1 or larger in their amount of information acquisition, which could be interpreted as "experimentation." We tested whether experimentation decreased over time by looking at the variations from period to period in the number of signals acquired, using the absolute of the first difference of variable NumAcDraws, $|\text{Firstdif}_{it}| = |\text{NumAcDraws}_{it} - \text{NumAcDraws}_{it+1}|$. We do not find any time trend in this variable, which suggests that average experimentation does not decrease over time.

are included. In these regressions, the point estimate γ and the difference $\gamma - \delta$ are typically positive, which goes in the direction of a greater effect of information on net profit over time. However, there was no specification where these were significantly different from zero.

Result 6. *Traders do not learn to reduce their overacquisition of information over time.*

While excessive information acquisition is clearly hurting trader profits, it provides a novel explanation why prediction markets work well as forecasting tools. Traders invest a lot in acquiring information about the topics they trade on, so there is more information to be aggregated by prediction markets. In our experiment, information costs of 3 ECU might be viewed as relatively high, and results may be different with lower information costs. To check this possibility, we ran an additional session (12 traders with 12 rounds, $N = 144$) with the cost of information acquisition reduced to a third (1 ECU instead of 3 ECU). We found that participants reacted to the lower cost by acquiring more information overall and again lost money in doing so.[13]

The pattern we observe—substantial information acquisition leading to average net losses—is very similar to the phenomenon of overdissipation, which is consistently observed in tournament and contest games (e.g., Anderson et al., 1998; Konrad, 2009; Sheremeta, 2010; Sheremeta, 2013). In these games such as all-pay auctions, contest participants typically exert considerably more effort or incur higher costs than predicted by Nash equilibrium, which in the sum often exceeds the value of the prize that participants fight over. In many respects, a financial market is similar to a contest: There is a fixed endowment of value (cash and assets) each round, and this value can be redistributed via trading. According to rational expectations theory, participants should anticipate that any trade offered will be disadvantageous and therefore should be rejected (Milgrom and Stokey, 1982). Hence it is very easy not to lose in this 'contest' by not trading. However, as we observe, traders typically trade anyway (see also Angrisani et al., 2008; Carrillo and Palfrey, 2011). They acquire a lot of information, presumably because information in this financial market is useful in obtaining more of the prize (buy if asset is valuable, sell if worthless). Thus, information acquisition in our financial market looks very much like the analogue of investing resources in contest games.

Contests are often designed to induce effort among participants. Analogously prediction markets appear to be effective at inducing information acquisition effort in traders, even beyond the point where it is rational for them (assuming they only care about net profits). Thus, besides aggregating existing information, prediction markets appear to be effective in promoting desirable behavior just like contests. And it is not even necessary to explicitly provide a prize: The prize is the part of the endowment that traders wager by trading.

[13]Thus, our evidence suggests that trader overacquisition of information is not driven by high information costs. Looking at another type of experimental market, Huber et al. (2011) similarly found that averagely informed traders lose money compared to uninformed traders for different levels of information costs.

Traders provide the prize themselves and no extra subsidy by the market maker is required, unlike other mechanisms such as incentivized iterated polls (Healy et al., 2010).

Our finding that traders systematically over-invest in information compared to the REE prediction given rational agents with preferences depending only on net returns suggests that existing theories might have to be adapted. The contest literature already attempted to explain the observed divergence from Nash predictions and tested augmented theories (see the recent surveys Sheremeta, 2013 or Sheremeta, 2015 and the references therein). One avenue is to extend the standard preferences defined over wealth or net returns to include utility from winning, relative payoffs, or a preference for figuring out the unknown state. However, a convincing model has to go beyond those contest theories, since most contests are not games of asymmetric information, and the uncertainty about asset values and differential information among traders is crucial in financial markets. Thus, the contest theories should be seen as starting points for augmented models rather than substitutes for the existing ones.

3.4.4 Calibration and predictive power of market prices

To get a feeling for the trading process, Figure 3.2 displays the price accuracy for each transaction in all of the 12 rounds for the first session over time. To make the graphs for state A and B comparable, we plotted $(10 - \text{price})$ instead of the price whenever the state of nature was B. The distance between the line and 10 therefore represents the absolute difference between transaction price and value of the asset. For comparison, we also plotted the full information Bayesian price prediction, i.e., the asset value a risk neutral Bayesian would assign given knowledge of all observed draws of all traders in the market (dashed line, see below for a more detailed explanation of the Bayesian full information posterior).

In most of the rounds, the market price implied a forecast of 60% to 80% that the underlying event occurs. Occasionally, as in round 10 of this specific session, the market predicted the wrong outcome (implied forecast less than 50%). The graphs suggest that transaction prices tend to increase in accuracy over time, especially in round 10 or 11. This observation is confirmed by a Wilcoxon signed-rank test comparing prediction accuracy of first and last transaction prices in the entire sample ($z = 3.427$, $p = 0.0006$). Figure C.1 in the appendix plots the bids and asks for the first session over time.

Calibration. There are several ways to assess the quality of prediction market prices as forecasting tools. First, one can assess their calibration, i.e. whether they are good estimates of the probability of the forecasted event. Page and Clemen (2013) investigated price calibration by testing whether prices are an unbiased estimate of the frequency of the event, given the observed price. For example, if a price of 2 is a well calibrated forecast, the outcome should be A in 20% of the cases where we observe such a price. Formally, well

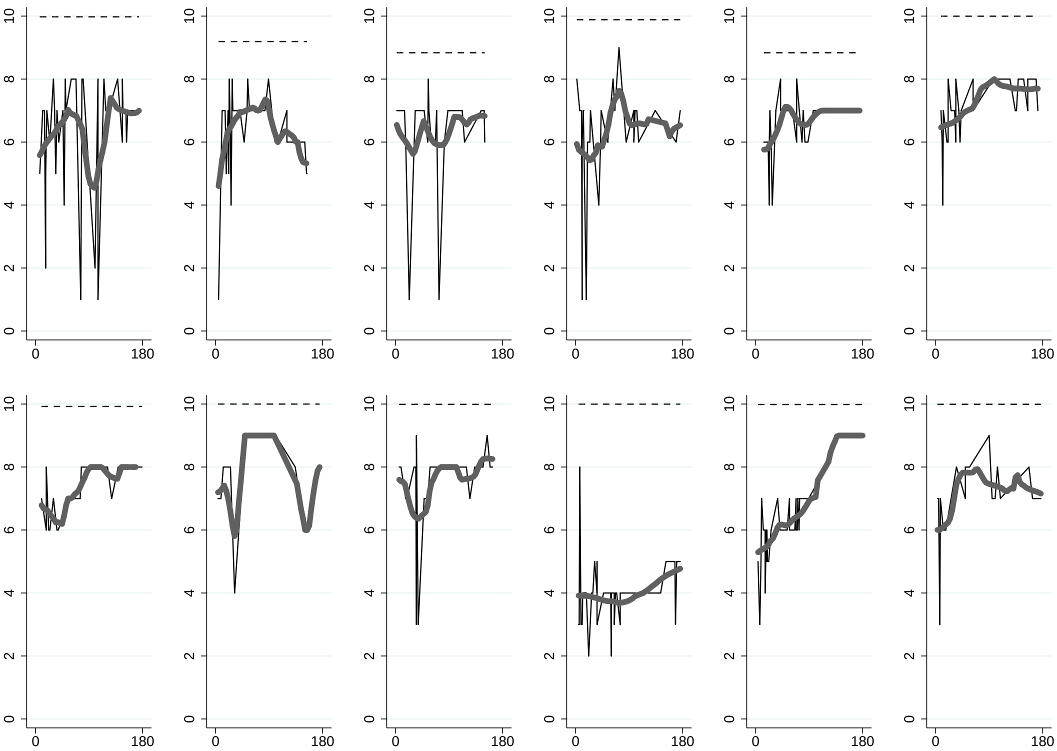

Figure 3.2: Transaction price accuracy (y-axis) over time (x-axis) in each of the 12 rounds of the first session in black, smoothed price accuracy graph in grey, and the Bayesian full information price prediction as dashed line.

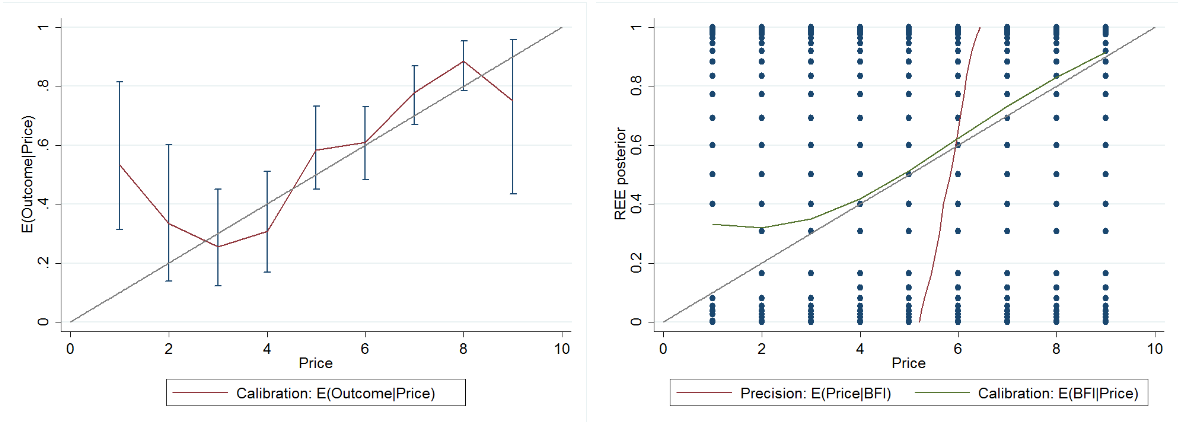

Figure 3.3: Calibration and precision of market prices. In the left panel, confidence intervals are given by percentile clustered bootstrap using markets as clusters. In the right panels, estimations are computed by local linear regression, using Epanechnikov kernel bandwiths of 2 (Price) and 0.2 (BFI).

calibrated market prices fulfill the following condition:

$$\text{Price}_{srt} = 10 \cdot \mathbb{E}[\theta_{sr}|\text{Price}_{srt}],\ \theta_{sr} \in \{0, 1\}.$$

We display evidence on calibration in Figure 3.3, left panel. It shows the frequency of outcome A for each of the transaction prices, $\text{Price}_{srt} \in \{1, 2, \ldots, 9\}$, where we cluster bootstrap standard errors within rounds.[14] The graph shows that the transaction prices are well calibrated for interior values ($\text{Price}_{srt} \in \{2, \ldots, 9\}$), where we cannot reject the hypotheses that frequencies (times 10) are equal to prices. Thus, interior prices can be considered unbiased forecasts of the outcome frequencies. However, for a price of 1, the calibration is off quite a bit, and the outcome frequency is significantly above the price.

Result 7. *Market prices are good estimates of the probability of the forecasted event for interior values i.e.,* $\text{Price}_{srt} \in \{2, \ldots, 9\}$.

An advantage of our experiment relative to field data is that we can measure the "objective probability" that A occurs given information in the form of the Bayesian posterior probability of A given all draws in the market, which we will abbreviate as BFI (Bayesian full information). This posterior is identical to the fully revealing REE price divided by 10.[15] We can now estimate the *precision* of the prices, which reflects how well market prices tend to follow the full information Bayesian posterior. To do so, we estimate $\mathbb{E}[\text{Price}_{srt}|\text{BFI}_{sr}]$. We can also measure the calibration of market prices relative to the BFI, $\mathbb{E}[\text{BFI}_{sr}|\text{Price}_{srt}]$, similarly to what we do with outcomes.

[14] We observed a very small number of transactions with prices 0 or 10 (44 out of 3952 transactions). Such prices are weakly dominated offer strategies, we therefore consider them as mistakes and did not include them in the calibration analysis and graph.

[15] More precisely, the posterior is identical to the fully revealing REE price for risk neutral traders, taking the information acquisition decisions as given. As we explain in the predictions section, the fully revealing REE predicts no information acquisition.

Figure 3.3, right panel, superimposes the calibration and precision curves of the market prices relative to the BFI. The calibration curve is very similar to the one observed for outcomes in the left panel. This is not surprising, because the full information posteriors are typically very close to 0 or 1 given the large amount of information in the market (see, for example, Figure 3.4 below). However we observe that the precision curve is very different from the calibration curve. Prices do not reflect all the information available in the market, contradicting the fully revealing REE price prediction. The precision graph illustrates that market prices underreact to information as represented by the BFI. For example, consider a Bayesian posterior given all information of 0.8: The average transaction price given this full information posterior is only around 6 while it should be 8. On the other hand, average transaction prices are slightly above 5 if the Bayesian posterior is 0.2. Consequently, market prices tend to lean towards 5 or 6, whereas the Bayesian posteriors are more extreme.

Underreaction to information. We can measure this underreaction formally. Denote the mean transaction price in a market by $\overline{\text{Price}}_{sr} = T_{sr}^{-1} \sum_{t=1}^{T_{sr}} \text{Price}_{srt}$, where T_{sr} is the number of transactions which took place on a given market. Then we can define an underreaction measure as follows:

$$\text{Underreaction}_{sr} = \begin{cases} (10 \cdot \text{BFI}_{sr} - \overline{\text{Price}}_{sr}) & \text{if } \text{BFI}_{sr} > 1/2, \\ 0 & \text{if } \text{BFI}_{sr} = 1/2, \\ (\overline{\text{Price}}_{sr} - 10 \cdot \text{BFI}_{sr}) & \text{if } \text{BFI}_{sr} < 1/2. \end{cases}$$

This underreaction measure is larger (positive) if mean transaction prices tend to forecast probabilities closer to 50% than the BFI. For example, if the BFI is 0.9, and the mean transaction price is 7, then underreaction is $10 \cdot 0.9 - 7 = 2$. On the other hand, if the BFI is 0.7 and the mean transaction price is 9, then underreaction is $10 \cdot 0.7 - 9 = -2$, i.e., prices overreact (are more extreme than the Bayesian posterior).

Using a t-test, we test whether underreaction equals zero (each observation is one market). This hypothesis is rejected at any significance level ($t = 18.5$, $df = 95$, $p < 0.001$). Mean underreaction is substantial at 3.2, thus information aggregation is not perfect.[16] The underreaction is also clearly visible in Figure 3.2.

Result 8. *Information aggregation is imperfect, and market prices substantially underreact to information compared to fully revealing REE prices.*

[16]We also tested whether the degree of underreaction differs for favorable versus unfavourable information in the market (we thank a reviewer for this suggestion). Specifically, we tested for underreaction for the cases $\text{BFI}_{sr} > 1/2$ and $\text{BFI}_{sr} < 1/2$ separately, with underreaction means 2.90 and 4.19, respectively (both significantly different from zero at all levels). The hypothesis of no underreaction difference between the two is rejected with a t-test at all significance levels ($t = 4.4$, $df = 91$, $p < 0.001$). Thus, the results indicate that assets are more overpriced given unfavorable information than underpriced given favorable information. Such difference in underreaction could, for example, be explained if the observed underreaction is driven by the differences in traders risk attitudes when facing gains (favorable information) and losses (unfavorable information), as suggested in an explanation for the disposition effect by Frazzini (2006).

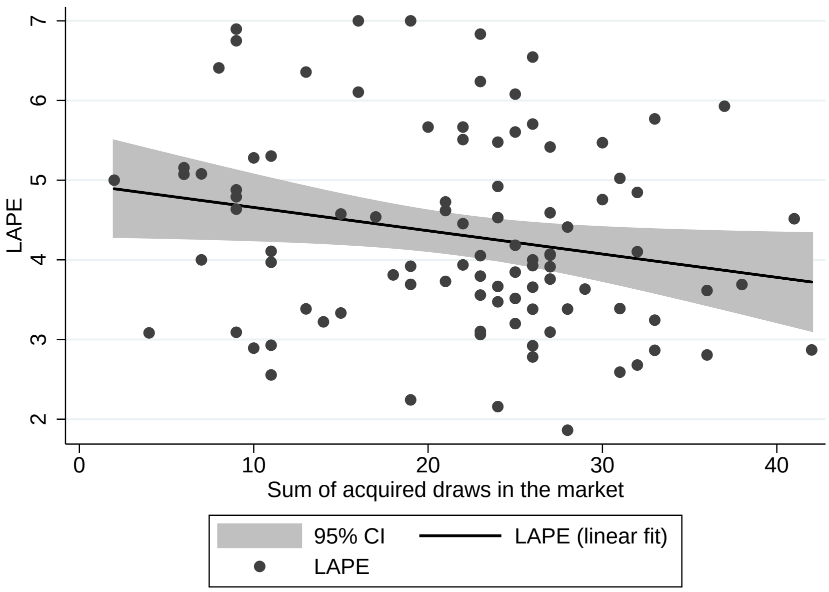

Figure 3.4: Scatter plot and regression line of the linear absolute prediction error depending on the number of acquired draws in the market.

Forecast error. Finally we can assess the quality of market prices as forecasting tools by measuring the absolute prediction error, i.e., the difference between the price and the final value of the asset. Let $t = 1, 2, \ldots, T_{sr}$ index the transactions in an experimental round r of session s. Then the linear absolute prediction error of the average transaction price is defined as

$$\text{LAPE}_{sr} = |T_{sr}^{-1} \sum_{t=1}^{T_{sr}} \text{Price}_{srt} - 10 \cdot \theta_{sr}|, \;\; \theta_{sr} \in \{0, 1\}, \tag{3.2}$$

which can take any value from 0 (perfect forecast) to 10 (worst forecast).[17]

Before discussing the regressions, Figure 3.4 plots LAPE for every market with the number of acquired draws among all traders in the market. The fitted line clearly shows a negative relationship: The more information in the market, the lower the prediction error of the average market price. The slope coefficient for the sum of draws in the market is $-.029$ ($t = -2.11$, $p = 0.038$).[18]

Table 3.4 presents regressions with LAPE as dependent variable, with each round of

[17]Some studies investigating information aggregation compare market prices to the fully revealing rational expectations price given all information in the market (e.g., Healy et al., 2010). Since information is endogenous in our setting, we prefer to use as a benchmark how close market prices are to the perfect forecast.

[18]We can think of at least two possible channels by which additional information in the market reduces the LAPE. First, with more informed traders, there will be fewer traders who are willing to accept disadvantageous transactions due to being uninformed, thus reducing the prediction error of prices. Second, informed traders will compete with each other when trying to make use of their additional information in transactions with uninformed traders. This competition will drive the price towards a price level reflecting the information held by informed traders.

each session (market) as one observation.[19] Column (2) estimates the effect of the number of informed traders, defined as traders with NumAcDraws > 0, and of the sum of acquired draws on the forecast error. Only the sum of acquired draws has a significant effect, and if the number of informed traders is held constant, then another signal in the market reduces the prediction error by about 0.05. Consequently, information acquisition in the aggregate can reduce the prediction error noticeably. For example, instead of giving the true state of nature, say, a 60% chance, 10 additional signals improve this prediction to 65% (typically more than 10 additional draws are acquired in the market, see Figure 3.4). Column (2) also tells us that, if we had the choice between getting more subjects informed while holding the sum of draws constant in the market, or holding the number of informed subjects constant but increasing the sum of draws, then the latter is associated with a lower prediction error while the former is not. Note, however, that this conclusion is drawn from a sample where most subjects are informed (about 60% acquire information, see Table 3.1).

Column (3) estimates the effect of the sum of acquired draws in the market separately for all treatments. More acquired draws significantly reduce the prediction error in the high endowment and in the public draws treatment (10% level), but not in the baseline treatment. In the high endowment treatment, a lower prediction error can be explained by the fact that richer traders—who tend to be the ones that acquire more information (Table 3.2)—also trade more (Table C.4 in the appendix). In the public draws treatment, initial information is weak, so we would expect an additional draw to have a greater impact all else equal, and this is exactly what the point estimates show.

Result 9. *More information acquisition improves the accuracy of the market price.*

3.5 Conclusion

We investigated how asset markets, and in particular prediction markets, not only aggregate existing information, but also induce traders to acquire new information. To this end, we designed an experiment where traders decide whether to acquire information at a cost, and subsequently trade an asset whose value depends on the imperfectly known realization of the state of nature. We find that traders acquire substantial amounts of information, and that they are able to realize higher gross investment returns when better informed. However, we find that informed traders only recover slightly more than half of the incurred information costs via trading. As a consequence, informed traders on average obtain lower profits net of information costs compared to uninformed traders.

This "overacquisition" of information may be a novel explanation why prediction markets typically beat other tools at forecasting uncertain outcomes. We find that more acquired

[19] Prior studies used the median price or the last transaction price rather than the mean price, and the squared prediction error rather than LAPE. Our results are very similar with these. Also see Table C.3 in the appendix.

Table 3.4: What drives prediction market forecast errors?

	(1) OLS	(2) OLS	(3) OLS
Dependent variable	LAPE	LAPE	LAPE
PubDrawsTreat	-0.267	-0.272	
	(0.311)	(0.310)	
HighEndowmentTreat	-0.470	-0.295	
	(0.298)	(0.309)	
SumAcDraws		-0.051**	
		(0.025)	
NumSubInformed		0.115	
		(0.087)	
SumAcDraws in BaselineTreat			-0.021
			(0.018)
SumAcDraws in PubDrawTreat			-0.033*
			(0.017)
SumAcDraws in HighEndowTreat			-0.029**
			(0.014)
Constant	4.559***	4.752***	4.914***
	(0.212)	(0.407)	(0.338)
R^2	0.03	0.07	0.05
Observations	96	96	96

Note: LAPE is the linear absolute prediction error. The baseline treatment is private draws and homogeneous endowment. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

information in the market leads to lower forecast errors on average. Thus, although traders make losses individually, the excess information in the market makes prices better predictors of the underlying outcome. Our overacquisition finding has striking similarities to the well established overdissipation phenomenon in the contest literature, where contestants tend to invest too many resources relative to the available prize in order to gain an edge. Being a zero-sum game, the prediction market is similar to a contest, and investing in information is one way to get an advantage over other participants. Interestingly, we find no evidence of learning to acquire less information over time, or to make better use of the acquired information over time. Hence, traders systematically pick suboptimal actions, which suggests that existing theories assuming rational behavior and preferences depending only on net returns—such as (noisy) rational expectations equilibrium—cannot explain the information acquisition behavior.

By investigating information acquisition on the individual level, we find that traders tend to invest more in information if their initial information about the asset value is inconclusive or when the available information publicly available. These findings add to our understanding of why prediction market prices have a high predictive power, as traders compensate for weak initial information by acquiring more. Moreover, in situations where the main source of initial information is public, traders tend to acquire more private information to gain an advantage in the transactions. Part of this newly acquired private information is aggregated and revealed in market prices, thus prediction markets typically improve on forecasting tools relying on public information. In line with previous studies, we find that market prices are mostly well calibrated, i.e., prices are good estimates of the probability of the outcome. However, we also find that the price only imperfectly reflects all the information available in the market. Compared to the Bayesian full information price prediction, market prices underreact substantially to information.

Overall, our results provide insights into why prediction markets work in practice, even though they do not always work as predicted by rational expectations theory. We confirm that prediction markets can be used as information aggregation mechanisms to inform decision makers about the probability of future events. Our main conclusion is that the prediction market role goes beyond mere aggregation of information, and that they induce traders to acquire information exceeding standard theory predictions, thus improving price accuracy.

In our experiment, we investigated one particular mechanism for forecasting and information aggregation. It is an open question whether other information aggregation mechanisms such as incentivized iterative polls or market scoring rules can induce similar amounts of information acquisition. Iterative polls, for example, rely on participants to share their information with others in a cooperative manner, and all participants receive the same payoff, which requires subsidization of the mechanism (see Healy et al., 2010). Consequently, iterative polls lack the adversarial contest nature of zero-sum prediction markets. Comparing information acquisition with such forecasting tools might further add to our knowledge on what motivates costly information acquisition.

Appendix A

Appendix to chapter 1

A.1 Alternative definition of rational expectations equilibrium without noise

This section gives an alternative definition of rational expectations equilibrium requiring market clearing and optimal demands of traders given the available information, instead of imposing conditions on the equilibrium price function directly. It furthermore shows that the main result can be proven using this definition of rational expectations equilibrium as well.

Before defining the equilibrium concept, we need to define endowments and budget constraints of traders, since risk neutrality would otherwise imply unbounded positions. Each trader $j \in [0,1]$ is initially endowed with one unit of the risky asset and $0 < \omega < \infty$ units of cash, where ω is enough to match the asset value in every state of nature. Traders cannot sell more than one unit of the asset and may not spend more than ω of their cash. The net demand function of trader j is denoted by $x_j(p, s_j)$, which may depend on the private information s_j and the price (and information contained therein) p. A rational expectations equilibrium in this setting is defined as follows, where both the $p(\boldsymbol{s})$ and $i(p, s_p)$ functions are known in equilibrium.

Definition 9. *A rational expectations equilibrium (REE) requires*

i. a measurable price function $p(\boldsymbol{s})$ that clears the market:

$$\int_0^1 x_j(p = p(\boldsymbol{s}), s_j)dj = 0 \text{ a.s.},$$

ii. all traders j use optimal demand strategies given the available information,

$$x_j(p, s_j) \in \arg\max_x \; \mathbb{E}_\theta[a(\theta, i(p, s_p))x | p = p(\boldsymbol{s}), s_j] \text{ s.t. } xp \leq \omega, \; -1 \leq x, \; \forall j \in [0,1],$$

iii. the policy maker sets an optimal policy given available information,

$$i(p, s_p) \in \arg\max_i \ \mathbb{E}_\theta[u(\theta, i)|p = p(\boldsymbol{s}), s_p].$$

With this definition, we get the same result as before.

Corollary 5. *A fully revealing (definition 2) rational expectations equilibrium exists if and only if condition 1 holds.*

Proof. Necessity: Existence of a fully revealing rational expectations equilibrium implies that condition 1 holds. If the equilibrium is fully revealing, then by definition 2,

$$|\{t \in \boldsymbol{S} : \Pr(\boldsymbol{s} = t|p(\boldsymbol{s}) = p, s_p) > 0\}| \leq 1 \ \forall s_p, \forall p.$$

Taking similar steps as in the proof of Theorem 2, if $\Pr(\boldsymbol{s} = t|p(\boldsymbol{s}) = p, s_p)$ in (1.17) is positive for exactly one $t \in \boldsymbol{S}$ for all p in the image of $p(\boldsymbol{s})$, then the antecedent of condition 1 is false and the condition holds.

If $\Pr(\boldsymbol{s} = t|p(\boldsymbol{s}) = p, s_p) > 0$ for more than one $t \in \boldsymbol{S}$, then (1.17) and full revelation imply that whenever $p(t) = p(t')$ for $t \neq t' \in \boldsymbol{S}$, so that $\Pr(p(t) = p) \cdot \Pr(p(t') = p) > 0$, then $\Pr(s_p|\boldsymbol{s} = t) \cdot \Pr(s_p|\boldsymbol{s} = t') = 0$. Thus, the consequent of condition 1 is true and the condition holds.

Sufficiency: Condition 1 implies that a fully revealing rational expectations equilibrium exists. We prove this by construction. Suppose traders conjecture the (accurate) price function $p(\boldsymbol{s}) = v(\boldsymbol{s})$. Theorem 2 shows that condition 1 implies that a fully revealing and accurate price function exists. It is in fact unique, since we assumed that the optimal policy is unique, hence $p(\boldsymbol{s}) = v(\boldsymbol{s})$ is the price function that fully reveals $\boldsymbol{s}$ to the policy maker. Clearly, this price function, which reveals the expected asset value given the combined trader information $v(\boldsymbol{s})$ to traders, also clears the market, since traders are indifferent between buying and selling ($\mathbb{E}_\theta[a(\theta, i(p, s_p))|p = p(\boldsymbol{s}), s_j] = \mathbb{E}_\theta[a(\theta, i(p, s_p))|p = p(\boldsymbol{s})] = p(\boldsymbol{s}) = v(\boldsymbol{s})$), thus for example $x_j(p = p(\boldsymbol{s}), s_j) = 0$ is optimal. Finally, the policy maker sets her optimal policy given $\boldsymbol{s}$, which leads to asset value $v(\boldsymbol{s})$ and makes the conjecture self-fulfilling. $\square$

A.2 Full revelation using Perfect Bayesian Nash equilibrium

This section shows that condition 1, which is necessary and sufficient for fully revealing rational expectations equilibria (Corollary 5), is also necessary and sufficient for full revelation using the Perfect Bayesian Nash equilibrium concept in a continuum economy with perfectly correlated information among traders.

To derive the Perfect Bayesian Nash equilibrium (PBNE), we need to specify the trading environment. The financial market consists of a continuum (of mass 1) of risk neutral traders,[1] who have an endowment in cash $0 < \omega < \infty$ and one unit of the risky asset each, where ω is enough to match the asset value in every state of nature. In this section, I shall only consider the cases where trader signals s_j are perfectly correlated, $s_j = c \in \mathbb{R}\ \forall j$, $\boldsymbol{s} = c$, drawn from density $f(s_j|\theta)$, which includes the perfect information case ($s_j = \theta\ \forall j$). I confine attention to the perfectly correlated information structure, since fully revealing REE in differential information economies are not generally implementable with a market mechanism as Bayesian Nash equilibrium (see, e.g., Vives, 2010). Thus, the role of the market in this section is not aggregation of trader information as in REE, since every trader has the same piece of information, but revelation of the information.

Denote the set of all possible excess demand functions for traders by $\mathcal{X}$. An excess demand function x_j specifies for each p and s_j by how many units trader j would like to change his initial holding of the risky asset. Unlike REE, we have to specify how prices are set in this trading environment, as a Walrasian auctioneer is not available. To be as general as possible, I merely require that prices are a function of the demand function profile. Any price setting algorithm that conditions on demand functions (e.g., market orders, limit orders, etc.) in determining the price fulfills the definition. Examples include double auctions, Dutch auctions, or automated market makers that react to orders.

Definition 10. *A generic price finding rule $\Gamma : \mathcal{X}^N \to \mathbb{R}$ maps all profiles of excess demand functions $\boldsymbol{x}(\boldsymbol{s}) := (x_1(p, s_1 = \boldsymbol{s}), x_2(p, s_2 = \boldsymbol{s}), \dots)$ into a price p^*.*

Price $p^* = \Gamma(\boldsymbol{x}(\boldsymbol{s}))$ need not be a market clearing price, since a market clearing price need not exist for all profiles of excess demand functions. However, the following definition requires market clearing in equilibrium.

Definition 11. *A Perfect Bayesian Nash equilibrium (PBNE) requires*

i. an expected utility maximizing excess demand function $x_j(p, s_j)$ using information set $\{p(\boldsymbol{s}) = p, s_j\}$ given $\boldsymbol{x}_{-j}$ s.t. $-1 \le x_j, x_j p \le \omega$ for all traders j at $t = 1$,

ii. market clearing in the financial market at $t = 1$, i.e., a price function $p(\boldsymbol{s}) = \Gamma(\boldsymbol{x}(\boldsymbol{s}))$ generated by a generic price finding rule, such that

$$\int x_j(p^* = p(\boldsymbol{s}), s_j = \boldsymbol{s})dj = 0 \text{ (a.s.)},$$

iii. policy maker beliefs $\mu(t|p^*) = \Pr(\boldsymbol{s} = t|p^* = p(\boldsymbol{s}), s_p)$, $\sum_{t\in \boldsymbol{S}} \mu(t|p^* = p(\boldsymbol{s}), s_p) = 1$, about financial market behavior, derived by Bayes' rule whenever possible,

[1] The continuum assumption implies that individual traders do not affect the price. Results are not generally the same with finitely many traders, because in such a strategic setting any single trader can change the message sent to the policy maker (price) by changing his own investment strategy, making this a market signaling game, which is beyond the scope of this paper.

iv. and an expected utility maximizing policy $i(s_p, p^*, \mu) \in I$ for the policy maker at $t = 2$.

Beliefs μ refer to the financial market as a whole, represented by the market clearing price p^*, rather than to actions of individual traders. This is a realistic way to model the interaction, as a policy maker typically cannot track investments by specific traders in anonymous markets, especially not in large markets as considered here. The price function $p(\boldsymbol{s})$ is known in equilibrium via Bayes' rule. There is no need to specify off equilibrium path beliefs, because no trader can influence the price alone, i.e., the price with a single deviation can never be off the equilibrium path.

The following proposition shows that condition 1 is necessary and sufficient for the possibility of a fully revealing PBNE.

Proposition 29. *If there is a continuum of traders and trader signals are perfectly correlated, then a fully revealing PBNE exists if and only if condition 1 is fulfilled.*

Proof. Necessity: If a fully revealing PBNE exists, then condition 1 holds. First, a fully revealing equilibrium implies that $p(\boldsymbol{s}) = v(\boldsymbol{s})$. Suppose not ($p(\boldsymbol{s}) \neq v(\boldsymbol{s})$ for some $\boldsymbol{s} \in \boldsymbol{S}$), then market clearing implies there exists at least one trader who can profit by deviating from the candidate strategy x_j to $\hat{x}_j > x_j$ if $v(\boldsymbol{s}) > p(\boldsymbol{s})$ or to $\hat{x}_j < x_j$ if $v(\boldsymbol{s}) < p(\boldsymbol{s})$, since $v(\boldsymbol{s})$ is the expected value of the asset given full revelation, and the deviation in a continuum does not change prices. Thus, full revelation implies either that $v(\boldsymbol{s})$ is invertible, so that $p(\boldsymbol{s})$ perfectly reveals the state, or that $v(\boldsymbol{s})$ is not invertible but $p(\boldsymbol{s}) = v(\boldsymbol{s})$ together with policy maker signal s_p always reveals $\boldsymbol{s}$, i.e., condition 1 holds.

Sufficiency: If condition 1 holds, then a fully revealing PBNE exists. By construction. Define $\overline{x}_j(p, \boldsymbol{s}) = \omega/p > 0$ as j's best response to $v(\boldsymbol{s}) > p$ and $\underline{x}_j(p, \boldsymbol{s}) = -1$ as j's best response to $v(\boldsymbol{s}) < p$. If condition 1 holds, then the following strategy for all j forms a fully revealing PBNE:

$$x_j(p, \boldsymbol{s}) = \begin{cases} \overline{x}_j(p, \boldsymbol{s}) > 0 & \text{if } v(\boldsymbol{s}) > p, \\ \underline{x}_j(p, \boldsymbol{s}) < 0 & \text{if } v(\boldsymbol{s}) < p, \\ 0 & \text{if } v(\boldsymbol{s}) = p. \end{cases}$$

Clearly, market clearing occurs if and only if $p = v(\boldsymbol{s})$ for all $\boldsymbol{s} \in \boldsymbol{S}$. Due to condition 1, the equilibrium candidate price function $p(\boldsymbol{s}) = v(\boldsymbol{s})$ is fully revealing. There is no profitable deviation, since x_j best responds to mispricing by going long if $p < v(\boldsymbol{s})$ and short if $p > v(\boldsymbol{s})$. At $p = v(\boldsymbol{s})$, no deviation changes the payoff, since individual traders cannot influence the price and the asset is priced at its expected value. The policy maker obtains her first best given all available information, so she has no incentive to deviate either. □

Thus, if traders hold the same information, then condition 1 determines the possibility of full revelation in both major equilibrium concepts used in financial economics, REE and PBNE, in large markets. However, if traders have differential information, then the fully revealing REE equilibrium is not implementable as PBNE using a generic price finding rule in a market mechanism, as is well known in financial economics (e.g., Vives, 2010).

Appendix B

Appendix to chapter 2

In this section, we study the robustness of our main results when the production technology is nonlinear. We assume that aggregate investment X translates into supply according to production function

$$x_{sup}(X) = X^\lambda = \left[\int \hat{x}_i di\right]^\lambda, \ \lambda > 0.$$

Although this looks as if only one firm produces the novel good, the production function is up to a constant factor identical to a situation where $1 \leq M < \infty$ firms receive an $1/M$-share of the investment and produce, so that aggregate supply is given by

$$x_{sup}(X) = M\left[\int \hat{x}_i/M di\right]^\lambda = M^{1-\lambda}\left[\int \hat{x}_i di\right]^\lambda.$$

For $0 < \lambda < 1$ the production function will be concave (decreasing returns to scale), and for $\lambda > 1$ it will be convex (increasing returns to scale). $\lambda = 1$ is the linear case considered throughout the main part of the paper. For $\lambda > 1$, we require $1/\lambda > \alpha$, otherwise the planner's problem may have a corner solution.

Consumer demand for given prices remains unchanged:

$$x_i(p) = \left(\frac{\alpha\theta_i}{p}\right)^{1/(1-\alpha)}.$$

The generalized market clearing condition and spot market price is

$$X^\lambda = s\left(\frac{\alpha}{p}\right)^{1/(1-\alpha)} \iff p = \alpha\left(\frac{s}{X^\lambda}\right)^{1-\alpha}.$$

The social optimum

We first determine the planner's solution for the optimal aggregate state dependent investment X^* in the novel good. In the aggregate, market clearing requires that $x_i = x_{sup}/s$, where $s = \int \theta_i di$, x_i is the symmetric consumption level for $\theta_i = 1$ types in the population, and x_{sup} is the aggregate supply (or production) of the novel good. The cost function for producing the novel good is $c(x_{sup}) = RX = Rx_{sup}^{1/\lambda}$, since every unit of investment X has an opportunity cost of R, and the marginal cost is $MC_2 = x_{sup}^{1/\lambda-1}R/\lambda$. In the social optimum, the marginal rate of substitution for a $\theta_i = 1$ consumer has to equal the ratio of marginal costs of production (investment) of the two goods,

$$\begin{aligned} MRS = \frac{MU_1}{MU_2} = -\frac{1}{\alpha x_i^{\alpha-1}} = -\frac{1}{\alpha(x_{sup}/s)^{\alpha-1}} \overset{!}{=} -\frac{MC_1}{MC_2} = -\frac{1}{x_{sup}^{1/\lambda-1}R/\lambda} \\ \iff x_{sup}^* = \left[\frac{\lambda\alpha}{R}s^{1-\alpha}\right]^{\frac{\lambda}{1-\lambda\alpha}} \iff X^* = {x_{sup}^*}^{1/\lambda} = \left[\frac{\lambda\alpha}{R}s^{1-\alpha}\right]^{\frac{1}{1-\lambda\alpha}}. \end{aligned} \tag{B.1}$$

Consequently, the optimal aggregate investment X^* depends nonlinearly on state s whenever $\lambda \neq 1$. The Pareto-optimal investment yields a market clearing price p^*,

$$x_{sup}^* = \left[\frac{\lambda\alpha}{R}s^{1-\alpha}\right]^{\frac{\lambda}{1-\lambda\alpha}} \overset{!}{=} x_i s = s\left(\frac{\alpha}{p}\right)^{\frac{1}{1-\alpha}} \iff p^* = \alpha s^{1-\alpha}\left(\frac{R}{\lambda\alpha s^{1-\alpha}}\right)^{\frac{\lambda-\lambda\alpha}{1-\lambda\alpha}},$$

which depends on state s whenever $\lambda \neq 1$. Moreover, $p^* < R \iff \lambda < 1$. Intuitively, if $\lambda < 1$, then firms can produce more for a given aggregate investment $X < 1$ compared to the linear case, since $X^\lambda > X \iff \lambda < 1$. Thus, in the social optimum, the planner produces more the smaller λ, which means p^* decreases when λ decreases.

Market investment

Given sufficient wealth, $\theta_i = 1$ types will invest such that $\mathbb{E}_s[p|\theta_i = 1] = R$ in any market equilibrium. For $\lambda < 1$ this implies underprovision of funds by the market compared to the planner's solution, because more than the investment/production that results in $p = R$ is efficient (see previous section).

As before, we assume there are two groups of mass 1/2 each, where a share $s_1 \in \{1-\beta, \beta\}$ and $s_2 \in \{1-\beta, \beta\}$ of consumers is interested in the novel good, respectively. Moreover, we assume s_1 and s_2 are independently distributed. If all consumers have enough wealth to invest, then the investment $\hat{x}_e$ in a symmetric equilibrium where all $\theta_i = 1$ types invest is

$$\begin{aligned} R = \mathbb{E}_s[p|\theta_i = 1] &= \alpha\mathbb{E}_s\left[\left(\frac{\int \theta_i di}{(\int \hat{x}_e di)^\lambda}\right)^{1-\alpha}\middle|\theta_i = 1\right] \\ &= \alpha\beta/2\left(\frac{\beta}{(\beta\hat{x}_e)^\lambda}\right)^{1-\alpha} + \alpha(1-\beta)/2\left(\frac{(1-\beta)}{((1-\beta)\hat{x}_e)^\lambda}\right)^{1-\alpha} + \alpha/2\left(\frac{1/2}{(\hat{x}_e/2)^\lambda}\right)^{1-\alpha} \\ \iff \hat{x}_e &= \left(\alpha/R\left[\beta^{(1-\lambda)(1-\alpha)+1}/2 + (1-\beta)^{(1-\lambda)(1-\alpha)+1}/2 + 2^{-1-(1-\lambda)(1-\alpha)}\right]\right)^{\frac{1}{\lambda(1-\alpha)}}. \end{aligned} \tag{B.2}$$

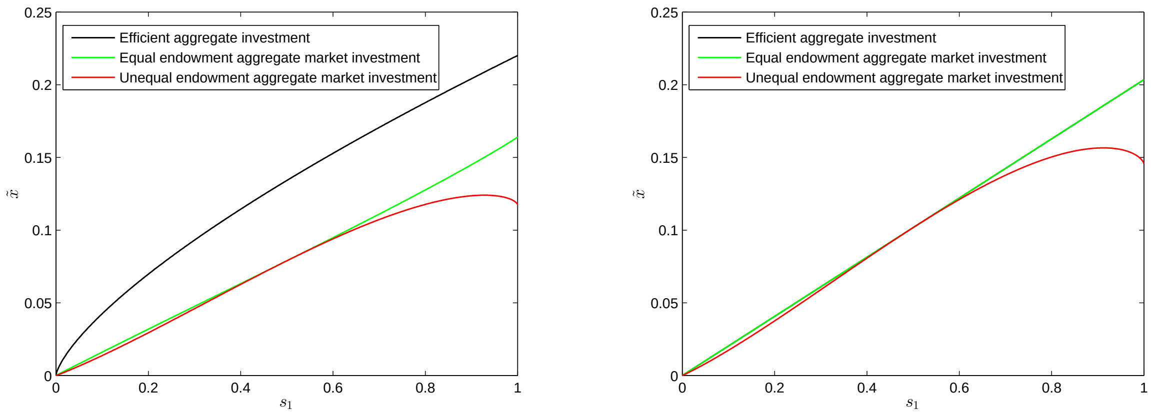

(a) $R = 1.1,\ \alpha = 0.8,\ \lambda = 0.9$ (b) $R = 1.1,\ \alpha = 0.8,\ \lambda = 1$

Figure B.1: Plot of aggregate investment X in states $s_1 = s_2 = \beta$, depending on the share of interested consumers s_1.

In a symmetric equilibrium where only half of the population has wealth to invest ("unequal wealth"), the symmetric investment $\hat{x}_u$ by wealthy $\theta_i = 1$ types fulfills

$$
\begin{aligned}
R &= \alpha \mathbb{E}_s \left[\left(\frac{\int \theta_i di}{(\int \hat{x}_u di)^\lambda} \right)^{1-\alpha} \middle| \theta_i = 1 \right] \\
&= \alpha\beta/2 \left(\frac{\beta}{(\beta \hat{x}_u/2)^\lambda} \right)^{1-\alpha} + \alpha(1-\beta)/2 \left(\frac{(1-\beta)}{((1-\beta)\hat{x}_u/2)^\lambda} \right)^{1-\alpha} \\
&+ \alpha\beta/2 \left(\frac{1/2}{(\beta \hat{x}_u/2)^\lambda} \right)^{1-\alpha} + \alpha(1-\beta)/2 \left(\frac{1/2}{((1-\beta)\hat{x}_u/2)^\lambda} \right)^{1-\alpha} \\
\iff \hat{x}_u &= \Bigg(\alpha/R \Big[\beta^{1+(1-\lambda)(1-\alpha)} 2^{\lambda(1-\alpha)-1} + (1-\beta)^{1+(1-\lambda)(1-\alpha)} 2^{\lambda(1-\alpha)-1} \\
&+ \beta^{1-\lambda(1-\alpha)} 2^{-1-(1-\lambda)(1-\alpha)} + (1-\beta)^{1-\lambda(1-\alpha)} 2^{-1-(1-\lambda)(1-\alpha)} \Big] \Bigg)^{\frac{1}{\lambda(1-\alpha)}}.
\end{aligned}
\tag{B.3}
$$

Thus, if all consumers invest, aggregate investment is $X_e = \int_0^1 \theta_i \hat{x}_e di$, and if group 2 cannot invest, aggregate investment is $X_u = \int_0^{1/2} \theta_i \hat{x}_u di$. In Figure B.1, we consider the two states $s_1 = s_2 = s$, where the realization of both random variables is the same, in order to make aggregate investment with $\hat{x}_e$ and $\hat{x}_u$ comparable. The figure plots the aggregate investment depending on the share of interested consumers when $s = s_1 = s_2$ for a specific parameter profile (R, α, λ) with concave and linear production technology. The plots depict all $\beta \in [0,1]$ by setting $\beta = s$.

The left figure shows the efficient aggregate investment (black line) for $\lambda = 0.9$, which is concave in the share of interested consumers s, since the average cost of production is increasing in s. It shows that the market invests less than what would be efficient in equilibrium. And the market investment if all consumers invest (green line) is weakly larger than the market investment if only one group invests, i.e., weakly more efficient (proven in Proposition 30 for all $\lambda < 1$).

The reason why aggregate investment tends to be larger when all consumers can invest is most easily seen in the linear case $\lambda = 1$. Also in this special case, aggregate investment when all consumers invest is larger compared to when not all consumers can invest (see the right plot in Figure B.1). Crowdinvestors aim to equalize the expected return of investment—which is equal to the market clearing price of the novel good p—with the risk free rate R. The expected market clearing price is given by

$$\mathbb{E}_s[p|\theta_i = 1] = \alpha \mathbb{E}_s \left[\left(\frac{s}{X(s)} \right)^{1-\alpha} \middle| \theta_i = 1 \right].$$

In the linear case, if all consumers invest, the price is state independent and equal to R. Thus, there is no price risk. If not all consumers can invest, however, then prices vary depending on the state: It is higher if more poor consumers have a preference for the good ($s_2 = \beta$), and lower if not ($s_2 = 1 - \beta$), keeping s_1 and therefore aggregate investment constant. Because the market clearing price is a concave function of s for a given aggregate investment, crowdinvestors in the economy with unequal wealth and price risk expect a lower price than crowdinvestors in the economy with equal wealth for the same aggregate investment X, i.e.,

$$\mathbb{E}_s[p_e|\theta_i = 1, X] > \mathbb{E}_s[p_u|\theta_i = 1, X],$$

which follows from the strict concavity of the price and Jensen's inequality. Thus, because their investment return expectations are more optimistic, crowdinvestors invest more in the economy where everyone invests. Intuitively, although crowdinvestors are risk neutral, the investment return is concave in the random variable, so they need higher returns with more uncertainty for the same expected return, leading to reduced investment if the investment return is risky. The intuition carries over to the nonlinear case $\lambda \neq 1$, because price risk in the economy where all consumers can invest is lower—since aggregate investment scales up monotonically with the share of interested consumers—than in the economy where only wealthy consumers can invest. This is reflected in the price expectations in (B.2) and (B.3), where the price can take three different values depending on the state when all consumers invest, but four different values if not all can invest.

Proposition 30. *Consider the states in which $s_1 = s_2$. If $\lambda < 1$, then (ex post) utilitarian welfare is weakly larger if both groups invest compared to the case where only one group can invest, and strictly larger if $\beta \neq 1/2$.*

Proof. From (B.1), if $\lambda < 1$, the efficient aggregate investment is larger than the investment leading to price R. Thus, we have to show that the aggregate investment when both groups invest is weakly larger than the aggregate investment when only one group can invest. The corresponding condition is

$$
\begin{aligned}
&(s/2+s/2)\left(\alpha/R\left[\beta^{(1-\lambda)(1-\alpha)+1}/2+(1-\beta)^{(1-\lambda)(1-\alpha)+1}/2+2^{-1-(1-\lambda)(1-\alpha)}\right]\right)^{\frac{1}{\lambda(1-\alpha)}}\\
&\geq s/2\Bigg(\alpha/R\Bigg[\beta^{1+(1-\lambda)(1-\alpha)}2^{\lambda(1-\alpha)-1}+(1-\beta)^{1+(1-\lambda)(1-\alpha)}2^{\lambda(1-\alpha)-1}\\
&+\beta^{1-\lambda(1-\alpha)}2^{-1-(1-\lambda)(1-\alpha)}+(1-\beta)^{1-\lambda(1-\alpha)}2^{-1-(1-\lambda)(1-\alpha)}\Bigg]\Bigg)^{\frac{1}{\lambda(1-\alpha)}}\\
&\iff \beta^{(1-\lambda)(1-\alpha)+1}+(1-\beta)^{(1-\lambda)(1-\alpha)+1}+2^{-(1-\lambda)(1-\alpha)}\\
&\geq 2^{-\lambda(1-\alpha)}\Bigg[\beta^{1+(1-\lambda)(1-\alpha)}2^{\lambda(1-\alpha)}+(1-\beta)^{1+(1-\lambda)(1-\alpha)}2^{\lambda(1-\alpha)}\\
&+\beta^{1-\lambda(1-\alpha)}2^{-(1-\lambda)(1-\alpha)}+(1-\beta)^{1-\lambda(1-\alpha)}2^{-(1-\lambda)(1-\alpha)}\Bigg]\\
&=\beta^{1+(1-\lambda)(1-\alpha)}+(1-\beta)^{1+(1-\lambda)(1-\alpha)}+\beta^{1-\lambda(1-\alpha)}2^{-(1-\alpha)}+(1-\beta)^{1-\lambda(1-\alpha)}2^{-(1-\alpha)}\\
&\iff 2^{-(1-\lambda)(1-\alpha)}\geq \beta^{1-\lambda(1-\alpha)}2^{-(1-\alpha)}+(1-\beta)^{1-\lambda(1-\alpha)}2^{-(1-\alpha)}\\
&\iff 2^{\lambda(1-\alpha)}\geq \beta^{1-\lambda(1-\alpha)}+(1-\beta)^{1-\lambda(1-\alpha)} \qquad \text{(B.4)}
\end{aligned}
$$

The right hand side of the inequality is a sum of concave functions, and it achieves its unique maximum at $\beta = 1/2$. Evaluating the RHS at its maximum, the inequality (B.4) changes to

$$2^{\lambda(1-\alpha)} \geq (1/2)^{1-\lambda(1-\alpha)} + (1/2)^{1-\lambda(1-\alpha)}.$$

At the maximum $\beta = 1/2$, the condition holds with equality, and $2^{\lambda(1-\alpha)} > \beta^{1-\lambda(1-\alpha)} + (1-\beta)^{1-\lambda(1-\alpha)}$ whenever $\beta \neq 1/2$ immediately follows. □

The proposition does not imply, however, that welfare is larger in every state if all consumers invest and $\lambda < 1$. If $s_1 = \beta$ and $s_2 = 1-\beta$, then $s = 1/2$. Group 1 is the wealthy one, so for β large enough aggregate investment will be larger in the economy where only one group can invest, because wealthy consumers overestimate aggregate demand in the economy. This is the only state where endowment inequality among consumers may be better in terms of welfare.

Since welfare from an ex post perspective may depend on the state, the main question is whether ex ante welfare (expectation over all four states) is still larger if all consumers invest. We investigate this question numerically in the next section.

Numerical welfare analysis

To compare welfare for equal and unequal wealth among consumers, we assume aggregate wealth in the economy is constant, but in the unequal distribution case half of the consumers holds investment endowment $2w$ whereas the other half holds zero. In the equal distribution case, every consumer has $w_i = w > 0$ to invest. Wealth w and income y are chosen so that budget constraints when investing in x or consuming x are never binding.

Ex ante utilitarian welfare in the market with equal wealth distribution (all consumers can invest) is

$$W_e = \mathbb{E}_s \left[s(x_i(p_e(s))^\alpha + y + R(w - \hat{x}_e) - p_e(s)(x_i(p_e(s)) - \hat{x}_e)) + (1-s)(y + Rw) \right],$$

where $p_e(s)$ denotes the market clearing price in state s. Similarly, ex ante utilitarian welfare in the market with unequal wealth (only half of consumers can invest) is

$$\begin{aligned} W_u = \mathbb{E}_s[s_1(x_i(p_u(s))^\alpha + y + R(2w - \hat{x}_u) - p_u(s)(x_i(p_u(s)) - \hat{x}_u)) \\ +(1-s_1)(y + 2Rw) + (1-s_2)y + s_2(x_i(p_u(s))^\alpha + y - p_u(s)x_i(p_u(s)))]/2. \end{aligned}$$

The following results assume $\beta > 1/2$, since there is no demand uncertainty for $\beta = 1/2$, and investment and welfare is always the same for equal and unequal wealth distribution. Result 1 is based on numerical calculations with the following parameter values in all possible combinations, where (following the Matlab syntax) $\{a : z : b\} := [a, b] \cap \{a + kz\}_{k=0,1,2,...}$ is the parameter grid.

$$\beta \in \{0.6 : 0.1 : 1\}, R \in \{1 : 0.1 : 2\}, \lambda \in \{0.1 : 0.1 : 1\}, \alpha \in \{0.1 : 0.1 : 0.9\}.$$ [1]

Result 1. *For $\lambda \leq 1$, ex ante utilitarian welfare is always strictly larger if all consumers invest compared to the case where only one group invests.*

Thus, the results of section 2.2 generalize to concave production technologies in the sense that market outcomes yield higher (ex ante) welfare when all consumers invest. The reason, as explained in the previous subsection, is that crowdinvestors tend to invest more if there is less price uncertainty, and it is efficient to invest more if the production function is strictly concave.

If the production technology is convex, i.e., $\lambda > 1$, then market investment tends to be lower if only one group can invest, and thus (for $\lambda >> 1$) tends to be closer to the efficient aggregate investment. However, aggregate investment still scales better with consumer preferences if all consumers can invest. This trade-off suggests that welfare is not unambiguously better for one or the other wealth distribution with $\lambda > 1$, which is confirmed in our next result. Result 2 is based on numerical calculations with the following parameter values in all possible combinations (which satisfy $1/\lambda > \alpha$, see above).

$$\beta \in \{0.6 : 0.1 : 1\}, R \in \{1 : 0.1 : 2\}, \lambda \in \{1.1 : 0.1 : 3\}, \alpha \in \{0.1 : 0.1 : 0.9\}, 1/\lambda > \alpha.$$ [2]

[1] We have not found a parameter profile where the results do not hold. Matlab scripts of the numerical calculations are available upon request.

[2] In this parameter profile, the equal wealth economy has larger ex ante total welfare for $\lambda < 1.9$, for

Result 2. *For $\lambda > 1$ and sufficiently close to 1, ex ante utilitarian welfare is larger if all consumers invest compared to the case where only one group invests. For $\lambda >> 1$, ex ante utilitarian welfare is larger if only one group invests.*

Thus, ex ante total welfare is still greater when all consumers invest if λ is close to but above 1. For very strong convexity in the production function, however, an unequal wealth distribution is superior in terms of ex ante welfare.

The intuition is as before: When not all consumers invest, the investment return is more risky, and crowdinvestors tend to invest less. Due to convexity of the production function, production for $X < 1$ is more expensive compared to a linear or concave production function, hence socially optimal production is less than in the linear/concave case. Thus, less investment/production is socially more desirable with convex production technology, which explains why the unequal wealth economy is superior in terms of ex ante welfare for large convexity. The trade-off is that aggregate investment reacts better to changes in consumer preferences when all consumers invest, which from an ex ante point of view is welfare improving. Thus, for small convexity the equal wealth economy still has larger ex ante welfare. Our numerical simulation suggests that the equal wealth economy fares better even for moderate convexity (up to $\lambda = 1.9$, see the footnote above).

$\lambda = 1.9$ it depends on the remaining parameter values, and for $\lambda > 1.9$ the unequal wealth economy is always superior in terms of ex ante welfare.

Appendix C

Appendix to chapter 3

C.1 Additional analyses and robustness checks

Table C.1: Description of variables

Variable Name	Description
AcquiredInformation	Dummy, subject acquired at least one draw in this round (NumAcDraws > 0)
Age	Subject age in years
Bid-ask spread	Difference between lowest sell offer and highest buy offer, measured every 5 seconds
ConclusiveInitDraws	Dummy, subject's two initial draws are either both black or both white in this round
EconBusMajor	Dummy, subject's field of study is economics or business/ management
Experience	Dummy, subject reported to have traded stocks/financial assets or placed bets with bookmakers/online
FirstPriceLAPE	Linear absolute prediction error (see (3.2)) of the first transaction price in this round
GrossProfit	Difference of subject's portfolio value at the end of the round and at its beginning, not accounting for information costs
HighEndowment	Dummy, subject received twice the initial endowment
HighEndowmentTreat	Dummy, all subjects have a 50% chance of receiving twice the initial endowment
LAPE	Linear absolute prediction error of the average transaction price in a round, see (3.2)
Male	Dummy, subject is male
NetProfit	Difference of subject's portfolio value at the end of the round and at its beginning, minus information costs (see (3.1))
NumAcDraws	Number of draws the subject acquired in a round (between 0 and 10)
NumSubInformed	Number of subjects with NumAcDraws > 0; sum of AcquiredInformation over all subjects
NumTrades	Number of trades the subject executed in a round
PubDrawsTreat	Dummy, initial information in this round is observed publicly
RiskAversion	Scale ranging from 0 to 10, inverted risk attitude scale from Dohmen et al. (2011)
SumAcDraws	Sum of acquired draws among all subjects in this round; sum of NumAcDraws
TransactionLAPE	Mean of LAPE (see (3.2)) of all transaction prices for this subject in this round

Table C.2: What drives information acquisition? Nonlinear models

	(1) Poisson	(2) Poisson FE	(3) Poisson FE	(4) Logit FE	(5) Poisson	(6) Logit
Dependent variable	NumAcDraws	NumAcDraws	NumAcDraws	InfoAcquired	NumAcDraws	InfoAcquired
PubDrawsTreat	0.113**	0.151***	0.122**	0.939***	0.092	0.386***
	(0.057)	(0.048)	(0.051)	(0.250)	(0.056)	(0.144)
HighEndowmentTreat	0.273***	0.245***	0.046	0.264	0.083	0.320*
	(0.052)	(0.059)	(0.069)	(0.291)	(0.097)	(0.192)
HighEndowment			0.381***	1.156***	0.311***	0.349
			(0.073)	(0.348)	(0.098)	(0.227)
ConclusiveInitDraws			-0.194***	-0.879***	-0.242***	-0.656***
			(0.058)	(0.208)	(0.072)	(0.172)
RiskAversion					-0.060	-0.187**
					(0.052)	(0.086)
Experience					-0.373**	-0.889**
					(0.158)	(0.384)
Age					-0.018**	0.012
					(0.008)	(0.019)
Male					0.195	0.512
					(0.145)	(0.358)
EconBusMajor					-0.341**	-0.502
					(0.156)	(0.379)
Constant	0.524***				1.618***	1.803**
	(0.040)				(0.400)	(0.818)
Subject Fixed Effects	no	yes	yes	yes	no	no
Cluster SE	no	yes	yes	yes	yes	yes
Log-(Pseudo-)likelihood	-2186.27	-1296.13	-1272.60	-234.21	-2038.14	-619.23
Observations	1080	996	996	612	1068	1068
Clusters		83	83	51	89	89

Note: This table replicates the regressions of Table 3.2 using nonlinear models rather than OLS. In the Poisson fixed effects regressions, observations of some subjects have to be dropped, because these subjects choose zero in all rounds. Similarly, observations of some subjects in the conditional fixed effects logit model have to be dropped, because these subjects always or never acquire information. The baseline treatment is private draws and homogeneous endowment. If applicable, standard errors are clustered at the subject level. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Table C.3: What drives prediction market forecast errors? (Results based on median and last transaction price)

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) OLS
Dependent variable	LAPE Med	LAPE Last	LAPE Med	LAPE Last	LAPE Med	LAPE Last
PubDrawsTreat	-0.343	-0.204	-0.371	-0.147		
	(0.349)	(0.433)	(0.349)	(0.429)		
HighEndowmentTreat	-0.368	-0.832**	-0.204	-0.514		
	(0.335)	(0.415)	(0.348)	(0.428)		
SumAcDraws			-0.059**	-0.065*		
			(0.028)	(0.034)		
NumSubInformed			0.157	0.075		
			(0.098)	(0.120)		
SumAcDraws Baseline					-0.020	-0.042*
					(0.020)	(0.025)
SumAcDraws PubDraw					-0.036*	-0.046**
					(0.019)	(0.023)
SumAcDraws HighEndow					-0.022	-0.060***
					(0.016)	(0.020)
Constant	4.515***	4.273***	4.565***	4.989***	4.838***	5.022***
	(0.239)	(0.296)	(0.458)	(0.564)	(0.380)	(0.465)
R^2	0.02	0.04	0.06	0.10	0.04	0.09
Observations	96	96	96	96	96	96

Note: LAPE is the linear absolute prediction error. LAPE Med is based on the median transaction price of the round, LAPE Last is based on the last transaction price of the round. The baseline treatment is private draws and homogeneous endowment. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

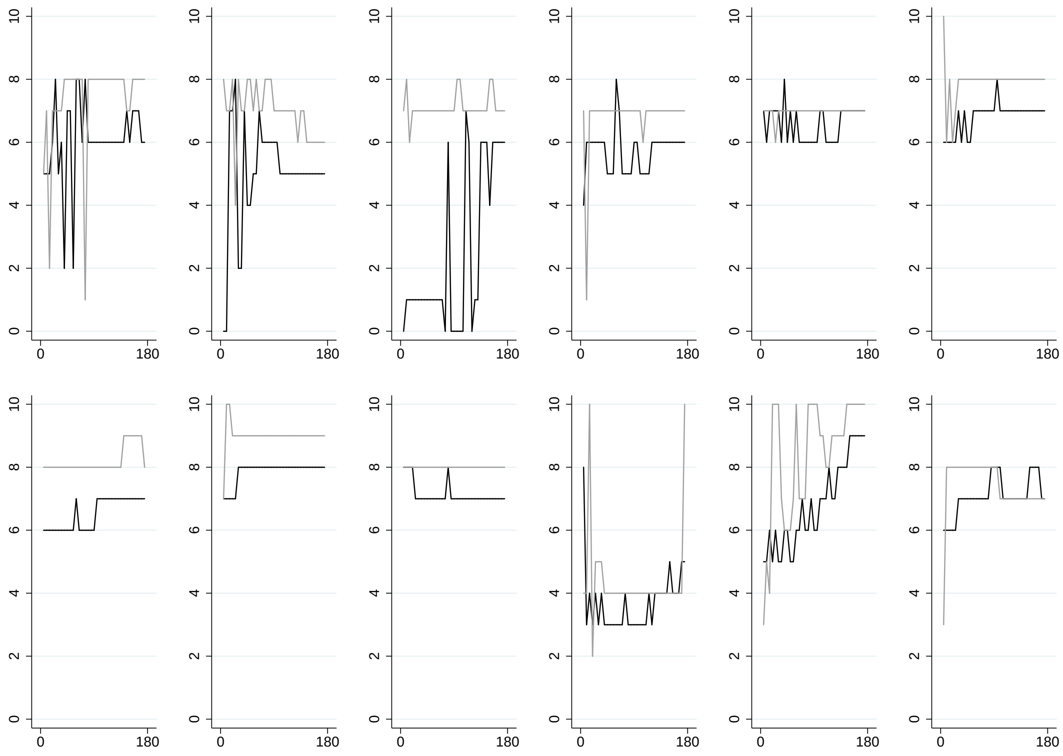

Figure C.1: Bids (black) and asks (grey) in five second intervals for each of the 12 rounds of the first session. As in Figure 3.2, which displays the corresponding transaction price chart, the values are normalized so that the actual value of the asset is always displayed as 10.

C.2 Information acquisition and trading activity

Traders' decisions to invest in information acquisition is motivated by the prospect of higher profits to extract from trading opportunities. We can therefore expect that informed and non-informed traders will be characterized by different trading activities. In particular, we hypothesize that traders who invested in information acquisition are likely to be willing to use this information and therefore trade more. Models with overconfidence could accomodate such a pattern, with overconfident traders being both confident in their ability to use information and in their tendency to trade (Barber and Odean, 2001).

Table C.4 displays OLS regressions with the number of trades per subject as dependent variable. Because every trade involves two parties, the first two columns cluster standard errors at the period-level, while the last two cluster on subject-level. First, columns (2)-(4) show that higher endowment induces more trade. This is consistent with wealth effects, but it might also be driven by the fact that traders with more endowment are on average better informed. In the latter case, we should observe differences in trading behavior depending on the acquired information, which we turn to now.

Information and number of trades. According to column (2), better informed traders are involved in significantly *fewer* transactions than less informed traders. However, this puzzling effect vanishes once we include subject fixed effects in column (3). This could suggest that subjects who tend to trade a lot typically acquire less information, or those who trade less frequently acquire more information. As a consequence, there is a negative relationship without fixed effects, which vanishes once we control for individual time invariant factors. Column (4) estimates the effect of the number of acquired draws separately for traders with high and low endowment. While the effect of more information on the number of trades is practically zero for low endowment traders, it is significantly positive for high endowment traders, with an additional draw increasing the number of trades by about 0.36 on average. Overall, information appears to induce more trade only for high endowment traders, and this finding is not just a wealth effect, which is captured separately by the HighEndowment dummy.

Result 10. *Larger endowment induces more trade, and traders with high endowment trade more if they are better informed.*

These findings suggest another mechanism contributing to prediction market forecast accuracy. First, traders with larger endowment acquire more information (Table 3.2). Thus, they are on average better informed. Second, as we just showed, they trade more. Consequently, the better informed traders have a larger weight in determining market prices. This mechanism, derived in Siemroth (2014), complements the widely studied role of prediction markets as information aggregation tools.

We also investigated the relationship of our demographic variables with four trading behavior variables: the number of trades per round and subject, the share of purchases (rather than sales) of assets, the share of offered (rather than accepted) trades,[1] and the number of trade offers (accepted or not).

Table C.5 displays the OLS regressions with standard errors clustered within rounds. Column (1) reveals that risk aversion and experience are significantly positively related to the number of trades, while age has significant negative relationship. Gender does not have a significant impact on the number of trades, which might be surprising, since overconfidence is typically more pronounced in men and found to increase trade (Barber and Odean, 2001). Column (2) shows that only risk aversion has a significant impact on the share of asset purchases (rather than asset sales). The direction is exactly as expected: More risk averse agents have a lower buy share, i.e., they sell their (risky) assets for cash. Column (3) looks at whether demographic variables can explain whether subjects offer rather than accept trades. Gender is a significant predictor here, with men being two percentage points more likely to offer rather than accept trades. Interestingly, better informed traders (larger

[1] Recall that, for each trade, there is one subject that offers the trade and one that accepts at the offered price.

NumAcDraws) are less likely to offer trades (and more likely to accept), which might suggest that they try to hide their information, since accepting an existing offer reveals less information than publicly posting a new price. This hypothesis is also consistent with the estimates in column (4), where the number of acquired draws and the number of offers made by the subject are strongly negatively related. Experience in stock markets, a major outside of finance/economics, and younger age, on the other hand, are associated with more offers per subject.

Table C.4: How does information acquisition affect the number of trades?

	(1) OLS	(2) OLS	(3) OLS	(4) OLS
Dependent variable	NumTrades	NumTrades	NumTrades	NumTrades
PubDrawsTreat	-0.488	-0.429	-0.438	-0.413
	(0.661)	(0.662)	(0.286)	(0.286)
HighEndowmentTreat	0.681	-0.517	-0.326	-0.317
	(0.547)	(0.581)	(0.325)	(0.327)
HighEndowment		2.760***	2.634***	1.727***
		(0.507)	(0.414)	(0.577)
NumAcDraws		-0.290***	0.082	
		(0.081)	(0.098)	
NumAcDraws with low endowment				-0.033
				(0.113)
NumAcDraws with high endowment				0.358**
				(0.150)
Constant	6.519***	7.009***	6.243***	6.440***
	(0.450)	(0.505)	(0.229)	(0.252)
Subject Fixed Effects	no	no	yes	yes
Cluster SE	Period	Period	Subject	Subject
R^2	0.01	0.04	0.07	0.08
Observations	1080	1080	1080	1080
Clusters	96	96	90	90

Note: NumTrades is the number of trades in a round for each trader. The baseline treatment is private draws and homogeneous endowment. If applicable, standard errors are clustered at the subject level. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Table C.5: Demographics and trading behavior

	(1) OLS	(2) OLS	(3) OLS	(4) OLS
Dependent variable	NumTrades	ShareBuy	ShareOffer	NumOffers
PubDrawsTreat	-0.329	-0.016	0.006**	4.542
	(0.612)	(0.015)	(0.003)	(2.796)
HighEndowmentTreat	-0.323	-0.001	-0.013**	4.480
	(0.545)	(0.027)	(0.005)	(2.952)
HighEndowment	2.661***	-0.017	0.036***	2.263
	(0.485)	(0.037)	(0.009)	(2.474)
NumAcDraws	-0.205***	0.010	-0.003**	-1.251***
	(0.077)	(0.008)	(0.001)	(0.445)
Age	-0.046**	-0.002	0.000	-0.164***
	(0.019)	(0.001)	(0.000)	(0.058)
RiskAversion	0.219**	-0.020***	0.001	0.479*
	(0.094)	(0.006)	(0.001)	(0.266)
Experience	1.012***	-0.028	0.011*	4.151***
	(0.351)	(0.025)	(0.006)	(1.356)
EconBusMajor	0.331	-0.011	0.007	-4.531**
	(0.312)	(0.026)	(0.005)	(1.981)
Male	0.307	-0.022	0.020***	-3.083*
	(0.265)	(0.024)	(0.005)	(1.711)
Constant	6.130***	0.646***	0.061***	15.523***
	(0.701)	(0.054)	(0.012)	(3.194)
Subject Fixed Effects	no	no	no	no
Cluster SE	Period	Period	Period	Period
R^2	0.05	0.02	0.05	0.04
Observations	1068	1041	1041	1068
Clusters	96	96	96	96

Note: NumTrades is the number of trades in a round for each trader. ShareBuy is the share of asset purchases (rather than sales) among trades per round. ShareOffer is the share of trades offered rather than accepted per round. NumOffers is the number of offers made (accepted or not) in a round for each traders. The baseline treatment is private draws and homogeneous endowment. Standard errors are clustered within each round. ***Significant at the 1% level; **significant at the 5% level; *significant at the 10% level.

Bibliography

Ackert, L. F., B. K. Church, and M. Shehata (1997): "Market behavior in the presence of costly, imperfect information: Experimental evidence," *Journal of Economic Behavior & Organization*, 33, 61–74.

Aghion, P. and P. Bolton (1997): "A theory of trickle-down growth and development," *Review of Economic Studies*, 64, 151–172.

Agrawal, A. K., C. Catalini, and A. Goldfarb (2011): "The geography of crowdfunding," NBER Working Paper 16820.

——— (2013): "Some simple economics of crowdfunding," NBER Working Paper 19133.

Ahlers, G. K., D. Cumming, C. Günther, and D. Schweizer (2012): "Signaling in equity crowdfunding," Working Paper.

Allen, B. (1981): "Generic existence of completely revealing equilibria for economies with uncertainty when prices convey information," *Econometrica*, 1173–1199.

Allen, F. (1990): "The market for information and the origin of financial intermediation," *Journal of Financial Intermediation*, 1, 3–30.

Allen, F. and D. Gale (1999): "Diversity of opinion and financing of new technologies," *Journal of Financial Intermediation*, 8, 68–89.

——— (2001): *Comparing Financial Systems*, MIT Press.

Anderson, S. P., J. K. Goeree, and C. A. Holt (1998): "Rent seeking with bounded rationality: An analysis of the all-pay auction," *Journal of Political Economy*, 106, 828–853.

Angrisani, M., A. Guarino, S. Huck, and N. C. Larson (2008): "No-trade in the laboratory," *The BE Journal of Theoretical Economics*, 11.

Banerjee, A. V. and A. F. Newman (1993): "Occupational choice and the process of development," *Journal of Political Economy*, 274–298.

Barber, B. M. and T. Odean (2001): "Boys will be boys: Gender, overconfidence, and common stock investment," *Quarterly Journal of Economics*, 261–292.

Belleflamme, P., T. Lambert, and A. Schwienbacher (2014): "Crowdfunding: Tapping the right crowd," *Journal of Business Venturing*, 29, 585–609.

Belleflamme, P., N. Omrani, and M. Peitz (2015): "The Economics of Crowdfunding Platforms," CORE Discussion Paper 2015/15.

Benveniste, L. M. and P. A. Spindt (1989): "How investment bankers determine the offer price and allocation of new issues," *Journal of Financial Economics*, 24, 343–361.

Bernanke, B. S. and M. Woodford (1997): "Inflation forecasts and monetary policy," *Journal of Money, Credit, and Banking*, 29, 653–684.

Bhattacharya, S. and A. V. Thakor (1993): "Contemporary banking theory," *Journal of Financial Intermediation*, 3, 2–50.

Birchler, U. W. and M. Facchinetti (2007): "Self-destroying prophecies? The endogeneity pitfall in using market signals as triggers for prompt corrective action," Working paper.

Bond, P. and I. Goldstein (2015): "Government intervention and information aggregation by prices," *Journal of Finance*, forthcoming.

Bond, P., I. Goldstein, and E. S. Prescott (2010): "Market-based corrective actions," *Review of Financial Studies*, 23, 781–820.

Boot, A. W. and A. V. Thakor (1997): "Financial system architecture," *Review of Financial Studies*, 10, 693–733.

Breon-Drish, B. (2015): "On Existence and Uniqueness of Equilibrium in a Class of Noisy Rational Expectations Models," *Review of Economic Studies*.

Cabrales, A., O. Gossner, and R. Serrano (2014): "The Inverse Demand for Information and the Appeal of Information Transactions," Working paper.

Carrillo, J. D. and T. R. Palfrey (2011): "No trade," *Games and Economic Behavior*, 71, 66–87.

Castro, V. and R. M. Sousa (2012): "How do central banks react to wealth composition and asset prices?" *Economic Modelling*, 29, 641–653.

Copeland, T. E. and D. Friedman (1992): "The market value of information: Some experimental results," *Journal of Business*, 241–266.

Coval, J. D. and A. V. Thakor (2005): "Financial intermediation as a beliefs-bridge between optimists and pessimists," *Journal of Financial Economics*, 75, 535–569.

Cowgill, B. and E. Zitzewitz (2015): "Corporate Prediction Markets: Evidence from Google, Ford, and Firm X," *Review of Economic Studies*, rdv014.

Davies, J. B., S. Sandström, A. Shorrocks, and E. N. Wolff (2011): "The Level and Distribution of Global Household Wealth," *Economic Journal*, 121, 223–254.

de Mesa, D. and D. C. Webb (1992): "Efficient credit rationing," *European Economic Review*, 36, 1277–1290.

Deck, C. and D. Porter (2013): "Prediction markets in the laboratory," *Journal of Economic Surveys*, 27, 589–603.

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

——— (1999): "On the uniqueness of fully informative rational expectations equilibria," *Economic Theory*, 13, 1–24.

Diamond, D. W. (1984): "Financial intermediation and delegated monitoring," *The Review of Economic Studies*, 51, 393–414.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner (2011): "Individual risk attitudes: Measurement, determinants, and behavioral consequences," *Journal of the European Economic Association*, 9, 522–550.

Dow, J., I. Goldstein, and A. Guembel (2015): "Incentives for information production in markets where prices affect real investment," Working Paper.

Dow, J. and G. Gorton (1997): "Stock Market Efficiency and Economic Efficiency: Is There a Connection?" *Journal of Finance*, 52, 1087–1129.

Edmans, A., I. Goldstein, and W. Jiang (2015): "Feedback Effects, Asymmetric Trading, and the Limits to Arbitrage," Working Paper.

Fama, E. F. (1970): "Efficient capital markets: A review of theory and empirical work," *Journal of Finance*, 25, 383–417.

Finocchiaro, D. and V. Q. Heideken (2013): "Do Central Banks React to House Prices?" *Journal of Money, Credit and Banking*, 45, 1659–1683.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10, 171–178.

Fortune (2014): "Why investors are pouring millions into crowdfunding," http://fortune.com/2014/04/17/why-investors-are-pouring-millions-into-crowdfunding.

Foucault, T. and T. Gehrig (2008): "Stock price informativeness, cross-listings, and investment decisions," *Journal of Financial Economics*, 88, 146–168.

Frazzini, A. (2006): "The disposition effect and underreaction to news," *Journal of Finance*, 61, 2017–2046.

Galor, O. and J. Zeira (1993): "Income distribution and macroeconomics," *Review of Economic Studies*, 60, 35–52.

Goldstein, I. and A. Guembel (2008): "Manipulation and the allocational role of prices," *Review of Economic Studies*, 75, 133–164.

Goldstein, I., E. Ozdenoren, and K. Yuan (2013): "Trading frenzies and their impact on real investment," *Journal of Financial Economics*, 109, 566–582.

Gorton, G. and A. Winton (2003): "Financial Intermediation," in *Handbook of the Economics of Finance*, ed. by M. H. George M. Constantinides and R. M. Stulz, Elsevier, vol. 1A, 431–552.

Greiner, B. (2004): "An online recruitment system for economic experiments," http://mpra.ub.uni-muenchen.de/13513/1/MPRA_paper_13513.pdf.

Gromb, D. and D. Vayanos (2010): "Limits of Arbitrage," *Annual Review of Financial Economics*, 2, 251–275.

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

Grossman, S. J. and J. E. Stiglitz (1980): "On the impossibility of informationally efficient markets," *American Economic Review*, 70, 393–408.

Grüner, H. P. (2003): "Redistribution as a selection device," *Journal of Economic Theory*, 108, 194–216.

Grüner, H. P. and R. Schils (2007): "The Political Economy of Wealth and Interest," *Economic Journal*, 117, 1403–1422.

Hayek, F. (1945): "The Use of Knowledge in Society," *American Economic Review*, 35, 519–530.

Healy, P. J., S. Linardi, J. R. Lowery, and J. O. Ledyard (2010): "Prediction markets: alternative mechanisms for complex environments with few traders," *Management Science*, 56, 1977–1996.

Hellwig, M. F. (1980): "On the aggregation of information in competitive markets," *Journal of Economic Theory*, 22, 477–498.

HPC Wire (2011): "Ford Motor Company Turns to Cloud-Based Prediction Market Software," http://www.hpcwire.com/2011/02/22/ford_motor_company_turns_to_cloud-based_prediction_market_software/.

Huber, J. (2007): "'J'-shaped returns to timing advantage in access to information–Experimental evidence and a tentative explanation," *Journal of Economic Dynamics and Control*, 31, 2536–2572.

Huber, J., M. Angerer, and M. Kirchler (2011): "Experimental asset markets with endogenous choice of costly asymmetric information," *Experimental Economics*, 14, 223–240.

Kim, D. H., D. R. Osborn, and M. Sensier (2005): "Nonlinearity in the Fed's monetary policy rule," *Journal of Applied Econometrics*, 20, 621–639.

Konrad, K. A. (2009): *Strategy and dynamics in contests*, Oxford University Press.

Leland, H. E. and D. H. Pyle (1977): "Informational asymmetries, financial structure, and financial intermediation," *Journal of Finance*, 32, 371–387.

Li, E. and J. S. Martin (2014): "Reputation acquisition in capital formation: An empirical investigation," Working Paper, University of Melbourne.

L'œillet, G. and J. Licheron (2012): "How does the European Central Bank react to oil prices?" *Economics Letters*, 116, 445–447.

Maloney, M. T. and J. H. Mulherin (2003): "The complexity of price discovery in an efficient market: the stock market reaction to the Challenger crash," *Journal of Corporate Finance*, 9, 453–479.

Milgrom, P. and N. Stokey (1982): "Information, trade and common knowledge," *Journal of Economic Theory*, 26, 17–27.

Mollick, E. (2014): "The dynamics of crowdfunding: An exploratory study," *Journal of Business Venturing*, 29, 1–16.

Noussair, C. N. and S. Tucker (2013): "Experimental research on asset pricing," *Journal of Economic Surveys*, 27, 554–569.

Ottaviani, M. and P. N. Sørensen (2007): "Outcome manipulation in corporate prediction markets," *Journal of the European Economic Association*, 5, 554–563.

——— (2009): "Surprised by the parimutuel odds?" *American Economic Review*, 99, 2129–2134.

Page, L. and R. Clemen (2013): "Do prediction markets produce well calibrated probability forecasts?" *Economic Journal*, 123, 491–513.

Palan, S. (2013): "A review of bubbles and crashes in experimental asset markets," *Journal of Economic Surveys*, 27, 570–588.

Pálvölgyi, D. and G. Venter (2015): "Multiple equilibria in noisy rational expectations economies," Working Paper.

Peress, J. (2004): "Wealth, information acquisition, and portfolio choice," *Review of Financial Studies*, 17, 879–914.

Piketty, T. (1997): "The dynamics of the wealth distribution and the interest rate with credit rationing," *Review of Economic Studies*, 64, 173–189.

——— (2014): *Capital in the Twenty-first Century*, Harvard University Press.

Plott, C. R. and S. Sunder (1988): "Rational expectations and the aggregation of diverse information in laboratory security markets," *Econometrica*, 1085–1118.

Prescott, E. S. (2012): "Contingent capital: the trigger problem," *Federal Reserve Bank of Richmond Economic Quarterly*, 98, 33–50.

Radner, R. (1979): "Rational expectations equilibrium: Generic existence and the information revealed by prices," *Econometrica*, 655–678.

Ravn, S. (2012): "Has the Fed Reacted Asymmetrically to Stock Prices?" *The BE Journal of Macroeconomics*, 12.

Reny, P. J. (2008): "Non-cooperative games (equilibrium existence)," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf and L. E. Blume, Palgrave Macmillan, 2nd ed.

Rigobon, R. and B. Sack (2003): "Measuring The Reaction of Monetary Policy to the Stock Market," *Quarterly Journal of Economics*, 118, 639–669.

Saez, E. and G. Zucman (2014): "Wealth inequality in the United States since 1913: Evidence from capitalized income tax data," NBER Working Paper 20625.

Sheremeta, R. M. (2010): "Experimental comparison of multi-stage and one-stage contests," *Games and Economic Behavior*, 68, 731–747.

——— (2013): "Overbidding and heterogeneous behavior in contest experiments," *Journal of Economic Surveys*, 27, 491–514.

——— (2015): "Behavioral Dimensions of Contests," in *Companion to the Political Economy of Rent Seeking*, ed. by R. D. Congleton and A. L. Hillman, Edward Elgar Publishing.

Shleifer, A. and R. W. Vishny (1997): "The limits of arbitrage," *Journal of Finance*, 52, 35–55.

Siemroth, C. (2014): "Why prediction markets work: The role of information acquisition and endogenous weighting," https://ub-madoc.bib.uni-mannheim.de/37452/1/Christoph_Siemroth_14%2D29.pdf, Working paper, University of Mannheim.

Subrahmanyam, A. and S. Titman (1999): "The going-public decision and the development of financial markets," *Journal of Finance*, 54, 1045–1082.

Sun, Y. (2006): "The exact law of large numbers via Fubini extension and characterization of insurable risks," *Journal of Economic Theory*, 126, 31–69.

Sundaresan, S. and Z. Wang (2015): "On the design of contingent capital with a market trigger," *Journal of Finance*, 70, 881–920.

Sunder, S. (1992): "Market for information: Experimental evidence," *Econometrica*, 667–695.

——— (1995): "Experimental asset markets: A survey," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel and A. E. Roth, Princeton University Press.

Surico, P. (2007): "The Fed's monetary policy rule and US inflation: The case of asymmetric preferences," *Journal of Economic Dynamics and Control*, 31, 305–324.

Surowiecki, J. (2005): *The Wisdom of Crowds*, Random House.

Vives, X. (2010): *Information and learning in markets: The impact of market microstructure*, Princeton University Press.

Weise, C. L. (1999): "The asymmetric effects of monetary policy: A nonlinear vector autoregression approach," *Journal of Money, Credit and Banking*, 85–108.

Williamson, S. D. (1986): "Costly monitoring, financial intermediation, and equilibrium credit rationing," *Journal of Monetary Economics*, 18, 159–179.

Wolfers, J. and E. Zitzewitz (2004): "Prediction Markets," *Journal of Economic Perspectives*, 18, 107–126.

Wolff, E. N. (2002): *Top heavy: The increasing inequality of wealth in America and what can be done about it*, New Press.

Woodford, M. (1994): "Nonstandard Indicators for Monetary Policy: Can Their Usefulness Be Judged from Forecasting Regressions?" in *Monetary Policy*, ed. by N. G. Mankiw, University of Chicago Press.

Eidesstattliche Erklärung

Hiermit erkläre ich, die vorliegende Dissertation selbstständig angefertigt und mich keiner anderen als der in ihr angegebenen Hilfsmittel bedient zu haben. Insbesondere sind sämtliche Zitate aus anderen Quellen als solche gekennzeichnet und mit Quellenangaben versehen.

Mannheim, 8.12.2015

Curriculum Vitae, Christoph Siemroth

2010 – 2016	PhD Studies in Economics, Center for Doctoral Studies in Economics, University of Mannheim. Employee at the Chair for Economic Policy, Prof. Hans Peter Grüner, University of Mannheim.
09/2014 – 12/2014	Visiting PhD student, University College London, United Kingdom.
2010 – 2012	M.Sc. in Economics, University of Mannheim.
2006 – 2010	B.A. in Philosophy & Economics, University of Bayreuth.
01/2009 – 05/2009	Visiting student, University of Helsinki, Finland.
1998 – 2006	Abitur, Gymnasium Luisenstift, Radebeul.