Three Essays on Information in Financial Markets and the Macroeconomy

Ilja Kantorovitch

TESI DOCTORAL UPF / Year 2021

THESIS SUPERVISOR
Alberto Martin, Vladimir Asriyan, and Victoria Vanasco
Department of Business and Economics, UPF

Universitat Pompeu Fabra Barcelona

This dissertation is dedicated to my parents, Larissa and Ephim.

Thanks

Above all, I would like to extend my deepest gratitude to my advisors, Alberto Martin, Vladimir Asriyan, and Victoria Vanasco. Without their help and guidance, this dissertation would not have been possible. Their limitless patience, insightful comments, and willingness to challenge me resulted in this work. I was lucky to have such great mentors. I hope to pass on the favor in the future.

I met and talked to many other great researchers. I want to say thank you to Edouard Schaal, Isaac Bailey, Geert Mesters, Luca Fornaro, Jaume Ventura, Davide Debortoli, Giacomo Ponzetto, Christian Brownlees, Andrea Caggese, Filippo Ippolito, Jan Eeckhout, Björn Richter, Priit Jeenas, Xavier Vives, Barbara Rossi, and many others. It was a privilege that I could learn from such great scholars. Special thanks to Libertad González Luna, who did a great job at preparing my cohort for the job market.

I have been supported by the Studienstiftung des deutschen Volkes throughout my studies, which allowed me to focus on my work. I met many inspiring people and had many opportunities, which had a significant impact on my development. I am thankful for this privilege.

Any organization lives from people who keep everything together. I highly appreciate the help and support from Marta Araque and Laura Agusti. Thanks to them, I never had to worry about things not working out. The UPF and the Ph.D. program are very lucky to have you!

Many in the UPF Ph.D. community helped me substantially. A big thank you for Ángelo Gutiérrez-Daza for helping me out with any numerical problem and providing some perspective. I am also grateful to Inês Martins Xavier for walks, talks, and an open ear. I met many more great people. Among them are Christian Höynck, Derrick Kanngiesser, Milena Djourelova, Sebastian Ellingsen, Mario Giarda, Flavio Hafner, Karolis Liaudinskas, Christoph Albert, Sampreet Goraya, Andrea Fabiani, Adrian Lerche, Felix Mauersberger, Ana Costa-Ramón, Madalen Castells Jauregui, Bjarni Einarsson, Adam Lee, Evangelia Spantidaki, Paul Soto, Dijana Zejcirovic, and Josep Gisbert. A big thank you to Andrea Chiavari, Marta Morazzoni, Milan Quentel, and Sofia Trommlerová for providing a sense of community. Special thanks to Thomas Woiczyk, who organized numerous activities that helped to escape from focusing only on research. You all made this experience less lonely and more fun.

I want to thank especially Christopher Evans and Adilzhan Ismailov. We went through the masters and Ph.D. together and in the end, we made it. You were great flatmates! I also want to thank Pablo Lara, who left after the first year but remained part of the friend group.

I want to thank my office mates over the years, Eda Gulsen, Ricardo Maertens, Juan Felipe Imbet, Konstantin Werner, and Sébastien Willis and towards the end, Janko Heineken. Sébastien, I appreciate you for all our conversations about research and life, tea breaks, and the occasional political and historical discussions. Special thanks to my co-author Janko Heineken, who has been a great friend and colleague throughout the Ph.D.. I will miss our walks around the campus and our conversations.

Also, a big thank you to the members of the Ph.D. band White Noise. It was great to make music with Juan Felipe Imbet, Kinga Tchórzewska, Shohei Yamamoto, and Nikolas Schöll. My only regret is that I did not join earlier. I will keep our final concert in fond memory!

Outside of the UPF Ph.D. bubble, I want to thank my friends who stayed close to me despite the distance. Andre Weltsch, Paulina Piaszczyk, Stefanie Faas, Mark Fischer, Derya Arastiran, and Sergej Kaiser always had an open ear when I needed one. In Barcelona, Ertan Kuntmann helped me gain a simpler view of life and enjoy what is now. Christoph Spiegel helped me to switch off from research. Special thanks to Jonas Kröger for suggesting to apply at UPF, which I did not consider before. A big thank you to Andre Weltsch, Svenja Friess, and Christian Decker for proofreading this dissertation.

I feel deep gratitude towards my family. They believed in me and motivated me throughout this long journey with its ups and downs. It would not have been possible without my parents, Larissa and Ephim, who gave everything to support me. I also want to thank my sister Anna and my brother-in-law Oliver for always being there for me. I am very grateful for everyone who helped me get through the job market, especially Kinga Tchórzewska. Also, my uncle Lev and aunt Tatiana provided very much appreciated moral support during the job market.

Finally, I want to thank my partner, Lisa Maria Leusch, a thousand times. Lisa joined me on this journey, which was not always easy, and shared happy and sad moments with me. It is the best feeling to have you by my side, and I am excited about our next chapter in Switzerland and beyond!

Abstract

This dissertation consists of three chapters investigating the role of information and beliefs in financial markets and the macroeconomy. In the first chapter, I develop a macroeconomic model in which financial markets aggregate dispersed information and determine the efficiency of capital allocation. I find that fundamental (productivity) booms lower capital misallocation by encouraging information acquisition. In contrast, non-fundamental (sentiment) booms increase capital misallocation by discouraging information acquisition. In the second chapter, I introduce a tractable model of a noisy financial market in which information acquisition is motivated by overconfidence in the form of correlation neglect. I study several applications. Finally, in the third chapter, Janko Heineken and I study the role of sentiment and disagreement in determining the asset characteristics of Bitcoin. We show that disagreement predicts negative returns far into the future, with vanishing effects towards the end of the sample.

Resum

Aquesta dissertació consta de tres capítols que investiguen el paper de la informació i les creences en els mercats financers i la macroeconomia. Al primer capítol, desenvolupo un model macroeconòmic en què els mercats financers agrupen informació dispersa i determinen l'eficiència de l'assignació de capital. Trobo que els fonaments (productivitat) fonamentals redueixen la desassignació de capital fomentant l'adquisició d'informació. En canvi, els booms no fonamentals (sentiment) augmenten la desassignació de capital desincentivant l'adquisició d'informació. Al segon capítol, introdueixo un model tractable d'un mercat financer sorollós en què l'adquisició d'informació està motivada per l'excés de confiança en forma de desatenciót de la correlació. Estudio diverses aplicacions. Finalment, al tercer capítol, Janko Heineken i jo estudiem el paper del sentiment i el desacord a l'hora de determinar les característiques dels actius de Bitcoin. Mostrem que el desacord prediu rendiments negatius en el futur, amb efectes de desaparició cap al final de la mostra.

Preface

The idea that markets aggregate dispersed information is long-standing in the literature, going back to Hayek (1945) and constituting an active research field nowadays (Vives 2008). However, this feature is mostly absent in macroeconomic models. The main reason for this gap is that such models usually require a specific structure and include non-optimizing agents called noise traders, who keep prices from being fully revealing and add and remove resources from the economy. Both features are difficult to reconcile with conventional macroeconomic general equilibrium models.

This dissertation suggests an alternative approach to information aggregation in financial markets that is most closely related to Albagli, Hellwig, and Tsyvinski (2021) but uses overconfidence in the form of correlation neglect to incentivize costly information acquisition. This approach has multiple attractive features. First, overconfidence and correlation neglect are behavioral biases that have been documented for traders, financial managers, and experiment participants. In contrast, noise trading is an abstract catch-all for a variety of phenomema.[1] Second, it facilitates the application to macroeconomic general equilibrium models as boundedly-rational traders observe resource constraints, which is demonstrated in Chapter 1. Third, the approach yields a tractable model and allows studying questions that are difficult to embed in other frameworks, for example, trader heterogeneity and funding constraints. The latter point is shown in Chapter 2.

Empirical research on the acquisition and aggregation of information faces natural challenges because traders' information sets are rarely directly observable. The proliferation of online discussions and textual analysis techniques has opened the door to constructing sentiment measures, reflecting heterogeneous information or beliefs. Janko Heineken and I use this approach to test the predictions of the differences-in-opinion theory in the presence of short-sale constraints on Bitcoin in Chapter 3. We find that meaningful measures of sentiment and disagreement can be extracted using this approach, which are highly correlated with the returns, turnover, and volatility of Bitcoin.

[1]See Biais et al. (2005), Allen and Evans (2005), and Ben-David, Graham, and Harvey (2013) for evidence of overconfidence and Brandts, Giritligil, and Weber (2015), Eyster and Weizsäcker (2016), Eyster et al. (2018), Enke and Zimmermann (2019), Grimm and Mengel (2020), and Chandrasekhar, Larreguy, and Xandri (2020) for evidence of correlation neglect.

In Chapter 1, I take the idea of financial markets as aggregators of dispersed information to a macroeconomic model to study the effects of booms on capital misallocation. I find that booms driven by different forces also have different effects on misallocation. Fundamental booms, e.g., driven by productivity growth, lower misallocation by encouraging information production. In contrast, non-fundamental booms, e.g., driven by sentiments, increase misallocation by discouraging information production. I also show that the distinction between both types of booms and busts is also crucial for economic policy. For instance, asset purchases can increase economic activity and lower capital misallocation during non-fundamental busts but increase capital misallocation during fundamental busts. Finally, looking through the lens of the model, the US dot-com boom of the late 1990s appears to have been driven by productivity, whereas the US housing boom of the mid-2000s seems to have been driven by sentiment.

In Chapter 2, I develop a model in which overconfidence in the form of correlation neglect incentivizes costly information acquisition in financial markets. Traders' information has two sources of noise, one idiosyncratic and the other correlated between traders. Traders are overconfident in that they overestimate the share of idiosyncratic noise in their private information, i.e., they partly neglect correlated noise. I find that an infinitesimal amount of overconfidence is sufficient to generate trade when the private signal is exogenous and free. However, substantial amounts of overconfidence are needed when traders acquire costly information. I show that the model can be integrated into macroeconomic models as in Chapter 1 and can be used to study trader heterogeneity. Finally, I consider an extension in which traders have limited resources for trading. Such funding constraints dampen the effect of new information on the price. Moreover, disagreement can affect the price level differently depending on the relative scarcity or abundance of trading capital.

In Chapter 3, Janko Heineken and I test the theoretical predictions of the differences-of-opinion literature in the case of Bitcoin, for which beliefs and disagreement are central. We analyze the extensive online discussion on Bitcoin to build a time-varying sentiment distribution, defining disagreement as the dispersion in the sentiment distribution. We confirm the theory's predictions as disagreement is associated with negative returns, high turnover growth, and volatility. Moreover, we find that disagreement predicts lower returns far into the future. However, this predictive effect vanishes towards the end of our sample when shorting instruments were introduced.

Contents

List of Figures

List of Tables

Chapter 1

EXUBERANT AND UNINFORMED: HOW FINANCIAL MARKETS (MIS-)ALLOCATE CAPITAL DURING BOOMS

1.1 Introduction

Financial markets play a central role in allocating capital to its most productive uses. Yet, they do not always fulfill this role well. The last three decades, for instance, have been characterized by successive booms and busts in financial markets.[1] These cycles have been difficult to justify on fundamental grounds alone (Martin and Ventura 2018). Against this backdrop, there are growing concerns that such booms lead to the deterioration of capital allocation, ultimately reducing aggregate productivity.[2] The general

[1]For example, the dot-com bubble in the US and the housing bubbles in the US and Southern Europe.

[2]For instance, Gopinath et al. (2017) and García-Santana et al. (2020) have found that the credit and asset price boom in Southern Europe, preceding the global financial crisis, had coincided with a rise in capital misallocation; Doerr (2018) provides such evidence for the US. Relatedly, Borio et al. (2015)

narrative is as follows: during booms, the perception is that all investments perform well. As a result, agents are less prone to produce information about specific investments, and markets eventually become less informative, thereby worsening the allocation of resources in the economy.[3] Though suggestive, this narrative is loose and cannot be fully evaluated without a theory of information production and its macroeconomic effects. The goal of this paper is to provide such a theory.

In this paper, I develop a tractable macroeconomic framework in which financial markets play the key role of aggregating information that is dispersed among economic agents. The framework's central feature is that information is endogenous, in a sense that agents can decide to engage in costly information production. The framework's novelty is to study the two-way feedback between macroeconomic conditions and agents' incentives to produce information.

I model a dynamic economy populated by firms with heterogeneous productivity and households, which consist of many traders. Households decide on borrowing and saving. Traders decide which firms to invest in, but they have imperfect information about firm productivity. To make their investment decisions, traders combine their private information with a public signal provided by financial markets, which effectively aggregates all traders' information.

The model is based on two core assumptions. First, traders agree on realizations of aggregate shocks but disagree about the distribution of firm productivity. Whereas the former part is for simplicity, the latter is central for motivating trade. In particular, traders' private information features both idiosyncratic and correlated noise. The idiosyncratic noise captures trader-specific information and drives disagreement. In contrast, the correlated noise stands for a common "sentiment" across traders.[4] Second, to incentivize information production in equilibrium by avoiding the well-known Grossman-Stiglitz paradox (see Grossman and Stiglitz 1980),[5] traders are assumed to be

show that credit booms tend to also coincide with misallocation of labor.

[3]There are several studies that point to a decline in information production and quality in explaining a worsening of investment efficiency, for example, Asea and Blomberg (1998), Keys et al. (2010), and Becker, Bos, and Roszbach (2020).

[4]From an economic standpoint, this sentiment is meant to capture a range of phenomena that drive asset prices away from their fundamental value, such as herding, network effects, social learning, extrapolative expectations, or bubbles (see Kindleberger and Aliber 2015; Shiller 2015, 2017).

[5]The Grossman-Stiglitz paradox states that no equilibrium exists in models of financial markets with costly information production without noise that keeps prices from being perfectly revealing. If prices

overconfident.[6] Formally, each trader believes the noise in her private information to be entirely idiosyncratic, allowing her to exploit mispricing due to sentiment. In a nutshell, each trader believes that she is not prone to sentiments even though she understands that everyone else is.

I find that information production crucially depends on the state of the economy. In particular, I study how information production reacts to two types of macroeconomic shocks: sentiment and productivity. Sentiment shocks, defined as waves of optimism or pessimism, formally drive the correlated noise in traders' private information. Sentiment and productivity shocks lead to similar co-movements in output, investment, and asset prices. However, they affect information production differently. Information is central in my model, as more precise information strengthens the correlation between the size of a firm and its productivity, thereby raising allocative efficiency. Consequently, an economy with higher information production allocates more capital to more productive firms and has higher aggregate productivity.

In particular, information production increases in productivity but is non-monotonic in sentiment. Productivity increases information production due to a scale effect. Since high productivity raises the optimal size of a firm, it also boosts the benefits of producing more precise information about it. From the viewpoint of an individual trader, producing more information is valuable if it significantly impacts the trader's investment decisions. However, if sentiment regarding a specific firm is too high or too low, producing more information is likely to not yield much. In particular, even without precise information, a trader knows not to invest in firms where sentiment is high (i.e., firms that are "overvalued") and to invest in firms where sentiment is low (i.e., firms that are "undervalued"). Thus, extreme sentiments discourage the production of information.

Finally, while productivity booms are endogenously amplified by information production's effect on capital allocation, sentiment booms may be dampened. Productivity booms *crowd in* information and improve allocative efficiency, thereby further increas-

reveal all information, traders have no reason to produce costly information. However, prices cannot be informative if traders do not produce information. Therefore, no equilibrium exists. Many models of informative financial markets (Grossman and Stiglitz 1980; Kyle 1985; Albagli, Hellwig, and Tsyvinski 2021) circumvent this problem by introducing so-called noise traders. These agents are non-optimizing, which makes them difficult to embed in a general equilibrium model.

[6]Some form of noise in asset prices is indispensable to motivate trade and information production in financial markets. I formalize and micro-found such noise by introducing correlated noise in traders' signals and assuming that traders are overconfident.

ing productivity. Sentiment booms, however, *crowd out* information and worsen allocative efficiency, thereby decreasing productivity. My finding is consistent with the empirical evidence that booms can fuel resource misallocation (e.g., Gopinath et al. 2017; Doerr 2018), suggesting that such booms are driven by sentiment. It also captures a dichotomy of booms as in Gorton and Ordoñez (2020) but stresses source of booms is the essential factor.

On the normative front, information production is too high or too low in the competitive equilibrium due to the presence of two externalities. First, there can be too much information because traders produce information to gain at the expense of other traders (*rent-extracting behavior*). Second, there can also be too little information because traders do not reap the benefits of improved capital allocation through collective information production (*information spillover*). Which effect dominates depends on whether the allocation of capital is important for aggregate productivity. For example, if firms produce similar goods, allocating capital to the most productive firms becomes exceedingly important. Yet, this is exactly when the competitive equilibrium features little information production.

Moreover, my model sheds light on two current policy debates. First, it suggests that policymakers should tax investment during sentiment-driven booms, which can be identified by increasingly synchronous asset price movements. This policy prescription of "leaning against the wind" is often criticized on informational grounds:[7] namely, it requires the policymaker to be able to distinguish sentiment- from productivity-driven booms in real-time (e.g., Mishkin et al. 2011). My model suggests that, although they look similar in many respects, both types of booms can be distinguished through their effects on information production. In particular, less informative asset prices display more synchronous movements, which can identify sentiment booms. In contrast, productivity booms lead to more asynchronous asset price movements.

A second policy debate refers to the effects of large-scale asset purchases by central banks. There is the widespread perception that asset purchases can distort prices and worsen the allocation of resources.[8] My model yields a simple yet robust insight:

[7] See Cecchetti et al. (2000).

[8] See da Silva and Rungcharoenkitkul (2017), DNB (2017), Fernandez, Bortz, and Zeolla (2018), Borio and Zabai (2018), Acharya et al. (2019), and Kurtzman and Zeke (2020). In particular, the Dutch central bank argues in their 2016 annual report (DNB 2017): "*The large-scale purchase programmes and the flood of liquid assets has set the risk compass in financial markets spinning, with misallocations as a*

whether this concern is justified depends on whether asset purchases reduce or aggravate the aggregate mispricing of assets. By reducing the asset supply in the hands of traders, asset purchases change the marginal trader's identity and thus raise equilibrium prices. If asset prices were initially depressed due to low productivity, the common perception laid out above is correct. By distorting prices upward, asset purchases discourage information production and thus worsen the allocation of capital. However, if asset prices were initially depressed due to negative sentiment, asset purchases *reduce* aggregate mispricing. Indeed, by undoing the effects of negative sentiment, asset purchases fuel information production, thereby improving the resource allocation.

Finally, the paper makes a methodological contribution by providing a tractable macroeconomic model of information production and aggregation, where financial market informativeness plays an important role for macroeconomic dynamics. With a few exceptions,[9] the role of financial markets as aggregators of dispersed information has received little attention in macroeconomics.[10] The primary reason is that most standard models of informative financial markets rely on non-optimizing agents, such as noise traders, which are not straightforward to reconcile with general-equilibrium analysis. Instead, my model relies on a small behavioral deviation – overconfidence – which means that traders do not adequately perceive the idiosyncratic and correlated components in their signals. This misperception motivates them to produce costly information as they believe in having an informational edge over the market. This simple assumption is grounded on empirical evidence,[11] and it avoids the Grossman-Stiglitz paradox.

1.1.1 Literature Review

A recent literature studies the link between information production and the business cycle (Veldkamp 2005; Ambrocio 2020; Farboodi and Kondor 2020; Chousakos, Gorton, and Ordoñez 2020; Gorton and Ordoñez 2020; Asriyan, Laeven, and Martin forthcom-

result."

[9]Some exceptions are Peress (2014), David, Hopenhayn, and Venkateswaran (2016), and Straub and Ulbricht (2018).

[10]The idea of markets as aggregators of dispersed information dates back to Hayek (1945): *"The economic problem of society is (...) rather a problem of how to secure the best use of resources known to any of the members of society, for ends whose relative importance only these individuals know. Or, to put it briefly, it is a problem of the utilization of knowledge which is not given to anyone in its totality.*"

[11]See for example Eyster et al. (2018), Grimm and Mengel (2020), and Enke and Zimmermann (2019).

ing).[12] In contrast to Gorton and Ordoñez (2020), my model shows that the source of fluctuations is important for the relationship between the cycle and information production.

This paper builds on the literature on informative financial markets (Grossman and Stiglitz 1980; Kyle 1985; Vives 2008; Albagli, Hellwig, and Tsyvinski 2021). In this literature, limits to arbitrage keep arbitrageurs from fully eliminating mispricing and, therefore, incentives to trade and produce information persist in equilibrium. My model's market microstructure is similar to Albagli, Hellwig, and Tsyvinski (2021). Whereas Albagli, Hellwig, and Tsyvinski (2021) use noise traders to keep prices from being fully revealing, I make the model more tractable by assuming instead that traders are overconfident. As I show, this innovation allows me to embed a noisy financial market into an otherwise standard macroeconomic model.

A strand of the literature uses the insight that prices can be informative to study the role of this information in economic decisions, as surveyed in Bond, Edmans, and Goldstein (2012). For example, secondary markets can be sources of information for managers (Holmström and Tirole 1993; Dow and Gorton 1997). Information is important in my model as a measure of allocative efficiency without any firms actively learning from prices. Similar to Dow, Goldstein, and Guembel (2017), I study the two-way feedback between the financial and real economy when traders produce information endogenously. A number of papers has brought this paradigm to macroeconomics (Peress 2014; David, Hopenhayn, and Venkateswaran 2016; Albagli, Hellwig, and Tsyvinski 2017; Straub and Ulbricht 2018; Asriyan 2021). My contribution is to study the effects of aggregate shocks on information production and the allocation of capital. From a normative perspective, I show under which conditions information production is likely to be too high or too low in the competitive equilibrium.

There is ample empirical evidence that asset prices are indeed informative. See Morck, Yeung, and Yu (2013) for a survey on the literature that uses "non-synchronicity" as a measure of price-informativeness.[13] Morck, Yeung, and Yu (2000) found that

[12] See also Van Nieuwerburgh and Veldkamp (2006), Angeletos, Lorenzoni, and Pavan (2010), Ordoñez (2013), Gorton and Ordoñez (2014), Fajgelbaum, Schaal, and Taschereau-Dumouchel (2017), and Straub and Ulbricht (2018) for related work.

[13] Non-synchronicity has been suggested by Roll (1988) as a measure of firm-specific information in asset prices. The main idea is that as the volatility of asset prices increasingly relates to firm-specific factors, prices also become increasingly informative about firms.

more developed countries have stock markets that are more informative. Focusing instead on the cross-section of firms, Durnev, Morck, and Yeung (2004) found that non-synchronicity is positively related to the efficiency of corporate investment. More recently, Bai, Philippon, and Savov (2016) and Farboodi et al. (2020) have shown that prices have become better predictors of corporate earnings in the US since the 1960s. The latter emphasize that this has been mainly the case for large growth firms. Finally, Bennett, Stulz, and Wang (2020) provide evidence that price informativeness increases firm productivity. Price informativeness is closely related to the allocative efficiency of financial markets in my model.

The results of my model are broadly consistent with empirical evidence on how price informativeness varies over the business cycle. Dávila and Parlatore (2021) proposed an identification procedure to estimate price informativeness from price and earnings data, which is closely related to information production in my model.[14] Comparing fluctuations around the corresponding trends for the US reveals a highly positive correlation between price informativeness and TFP growth, as can be seen in Figure 1.1. From 1995 to 2001, price informativeness and TFP growth were increasing, pointing to a productivity-driven expansion. In contrast, the housing boom from 2002 to 2008 eventually even led to a decline in TFP and a steep fall in price informativeness relative to trend, which indicates a sentiment-driven boom during these years. This interpretation is in line with Borio et al. (2015), who suggested that TFP growth slowed between 2002 and 2008 *because of the financial boom*, not despite it.

In my model, traders suffer from correlation neglect. This bias has been studied in the literature and documented repeatedly in experimental settings (Brandts, Giritligil, and Weber 2015; Eyster et al. 2018; Enke and Zimmermann 2019; Grimm and Mengel 2020; Chandrasekhar, Larreguy, and Xandri 2020). When receiving information from multiple sources, neglecting correlated noise in the signals can lead to an overly precise posterior. Therefore, correlation neglect leads to overconfidence, which plays a central role in the literature on behavioral biases, especially in relation to financial markets (Glaser and Weber 2010; Daniel and Hirshleifer 2015).

Finally, a broad literature studies the role of sentiments in macroeconomics (for a

[14] Intuitively, relative price informativeness is the weight an otherwise uninformed observer puts on the information embodied in the price relative to her prior. In my model, the weight only varies due to changes in the information production by traders.

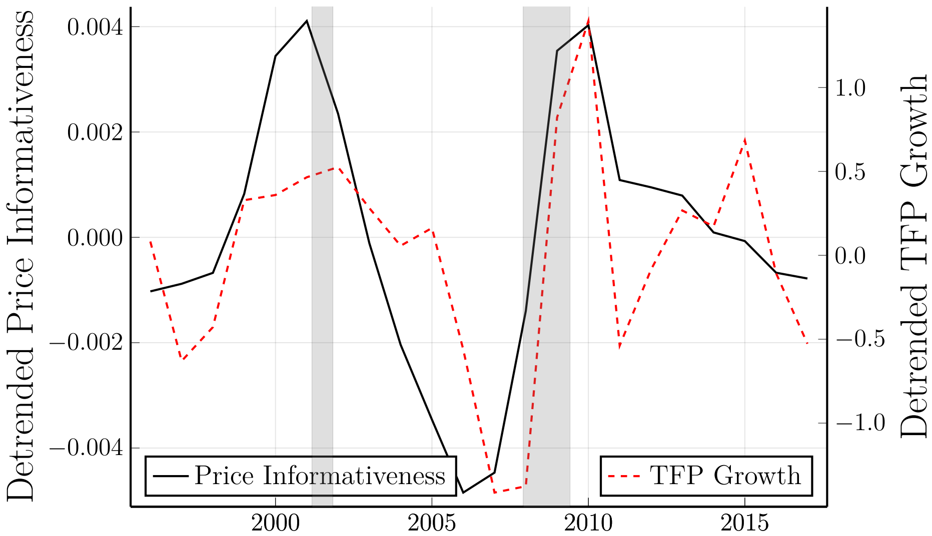

Figure 1.1: Detrended Price Informativeness and TFP Growth.

Notes: Price informativeness (black) as measured in Dávila and Parlatore (2021) and utilization-adjusted TFP growth (red) taken from the Federal Reserve of San Francisco following Basu, Fernald, and Kimball (2006). Grey bars indicate recessions following the NBER dating methodology. The time series have been detrended using a cubic time trend and smoothed with a two-year moving average. Through the lens of the model, productivity drove the expansion until 2001, as indicated by a rise in information and TFP growth. In contrast, sentiment drove the expansion from 2002 to 2008 as indicated by the decline in information and TFP growth.

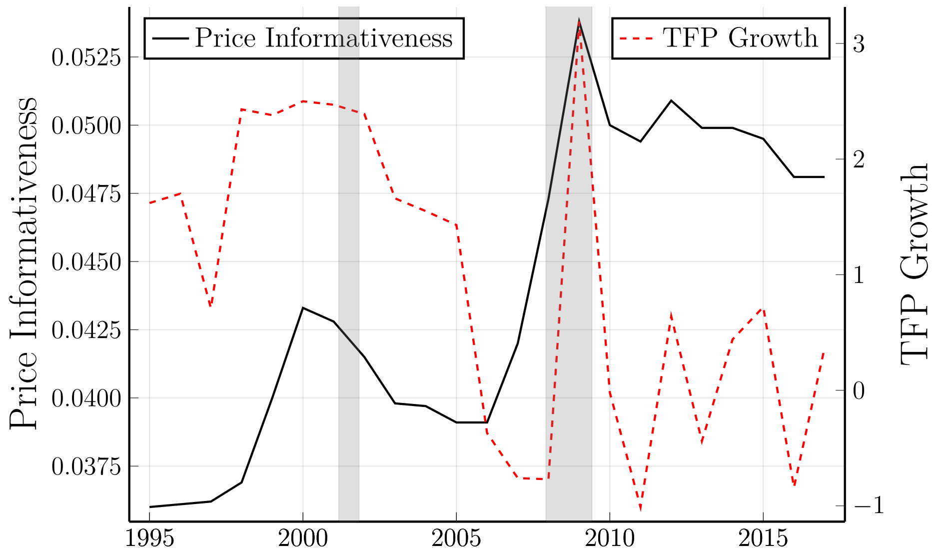

Figure 1.2: Price Informativeness and TFP Growth.

Notes: Price informativeness (black) as measured in Dávila and Parlatore (2021) and utilization-adjusted TFP growth (red) following Basu, Fernald, and Kimball (2006) taken from the Federal Reserve of San Francisco. Grey bars indicate recessions following the NBER dating methodology. Raw time series.

survey, see Nowzohour and Stracca 2020). There are different definitions of sentiments, ranging from self-fulfilling beliefs (Martin and Ventura 2018; Asriyan, Fuchs, and Green 2019) to news and noise shocks (Angeletos, Lorenzoni, and Pavan 2010; Schmitt-Grohé and Uribe 2012). In my model, sentiments are waves of non-fundamental optimism or pessimism. When a positive sentiment shock hits, agents become optimistic about productivity and vice versa.

1.2 Model

1.2.1 Households and Traders

The model is populated by overlapping generations of households indexed by $i \in [0, 1]$. As is common in the New Keynesian literature, I assume that each household i consists of a unit mass of traders indexed by $ij \in [0, 1] \times [0, 1]$ (for example, see Blanchard and Galí 2010). Households pool resources, borrow on behalf of traders, and distribute consumption equally, whereas traders individually maximize the utility for the household given by

$$U_{it} = C_{it,t} + \delta \mathbb{E} \left\{ C_{it,t+1} \right\} - \int_0^1 IA(\beta_{ijt}) dj, \tag{1.1}$$

where $C_{it,t}$ is youth consumption, $C_{it,t+1}$ is old age consumption, $\delta \in (0, 1)$ is the discount factor, and $\int_0^1 IA(\beta_{ijt})dj$ are information production costs, which are introduced in more detail in a later section.

When young, traders each supply one unit of labor inelastically, receive wage W_t and buy shares of intermediate good firms in a competitive financial market. To avoid unbounded demands by risk-neutral traders, demand for each stock is limited to the interval $[\kappa_L, \kappa_H]$ where $\kappa_L \leq 0$ and $\kappa_H > 1$.[15] Traders also choose the precision β_{ijt} of a noisy signal of firm productivity to inform their trading decision subject to a utility cost $IA\left(\beta_{ijt}\right)$. Finally, the household lends and borrows through risk-free bonds with return R_{t+1}.

[15] See Dow, Goldstein, and Guembel (e.g., 2017) and Albagli, Hellwig, and Tsyvinski (2021) for similar approaches and Appendix 1.B for a further elaboration.

1.2.2 Technologies

Final Good Sector

There are many identical final good firms owned by households. The production function for the final good, which also serves as the numéraire, is Cobb-Douglas over labor and a CES-aggregate of intermediate goods. Aggregate output is

$$Y_t = L^{1-\alpha} \left(\int_0^1 Y_{jt}^{\frac{\theta-1}{\theta}} dj \right)^{\frac{\alpha\theta}{\theta-1}}, \tag{1.2}$$

where $\theta \in (0, \infty)$ is the elasticity of substitution between varieties and α is the share of intermediate goods. Y_{jt} is an intermediate good produced by firm j. The final good can be consumed or invested in firm capital. L the labor supply and normalized to one.

Intermediate Good Sector

For each generation, there is a unit mass of intermediate good firms $j \in [0, 1]$ with production function

$$Y_{jt} = A_{jt-1}^{\frac{\theta}{\theta-1}} K_{jt}, \tag{1.3}$$

where K_{jt} is firm capital and $\ln(A_{jt-1}) \overset{iid}{\sim} \mathcal{N}(a_{t-1}, \sigma_a^2)$ is firm productivity. Note that time subscript $t-1$ is used as agents learn about firm productivity in the period prior to production. Capital takes time to build, such that investment takes place in t but production in $t+1$, and depreciates fully after production. Each firm sells a unit mass of claims to total firm-revenue to households and finances capital investment with the proceeds:[16]

$$P_{jt} = K_{jt+1}. \tag{1.4}$$

Information Structure

Trader ij is only active in the market for shares of firm j, for which she is an expert as she receives the signal

$$s_{ijt} = a_{jt} + \frac{\eta_{ijt} + \varepsilon_{jt}}{\sqrt{\beta_{ijt}}}, \tag{1.5}$$

[16] See Appendix 1.C for a micro-foundation and further discussion.

where $a_{jt} \overset{iid}{\sim} \mathcal{N}(a_t, \sigma_a^2)$ is firm productivity, $\eta_{ijt} \overset{iid}{\sim} \mathcal{N}(0,1)$ is idiosyncratic noise, $\varepsilon_{jt} \overset{iid}{\sim} \mathcal{N}(\varepsilon_t, \sigma_\varepsilon^2)$ is correlated noise, interpreted as *sentiment*, and β_{ijt} is a information precision parameter chosen by trader ij.[17] Both idiosyncratic and correlated noise are iid over time and across markets; idiosyncratic noise is also iid between traders. A high realization of η_{ijt} means that trader ij is optimistic about firm j relative to other traders in the same market. Similarly, a high realization of ε_{jt} means that all traders in market j are too optimistic.

Assumption 1.1 (***Overconfidence***)**.** *Trader ij believes the information structure to be*

$$
\begin{aligned}
s_{ijt} &= a_{jt} + \frac{\eta_{ijt}}{\sqrt{\beta_{ijt}}} \\
s_{-ijt} &= a_{jt} + \frac{\eta_{-ijt} + \varepsilon_{jt}}{\sqrt{\beta_{-ijt}}}.
\end{aligned}
$$

Following Assumption 1.1, traders believe that sentiment ε_{jt} drives the beliefs of all traders but not their own beliefs. As a result, traders are overconfident and willing to produce costly information to exploit mispricing induced through sentiment shocks ε_{jt}.[18] Finally, trader ij chooses the precision of her private signal β_{ijt} subject to a convex cost function $IA(\beta_{ijt})$ with standard properties $IA(0) = 0$, $IA'(0) = 0$, $IA''(\cdot) > 0$.

Aggregate Shocks

Two classes of shocks drive the economy. *Aggregate productivity shocks* move the mean of the distribution of firm-specific productivity shocks, $a_{jt} \sim \mathcal{N}(a_t, \sigma_a^2)$, and *aggregate sentiment shocks* drive the mean of firm-specific sentiment shocks, $\varepsilon_{jt} \sim \mathcal{N}(\varepsilon_t, \sigma_\varepsilon^2)$, similar to Angeletos, Lorenzoni, and Pavan (2010). The sentiment shock ε_t is meant to capture a range of phenomena that lead to non-fundamental price movements in financial markets, e.g., herding, informational cascades, social learning, bubbles, liquidity trading (see Kindleberger and Aliber 2015; Shiller 2015, 2017). I study economy-wide sentiment shocks as they affect cross-sectional misallocation of capital only through

[17] See section 1.7.2 for the effect of uncertainty about aggregate shocks.

[18] This assumption is necessary to avoid the Grossman-Stiglitz paradox (Grossman and Stiglitz 1980). It states that informationally efficient markets are impossible in the absence of noise when information is costly. In that case, markets would already reveal all information and, therefore, destroy the incentive to produce costly information in the first place.

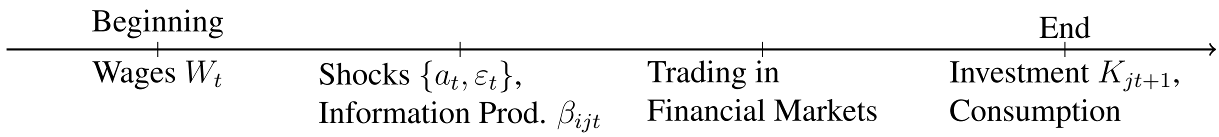

Figure 1.3: Intraperiod Timing.

their effect on information production.[19]

For simplicity, traders perfectly observe aggregate shocks $\{a_t, \varepsilon_t\}$ before their information production decision, but firm-specific shocks $\{a_{jt}, \varepsilon_{jt}\}$ need to be learned. The laws of motion for the aggregate shocks are irrelevant for this setup, as the dynamic model is a repetition of static problems. It follows that the information set of trader ij consists of the private signal s_{ijt}, share prices $\{P_{jt}\}$ for all markets $j \in [0, 1]$, and the mean and variances of firm-specific shocks $\{a_t, \varepsilon_t\}$, i.e., $\mathcal{I}_{ij} = \{s_{ijt}, \{P_{jt}\}, a_t, \varepsilon_t\}$. In other words, traders have rational beliefs about aggregates, but disagree about the productivity of intermediate firms based on public information in the forms of prices and private signals.

1.2.3 Timing

The timing is laid out in Figure 1.3. At the beginning of each period, young traders work in the final good sector and receive wage W_t. Then, traders choose the precision of their signal and the financial market opens. At the end of the period, both investment and consumption take place.

1.2.4 Notation

Traders think that their private signals do not contain correlated noise ε_{jt} as in Assumption 1.1. Therefore, expectations that condition on private signals are distorted and denoted by $\tilde{\mathbb{E}}(\cdot)$.

[19] Sector-specific sentiment shocks lead directly to an increase in capital misallocation, as aggregate output could be increased by reallocating capital away from the shocked sector. In this case, the results still go through on the sector level, as a sector-specific shock leads to an increase of capital misallocation inside the shocked sector. A more detailed analysis can be found in Appendix 1.D.

The determinants of functions are usually omitted to save on notation. For example, firm j's revenue is denoted by Π_{jt+1} instead of $\Pi\left(A_{jt}, K_{jt+1}, Y_{t+1}\right)$. Moreover, A_{jt} is indexed by t instead of $t+1$, as traders can learn about firm productivity in period t.

1.2.5 The Household's Problem and The Trader's Problem

Household i takes interest rate R_{t+1} as given and decides how much to borrow or lend. Furthermore, households are also prone to the behavioral bias of Assumption 1.1, in the that they each household believes that all its traders indeed have signals that are free of sentiment. However, households do not observe the private signals of traders. The household's problem is

$$\max_{B_{it+1}} \quad C_{it,t} + \delta\tilde{\mathbb{E}}_t\left\{C_{it,t+1}\right\} - \int_0^1 IA\left(\beta_{ijt}\right)dj \tag{P1.1}$$

$$s.t. \quad C_{it,t} = W_t - \int_0^1 x_{ijt}P_{jt}dj - B_{it+1} \tag{1.6}$$

$$C_{it,t+1} = \int_0^1 x_{ijt}\Pi_{jt+1}dj + R_{t+1}B_{it+1} \tag{1.7}$$

$$C_{it,t}, C_{it,t+1} \geq 0. \tag{1.8}$$

Households optimally choose how to much lend or borrow subject to the budget constraints. The first constraint (1.6) states that consumption during youth is equal to wages W_t minus the costs of buying stocks $\int_0^1 x_{ijt}P_{jt}dj$ and saving through the bond market B_{it+1}. Constraint (1.7) states that old age consumption is equal to revenue $\int_0^1 x_{ijt}\Pi_{jt+1}dj$ plus income from lending on the bond market $R_{t+1}B_{it+1}$. Although household i is overly optimistic about the return of its portfolio due to overconfidence, each household correctly values the portfolio of all other households. Therefore, limiting borrowing by the natural borrowing constraint as in (1.8) rules out defaulting on any borrowing through bonds.

Household i's optimal saving decision is given by

$$B_{it+1}\begin{cases} = -\frac{\int_0^1 x_{ijt}\Pi_{jt+1}dj}{R_{t+1}} & \text{if } R_{t+1} < \frac{1}{\delta} \\ \in \left[-\frac{\int_0^1 x_{ijt}\Pi_{jt+1}dj}{R_{t+1}}, W_t - \int_0^1 x_{ijt}P_{jt}dj\right] - & \text{if } R_{t+1} = \frac{1}{\delta} \\ = W_t - \int_0^1 x_{ijt}P_{jt}dj & \text{if } R_{t+1} > \frac{1}{\delta} \end{cases}. \tag{1.9}$$

If the interest rate R_{t+1} is below $\frac{1}{\delta}$, it is optimal to borrow as much as possible. If the interest is equal to $\frac{1}{\delta}$, household i is indifferent between borrowing and saving. Finally, if the interest rate is above $\frac{1}{\delta}$, then it is optimal to save as much as possible. Plugging (1.6) and (1.7) into (P1.1) and using the solution for the saving decision (1.9) yields trader ij's problem

$$\max_{\beta_{ijt}} \quad \tilde{\mathbb{E}}_t \left\{ \lambda_t \max_{x_{ijt}} \tilde{\mathbb{E}} \left\{ x_{ijt} \left(\frac{1}{R_{t+1}} \Pi_{jt+1} - P_{jt} \right) | s_{ijt}, P_{jt} \right\} \right\} - IA\left(\beta_{ijt}\right) \quad \text{(P1.2)}$$

$$s.t. \quad x_{ijt} \in [\kappa_L, \kappa_H] \tag{1.10}$$

$$\beta_{ijt} \geq 0, \tag{1.11}$$

where $\lambda_t = \max\{1, \delta R_{t+1}\}$ and terms that do not depend on the decision by trader ij were dropped. The problem is split into two parts, which are solved in reverse chronological order. Given information production β_{ijt} and realizations of the private signal s_{ijt} and price P_{jt}, trader ij chooses demand x_{ijt} for share j subject to the position limits (1.10). Using the solution to the trading problem, trader ij decides on the information precision β_{ijt} to increase the likelihood of trading profitably subject to a non-negativity constraint. Trader ij can use the household i's pooled resources and borrow through the household for trading. The term λ_t reflects that the value of an additional unit of wealth during youth may be above one.

1.3 Equilibrium Characterization

1.3.1 Input Markets

Wages and intermediate good prices are determined competitively,

$$W_t = \frac{\partial Y_t}{\partial L} = (1-\alpha) Y_t \tag{1.12}$$

$$\rho_{jt} = \frac{\partial Y_t}{\partial Y_{jt}} = \alpha Y_t^{\alpha_Y} Y_{jt}^{-\frac{1}{\theta}}, \tag{1.13}$$

where $\alpha_Y = \frac{\alpha\theta - \theta + 1}{\alpha\theta}$. Wages are equal to a share $(1-\alpha)$ of output. The price for intermediate good j is downward sloping in the quantity produced of the same good.

Finally, the revenue of intermediate good firm j is given by

$$\Pi_{jt+1} = \rho_{jt+1} Y_{jt+1}. \tag{1.14}$$

1.3.2 Trader's Decisions

Trading If price P_{jt} exceeds expectations of revenue Π_{jt+1} using the interest rate on bonds R_{t+1} as the benchmark rate, trader ij sells $-\kappa_L$ shares; when these values coincide trader ij is indifferent between buying and selling. When expectations exceed the price, trader ij buys κ_H shares:

$$x\left(s_{ijt}, P_{jt}\right) = \begin{cases} \kappa_L & \text{if } \frac{1}{R_{t+1}} \tilde{\mathbb{E}} \left\{ \Pi_{jt+1} | s_{ijt}, P_{jt} \right\} < P_{jt} \\ \in \left[\kappa_L, \kappa_H\right] & \text{if } \frac{1}{R_{t+1}} \tilde{\mathbb{E}} \left\{ \Pi_{jt+1} | s_{ijt}, P_{jt} \right\} = P_{jt} \\ \kappa_H & \text{if } \frac{1}{R_{t+1}} \tilde{\mathbb{E}} \left\{ \Pi_{jt+1} | s_{ijt}, P_{jt} \right\} > P_{jt} \end{cases}. \tag{1.15}$$

Information Production As laid out in (1.15), the trading decision is driven by the realization of the private signal s_{ijt} relative to price P_{jt}. Consequently, trader ij chooses information precision β_{ijt} to improve her ability to identify profitable trading opportunities. A central object in this context is the subjective probability of buying conditional on realizations of productivity a_{jt} and sentiment ε_{jt}, trader ij's information choice β_{ijt}, and the symmetric choice of all other traders in the market β_{jt}. Taking expectations with respect to the realizations of idiosyncratic noise, η_{ijt}, yields the probability of buying,

$$\mathcal{P}\left\{x_{ijt} = \kappa_H | a_{jt}, \varepsilon_{jt}, \beta_{ijt}, \beta_{jt}\right\} = \int_{-\infty}^{\infty} \phi\left(\eta_{ijt}\right) 1_{\frac{1}{R_{t+1}} \tilde{\mathbb{E}}\left\{\Pi_{jt+1} | s_{ijt}, P_{jt}\right\} > P_{jt}} d\eta_{ijt}, \tag{1.16}$$

where $\phi(\cdot)$ is the standard-normal pdf.[20]

The first-order condition for the information production decision is obtained after plugging (1.15) into (P1.2). Evaluating the expectations with respect to the realizations of the idiosyncratic noise η_{ijt} and taking the symmetric information production decisions of all other traders as given ($\beta_{-ijt} = \beta_{jt}$), leads to the first-order condition:

[20] A more detailed derivation can be found in Appendix 1.A.

$$
\widetilde{MB}\left(\beta_{ijt},\beta_{jt}\right) = \lambda_t \tilde{\mathbb{E}}_t \left\{ \left(\kappa_H - \kappa_L\right) \underbrace{\frac{\partial \mathcal{P}\left\{x_{ijt} = \kappa_H | a_{jt}, \varepsilon_{jt}, \beta_{ijt}, \beta_{jt}\right\}}{\partial \beta_{ijt}}}_{\text{Change in the Probability of Buying}} \underbrace{\left(\frac{1}{R_{t+1}} \Pi_{jt+1} - P_{jt}\right)}_{\text{Rents}} \right\}
$$

$$
= IA'\left(\beta_{ijt}\right). \tag{1.17}
$$

The marginal benefit of producing information consists of two parts. First, the probability of buying in state $(a_{jt}, \varepsilon_{jt})$ given information choices $(\beta_{ijt}, \beta_{jt})$. Second, trading rents given by the difference between the net present value of firm revenue minus the price of the stock.

1.3.3 Financial Market

Market-Clearing At the symmetric equilibrium ($\forall j : \beta_{ijt} = \beta_{jt}$), traders buy κ_H shares whenever their private signals are above some threshold, $\hat{s}\left(P_{jt}\right)$, are indifferent between buying and selling when their private signals coincide with the threshold, and sell otherwise. After normalizing the supply of shares in each market j to one, the market-clearing condition becomes

$$
\kappa_H \left(1 - \Phi\left(\sqrt{\beta_{jt}}\left(\hat{s}\left(P_{jt}\right) - a_{jt}\right) - \varepsilon_{jt}\right)\right) - \kappa_L \Phi\left(\sqrt{\beta_{jt}}\left(\hat{s}\left(P_{jt}\right) - a_{jt}\right) - \varepsilon_{jt}\right) = 1, \tag{1.18}
$$

where $\Phi(\cdot)$ is the standard normal cdf. The threshold $\hat{s}\left(P_{jt}\right)$ can be solved for directly,

$$
\hat{s}\left(P_{jt}\right) = a_{jt} + \frac{\varepsilon_{jt} + \Phi^{-1}\left(\frac{\kappa_H - 1}{\kappa_H + \kappa_L}\right)}{\sqrt{\beta_{jt}}}. \tag{1.19}
$$

Price Signal Traders learn from prices, which is equivalent to observing a noisy signal of the form

$$
z_{jt} = \hat{s}\left(P_{jt}\right) - \frac{\Phi^{-1}\left(\frac{\kappa_H - 1}{\kappa_H + \kappa_L}\right)}{\sqrt{\beta_{jt}}} = a_{jt} + \frac{\varepsilon_{jt}}{\sqrt{\beta_{jt}}}. \tag{1.20}
$$

When the price P_{jt} is high, traders realize that this can be due to two reasons: either firm j is productive (high a_{jt}) or other traders are very optimistic (high ε_{jt}). Therefore, prices are a noisy signal of firm productivity. The combination of dispersed information and position limits for asset demand ensure that the signal is normally distributed as $z_{jt} \sim \mathcal{N}(a_{jt}, \sigma_\varepsilon^2/\beta_{ijt})$ for all values of κ_L and κ_H. I call z_{jt} the *price signal* and expectations condition on z_{jt} instead of P_{jt}.

A crucial object in my analysis is the precision of the price signal $\beta_{jt}\sigma_{\varepsilon}^{-2}$, also referred to as *price informativeness* in the literature. If $\beta_{jt}\sigma_{\varepsilon}^{-2}$ is high, financial markets efficiently aggregate information and asset prices are informative about firm productivity. As a result, productive firms receive on average more capital, which improves the capital allocation through financial markets. I focus on the endogenous component β_{jt}.

As is evident now, the values of κ_H and κ_L do not matter for the price signal z_{jt}. They only pin down the identity of the marginal trader, which has a predictable effect on the price. For instance, the marginal trader is relatively optimistic for $\kappa_H - \kappa_L > 2$, which means that the price is set by a trader who received a private signal with positive idiosyncratic noise ($\eta_{ijt} > 0$). As a result, the price would be upward biased.[21] Choosing $\kappa_H = 2$ and $\kappa_L = 0$ ensures that the choice of position limits does not introduce a bias in share prices as the marginal trader has unbiased beliefs ($\eta_{ijt} = 0$).

The following proposition shows that the described equilibrium is unique. Moreover, the price P_{jt} is equal to the valuation of the *marginal trader* who is just indifferent between buying or not buying and who observed the private signal $s_{ijt} = z_{jt}$. Any trader who is more optimistic than the marginal trader ($s_{ijt} > z_{jt}$) buys two shares, whereas more pessimistic traders buy nothing.

Proposition 1.1. *Observing P_{jt} is equivalent to observing the signal (1.20) whenever K_{jt+1} is non-decreasing in z_{jt}. In the unique equilibrium, in which demand $x(s_{ijt}, P_{jt})$ is non-increasing in P_{jt}, the price is equal to the valuation of the trader with the private signal $s_{ijt} = z_{jt}$,*

$$P(z_{jt}) = \frac{1}{R_{t+1}} \tilde{\mathbb{E}} \{\Pi_{jt+1} | s_{ijt} = z_{jt}, z_{jt}\}. \tag{1.21}$$

1.3.4 Bond and Capital Market

The net supply of bonds is equal to zero, $\int_0^1 B_{it+1} di = 0$. Moreover, as all households are ex-ante identical, positions in bond markets are zero for all households, $\forall i : B_{it+1} = 0$. There is no excess demand or supply for bonds whenever the return on bonds R_{t+1} is equal to the return that traders expect to earn on the stock market. This is the case

[21] This mechanism plays an important role in Fostel and Geanakoplos (2012) and Simsek (2013) and is treated more in-depth in Appendix 1.B.

whenever

$$R_{t+1} = \frac{\int_0^1 \tilde{\mathbb{E}} \left\{ \Pi_{jt+1} | s_{ijt} = z_{jt}, z_{jt} \right\} dj}{\int_0^1 P_{jt} dj}, \tag{1.22}$$

which is derived by integrating (1.21) on both sides.

The aggregate value of the stock market is equal to the aggregate capital stock as all revenue from financial markets is invested by firms as follows from aggregating (1.4),

$$\int_0^1 P_{jt} dj = K_{t+1}. \tag{1.23}$$

1.3.5 Equilibrium Definition

In equilibrium, all traders choose the same information precision for all markets ($\forall ij$: $\beta_{ijt} = \beta_t$) and expect all other traders to choose the same.

Definition 1.1. *A competitive equilibrium consists of prices* $\{W_t, \rho_{jt+1}, P_{jt}, R_{t+1}\}$ *and allocations* $\{B_{it+1}, x_{ijt}, \beta_{ijt}, K_{jt+1}\}$ *such that:*

1. *Given prices* $\{W_t, \rho_{jt+1}, P_{jt}, R_{t+1}\}$ *and allocations* $\{x_{ijt}, \beta_{ijt}\}$*,* B_{it+1} *solves the household's problem P1.1.*

2. *Given prices* $\{P_{jt}, R_{t+1}\}$ *and allocations* $\{B_{it+1}, \beta_{jt}, K_{jt+1}\}$*,* $\{x_{ijt}, \beta_{ijt}\}$ *solve the trader's problem P1.2.*

3. *Prices are such that markets for labor, intermediate goods, shares, bonds, and capital clear, i.e., (1.12), (1.13), (1.18), (1.22) and (1.23) hold.*

1.4 Properties of the Equilibrium

In the following, I work out the properties of the equilibrium abstracting from the information production decision until the next section. I focus on how the allocation of capital can be expressed in terms of beliefs of the marginal trader and how these beliefs respond to both idiosyncratic and aggregate shocks. Next, I demonstrate how the allocation of capital through the stock market determines total factor productivity, which depends on the information choice. Finally, I show that the market allocation is distorted and derive the constrained efficient allocation.

As shown in (1.21), the beliefs of the marginal trader determine share prices. Therefore, they play a central role for the allocation of capital both in the cross-section and aggregate. The marginal trader's expectations are a weighted sum of the realization of both idiosyncratic and aggregate productivity and sentiment shocks,

$$\ln \tilde{\mathbb{E}} \left\{ A_{jt} | s_{ijt} = z_{jt}, z_{jt} \right\} = \omega_p \left(\beta_{jt} \right) a_t + \omega_a \left(\beta_{jt} \right) a_{jt} + \omega_\varepsilon \left(\beta_{jt} \right) \left(\varepsilon_{jt} - \varepsilon_t \right) + \omega_{s\varepsilon} \left(\beta_{jt} \right) \varepsilon_t + \frac{1}{2} \mathbb{V}_{jt}. \tag{1.24}$$

The weights $\{\omega_p(\beta_{jt}), \omega_a(\beta_{jt}), \omega_\varepsilon(\beta_{jt}), \omega_{s\varepsilon}(\beta_{jt})\}$ depend on information production β_{jt}. $\mathbb{V}_{jt}$ is posterior uncertainty of the marginal trader.[22]

The first two terms capture the effect of aggregate and idiosyncratic productivity shocks. If traders do not produce information ($\beta_{jt} = 0$), traders rely solely upon their prior a_t ($\omega_p(0) = 1$ and $\omega_a(0) = 0$). As traders produce more information, they shift weight from their prior to the realization of firm productivity ($\lim_{\beta_{jt}\to\infty} \omega_a(\beta_{jt}) = 1$). This leads to a higher sensitivity of the allocation of capital to firm-specific productivity shocks and improves the allocative efficiency of financial markets.

In contrast to the weights on productivity shocks, the weights on sentiment shocks are hump-shaped in β_{jt}. If traders do not produce information, they do not have a signal to learn from and, therefore, their expectations cannot be moved by noise ($\omega_\varepsilon(0) = \omega_{s\varepsilon}(0) = 0$). For perfect information, traders receive signals that do not contain noise in the first place ($\lim_{\beta_{jt}\to\infty} \omega_\varepsilon(\beta_{jt}) = \omega_{s\varepsilon}(\beta_{jt}) = 0$). If β_{jt} goes to either extreme, both idiosyncratic and aggregate sentiment shocks do not affect the beliefs of traders.

The aggregate sentiment shock ε_t moves the beliefs of traders although ε_t is common knowledge. This effect stems from the behavioral bias in Assumption 1.1. Traders correct the price signal z_{jt} for the aggregate sentiment shock but mistakenly believe that their private signal s_{ijt} is unaffected by sentiment and, therefore, do not correct their private signal in a similar way.

1.4.1 Capital Allocation and TFP

The results so far can be combined to derive the allocation of capital and total factor productivity in equilibrium as captured in the following proposition.

[22]See Appendix 1.A for derivations.

Proposition 1.2 (***Market Allocation***). *Under the market allocation,*

(i) firm capital is given by

$$K_{jt+1} = \frac{\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\}^{\theta}}{\int_0^1 \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\}^{\theta} dj} K_{t+1}. \tag{1.25}$$

(ii) the aggregate production function is

$$Y_t = A\left(a_{t-1}, \beta_{t-1}\right) K_t^{\alpha} \tag{1.26}$$

with total factor productivity

$$\ln A\left(a_{t-1}, \beta_{t-1}\right) = \underbrace{\frac{\alpha\theta}{\theta - 1}\left(a_{t-1} + \frac{\sigma_a^2}{2}\right)}_{exogenous} + \underbrace{\kappa^a\left(\beta_{t-1}\right)\sigma_a^2 - \kappa^{\varepsilon}\left(\beta_{t-1}\right)\sigma_{\varepsilon}^2}_{allocative\ efficiency}, \tag{1.27}$$

where $\kappa^a(\beta_{t-1})$ is increasing in β_{t-1} and $\kappa^{\varepsilon}(\beta_{t-1})$ is hump-shaped in β_{t-1}.

(iii) $A\left(a_{t-1}, \beta_{t-1}\right)$ is taking its minimum for some $\beta_{t-1} > 0$ if $\sigma_{\varepsilon}^2 > 1$.

(iv) $A\left(a_{t-1}, \beta_{t-1}\right)$ is monotonically increasing in β_{t-1} if $\sigma_{\varepsilon} \leq 1$.

The proposition's first part highlights that more capital is allocated to firms with higher realizations of the price signal z_{jt} whether it is driven by sentiment or productivity. Moreover, firm capital for all firms is proportional to aggregate investment K_{t+1}. Consequently, total factor productivity (TFP) has both an exogenous and endogenous component. The exogenous component is related to the realization of the aggregate productivity shock a_t, which mechanically increases the productivity of all firms. The endogenous component captures instead the allocational efficiency of financial markets, which is determined by aggregate information production β_t.

However, the market does not allocate capital efficiently given the available information. As traders are overconfident, expectations in (1.25) condition also on the private signal s_{ijt}, although it is uninformative after observing z_{jt}. In other words, P_{jt} behaves as if the precision of the market signal z_{jt} was $\beta_{jt} + \beta_{jt}\sigma_{\varepsilon}^{-2}$, although its true precision is $\beta_{jt}\sigma_{\varepsilon}^{-2}$. Therefore, the price *overreacts* to the price signal z_{jt}.[23]

[23] This distortion has been studied intensively in Albagli, Hellwig, and Tsyvinski (2011a, 2015, 2021) and is called the "information aggregation wedge." Its general equilibrium implications are studied in Albagli, Hellwig, and Tsyvinski (2017). In contrast to this paper, their model features a combination of

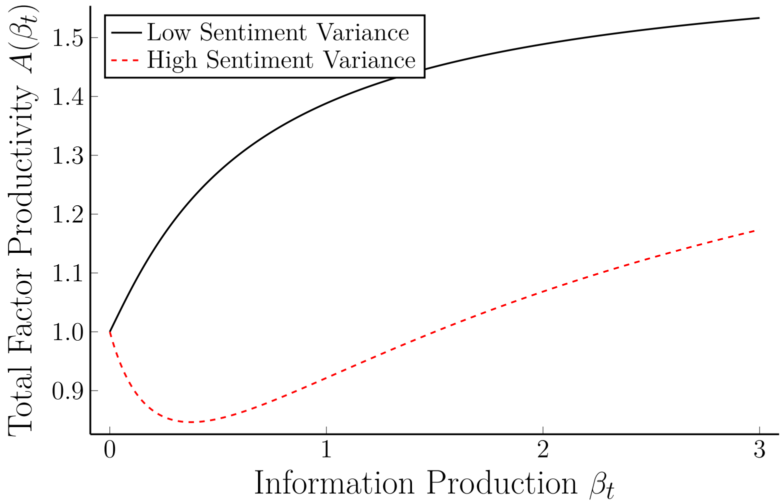

Figure 1.4: Total Factor Productivity and Information Production.

Notes: Total factor productivity as defined in (1.27). If the variance of sentiment shocks σ_ε^2 is sufficiently large, financial markets may worsen allocative efficiency relative to the case in which capital is equally distributed between firms ($\beta_t = 0$).

This distortion can be so severe that an increase in information production β_t leads to a decrease in TFP, as stated in Proposition 1.2 (iii) and seen in Figure 1.4. As traders produce more precise information, they also wrongly put more weight on their private signal. The overall effect on TFP depends on the balance between the beneficial effect of an increase in price informativeness $\beta_t \sigma_\varepsilon^{-2}$ and an increased weight on the private signal.

This price distortion leads to ex-ante misallocation of capital, i.e., output can be increased by reallocating capital between firms given the same publicly available information $\{z_{jt}\}$. A social planner would use the available information efficiently, leading to the *constrained efficient allocation* summarized in the following proposition.

---

rational and noise traders. Therefore, the information aggregation wedge does not require a behavioral price-setting traders. Furthermore, it arises in any informative financial market model in which traders learn from both a heterogeneous private signal and the price. It does not arise in models in which the information set of informed agents is homogeneous (Grossman and Stiglitz 1980) or in models where traders do not observe the price before submitting market orders (Kyle 1985). In the former case, informed agents cannot learn anything from the price, and in the latter, it is not possible to learn from the price before trading. Both of these models restrict the analysis to linear models, whereas non-linearity arises naturally in macroeconomic models; therefore, a different model is used here. See also Vives (2017) for an in-depth analysis in a linear setting.

Proposition 1.3 (***Constrained Efficient Allocation***)**.** *Under the constrained efficient allocation,*

(i) firm capital is

$$K_{jt+1}^{eff} = \frac{\mathbb{E}\left\{A_{jt}|z_{jt}\right\}^{\theta}}{\int_0^1 \mathbb{E}\left\{A_{jt}|z_{jt}\right\}^{\theta} dj} K_{t+1}. \tag{1.28}$$

(ii) total factor productivity is

$$\ln A_{t-1}^{eff} = \ln\left(\int_0^1 \mathbb{E}\left\{A_{jt-1}|z_{jt-1}\right\}^{\theta} dj\right)^{\frac{\alpha}{\theta-1}} = \underbrace{\frac{\alpha\theta}{\theta-1}\left(a_{t-1} + \frac{\sigma_a^2}{2}\right)}_{\text{exogenous}} + \underbrace{\frac{\alpha\theta}{2}\omega_a^{eff}\sigma_a^2}_{\text{allocative efficiency}}, \tag{1.29}$$

where $\omega_a^{eff} = \frac{\beta_{t-1}\sigma_\varepsilon^{-2}}{\sigma_a^{-2}+\beta_{t-1}\sigma_\varepsilon^{-2}}$. A_{t-1}^{eff} *is monotonically increasing in* β_{t-1}.

(iii) $A^{eff}\left(a_{t-1}, \beta_{t-1}\right) \geq A\left(a_{t-1}, \beta_{t-1}\right)$*, with strict inequality for* $\beta_{t-1} \in (0, \infty)$.

The *constrained efficient allocation* assigns the correct precision $\beta_t\sigma_\varepsilon^{-2}$ to the price signal z_{jt}. Relative to the market allocation, the constrained efficient allocation redistributes capital from firms that were previously too large to firms that were too small, as seen in Figure 1.5. Moreover, total factor productivity $A^{eff}\left(a_{t-1}, \beta_{t-1}\right)$ is monotonically increasing in aggregate information production β_{t-1} under the constrained efficient allocation, because the distortion due to traders' overconfidence is removed.

The following corollary provides conditions under which the market and constrained efficient allocation coincide.

Corollary 1.1. *The market allocation and constrained efficient allocation* ($K_{jt} = K_{jt}^{eff}$) *coincide if*

(i) symmetric information production β_t *goes to zero or infinity.*

(ii) the variance of firm-specific productivity shocks σ_a^2 *goes to zero or infinity.*

(iii) the variance of firm-specific sentiment shocks σ_ε^2 *goes to zero.*

As Corollary 1.1 shows, the behavioral bias disappears both when households have perfect information or when households have no information at all ($\beta_{jt} \in \{0, \infty\}$), as in both cases traders put zero weight on their private signal. There is also no distortion if the prior is arbitrarily noisy ($\sigma_a^2 \to \infty$), as in that case both the market and the efficient allocation put full weight on the price signal z_{jt}. If the prior is arbitrarily precise ($\sigma_a^2 \to 0$), the weight is zero for both. Finally, if the variance of sentiment shocks goes to

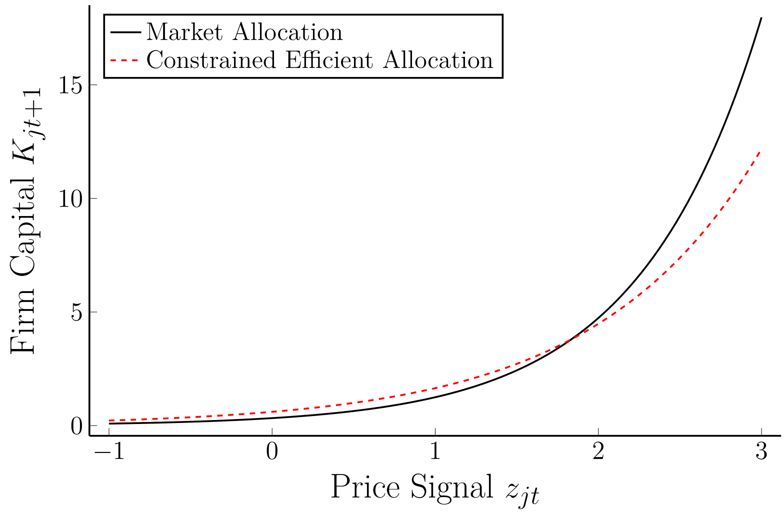

Figure 1.5: Market and Constrained Efficient Capital Allocation.

Notes: Market allocation of capital K_{jt} as in (1.25) and the constrained efficient allocation K_{jt}^{eff} as in (1.28).

zero, financial markets perfectly aggregate information as the price signal z_{jt} converges to firm productivity a_{jt} and private signals become irrelevant.

1.4.2 Aggregate Investment

Aggregate investment is in one of two regions. In the first region, traders consume during youth and investment is pinned down by $R_{t+1} = \frac{1}{\delta}$. In the second region, the interest rate is so high ($R_{t+1} > \frac{1}{\delta}$) that traders exhaust their wages for investment. Finally, $R_{t+1} < \frac{1}{\delta}$ cannot arise in equilibrium as investment would collapse to zero and the interest rate R_{t+1} would go to infinity. Taken together, aggregate investment is equal to

$$K_{t+1} = \min \left\{ \left(\alpha\delta A_t^{\alpha_Y} \left(\int_0^1 \tilde{\mathbb{E}} \left\{ A_{jt} | s_{ijt} = z_{jt}, z_{jt} \right\}^{\theta} dj \right)^{\frac{1}{\theta}} \right)^{\frac{1}{1-\alpha}}, W_t \right\}. \tag{1.30}$$

Aggregate shocks and information production determine investment in the elastic region. Aggregate productivity and sentiment shocks increase investment, as traders expect all firms to be more productive. An increase in aggregate information production

β_t has ambivalent effects, as it may increase or decrease TFP A_t and the average expectations of firm productivity $\int_0^1 \tilde{\mathbb{E}} \left\{ A_{jt} | s_{ijt} = z_{jt}, z_{jt} \right\}^\theta dj$ may be hump-shaped in β_t.

1.5 Main Results

As laid out in the prior section, the model has several sources of non-monotonicity. Not only may TFP be locally decreasing in aggregate information production β_t, but also aggregate investment K_{t+1} may be non-monotonic in β_t. These pathological cases are not due to a friction that can easily be removed but, rather, arise through the imperfect aggregation of information in a market with dispersed information.

Economic intuition tells us that better information usually leads to better economic outcomes. Indeed, the model allows for this intuition to hold by restricting the parameter space. As Corollary 1.1 shows, the distortion vanishes as the variance of firm-specific sentiment shocks σ_ε^2 goes to zero. As follows from Propositions 1.2 (iv) and later from the proof of Proposition 1.10, total factor productivity $A(a_t, \beta_t)$ and aggregate investment K_{t+1} are increasing in β_t for a neutral stance of sentiment ($\varepsilon_t = 0$) when $\sigma_\varepsilon^2 \leq 1$.[24] For the following analysis, I assume that more information production has beneficial effects as captured in the following Assumption.

Assumption 1.2. $\sigma_\varepsilon^2 \leq 1$, *such that*

(i) $\frac{\partial A(a_t, \beta_t)}{\partial \beta_t} \geq 0$.

(ii) $\left.\frac{\partial K_{t+1}(\beta_t)}{\partial \beta_t}\right|_{\varepsilon_t = 0} \geq 0$.

1.5.1 Aggregate Shocks and Information Acquisition

Recent experiences during stock and credit booms have raised concerns about increasing capital misallocation during these episodes (Gopinath et al. 2017; Doerr 2018; Gorton and Ordoñez 2020). My model can be used as a laboratory to think about the effects of productivity and sentiment shocks that may drive booms and their effects on the incentive to produce information, thereby affecting allocative efficiency.

Sentiment Shocks The following proposition starts with the effect of aggregate sentiment shocks.

[24] This threshold coincides with traders putting at most the same weight on their private and public signal. For $\sigma_\varepsilon^2 < 1$, traders put strictly more weight on the public signal than their private signal.

Proposition 1.4. *There exists a threshold $\bar{\varepsilon}$, such that*

(i) information production is increasing in the sentiment shock if $\varepsilon_t < \bar{\varepsilon}$,

(ii) information production is decreasing in the sentiment shock if $\varepsilon_t > \bar{\varepsilon}$,

(iii) the threshold $\bar{\varepsilon}$ is negative for $\theta > \frac{1}{1-\alpha}$ and positive for $\theta < \frac{1}{1-\alpha}$.

Proposition 1.4 shows that the effect of small sentiment shocks ($\varepsilon_t \approx 0$) on information production is ambiguous and depends on the parameters of the model. However, sentiment shocks always crowd out information production once they are sufficiently large. Moreover, note that aggregate sentiment shocks do not affect price informativeness directly but only through information production.

At first, it may seem surprising that *aggregate* sentiment shocks crowd out information production, especially as in my model, *firm-specific* sentiment shocks incentivize information production in the first place. This is because knowledge about an *aggregate* sentiment shock changes the incentive to produce *firm-specific* information. In particular, there are two direct channels through which sentiment shocks affect the incentive to produce information.

1. Sentiment shocks make valuations more extreme. As a result, trading becomes less *information-sensitive*. A relatively imprecise yet unbiased signal is sufficient to identify grossly mispriced firms and trade accordingly. Moreover, sentiment shocks make subtle mispricing rarer, for which precise information is helpful as shown in Figure 1.6. This effect *crowds out* information production for positive and negative sentiment shocks equally. Moreover, firms with such subtle mispricing must appear relatively unproductive in an otherwise exuberant market and consequently attract less capital, as in Figure 1.7. This *relative size* effect *crowds out* information production for positive sentiment shocks, as learning about smaller firms is unattractive.

2. Aggregate sentiment shocks increase aggregate investment K_{t+1}, which leads to an increase in the *absolute size* of all firms, encouraging more information production.

To further build intuition for this result, I use (1.25) in (1.17) to rewrite the marginal benefit of information production evaluated at the symmetric equilibrium,

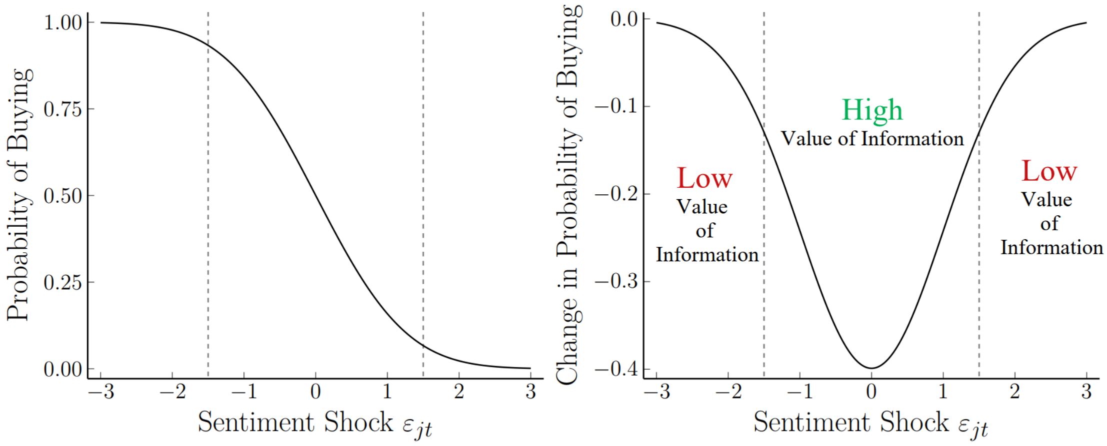

Figure 1.6: Probability of Buying and Sentiment Shocks.

Notes: Left panel: The probability of buying depending on the realization of the firm-specific sentiment shock ε_{jt}. Right panel: The derivative of the probability of buying. The trading decision is most information-sensitive, i.e., varies most with the realization of the sentiment shock ε_{jt}, around $\varepsilon_{jt} = 0$.

$$\widetilde{MB}\left(\beta_{ijt}, \beta_{jt}\right)\Big|_{\beta_{ijt}=\beta_{jt}} \propto$$
$$\tilde{\mathbb{E}}\left\{\underbrace{\underbrace{\frac{\partial \mathcal{P}\left\{x_{ijt}=2\right\}}{\partial \beta_{ijt}}\Bigg|_{\beta_{ijt}=\beta_{jt}}}_{\textit{Information-Sensitivity}} \left(\frac{K_{jt+1}}{K_{t+1}}\right)^{\frac{\theta-1}{\theta}}}_{\textit{Relative Size}} \underbrace{K_{t+1}^{\alpha}}_{\textit{Absolute Size}} \left(A_{jt} - \tilde{\mathbb{E}}\left\{A_{jt} | s_{ijt} = z_{jt}, z_{jt}\right\}\right)\right\}. \tag{1.31}$$

The *information sensitivity* channel materializes through the interaction of the change in the buying probability with the distribution of firm-specific sentiment shocks ε_{jt}. In the symmetric equilibrium ($\beta_{ijt} = \beta_{jt}$), traders expect to buy whenever they are more optimistic than the marginal trader, i.e., $s_{ijt} \geq z_{jt} \iff \eta_{ijt} \geq \varepsilon_{jt}$. The resulting probability of buying is $\Phi(-\varepsilon_{jt})$ where $\Phi(\cdot)$ is the standard-normal cdf. Consequently, the derivative of the buying probability with respect to the realization of the firm-specific sentiment shock ε_{jt} is $-\phi(\varepsilon_{jt})$ where $\phi(\cdot)$ is the standard-normal pdf. As shown in Figure 1.6, the trading decision is most elastic for relatively small realizations of the firm-specific sentiment shock ε_{jt}. However, aggregate sentiment shocks push the distribution of ε_{jt} to the more inelastic regions toward the extremes.

Formally, this effect can be captured by multiplying the change in the buying probability with the distribution of sentiment shocks,

$$\phi\left(\varepsilon_{jt}\right) f\left(\varepsilon_{jt}\right) \propto \exp\left\{-\frac{\varepsilon_t^2}{2\left(1+\sigma_\varepsilon^2\right)}\right\} \tilde{f}(\varepsilon_{jt}), \tag{1.32}$$

where $f\left(\varepsilon_{jt}\right)$ is the pdf of $\varepsilon_{jt} \sim \mathcal{N}\left(\varepsilon_t, \sigma_\varepsilon^2\right)$ and $\tilde{f}\left(\varepsilon_{jt}\right)$ is the pdf of ε_{jt} as if its distribution was $\mathcal{N}\left(\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}, \frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\right)$. The *information sensitivity* channel is captured by the term $\exp\left\{-\frac{\varepsilon_t^2}{2(1+\sigma_\varepsilon^2)}\right\}$, which is symmetrically decreasing around zero. Somewhat surprisingly, the decline in *information sensitivity* does not depend on the actual pass-through of sentiment shocks to expectations. The reason can be found in the trading decision, which does not depend on the actual mispricing caused by sentiment shocks but only on the realization of the firm-specific sentiment shock ε_{jt}. Therefore, aggregate sentiment shocks can discourage information production, even if they do not significantly affect actual prizes.

The additional effect of a decline in information sensitivity on the *relative size* of firms, for which information remains valuable, is captured by taking expectations of the relative firm-size $\left(\frac{K_{jt+1}}{K_{t+1}}\right)^{\frac{\theta-1}{\theta}}$ with the density $\tilde{f}\left(\varepsilon_{jt}\right)$,

$$\int_0^1 \tilde{f}(\varepsilon_{jt}) \left(\frac{K_{jt+1}}{K_{t+1}}\right)^{\frac{\theta-1}{\theta}} dj \propto \exp\left\{-\left(\theta-1\right)\omega_{s\varepsilon}\varepsilon_t\right\}, \tag{1.33}$$

where $\omega_{s\varepsilon} = \frac{\sqrt{\beta_t}}{\sigma_a^{-2}+\beta_t\left(1+\sigma_\varepsilon^{-2}\right)}$. For a positive sentiment shock, information production becomes effectively directed toward smaller firms, weakening the incentive to produce information. This channel is illustrated in Figure 1.7 and formally captured by the term $\exp\left\{-\left(\theta-1\right)\omega_{s\varepsilon}\varepsilon_t\right\}$.

The *relative size* effect is increasing in the elasticity of substitution and in the pass-through of aggregate sentiment shocks $\omega_{s\varepsilon}$, which is non-monotonic in information production β_t. If intermediate goods are close substitutes, firms that are perceived as unproductive attract very little capital. Moreover, if aggregate sentiment shocks have a large effect on expectations, underpriced firms will be even smaller, making information production even less attractive.

The *absolute size* effect is captured by changes in aggregate investment. Restricting

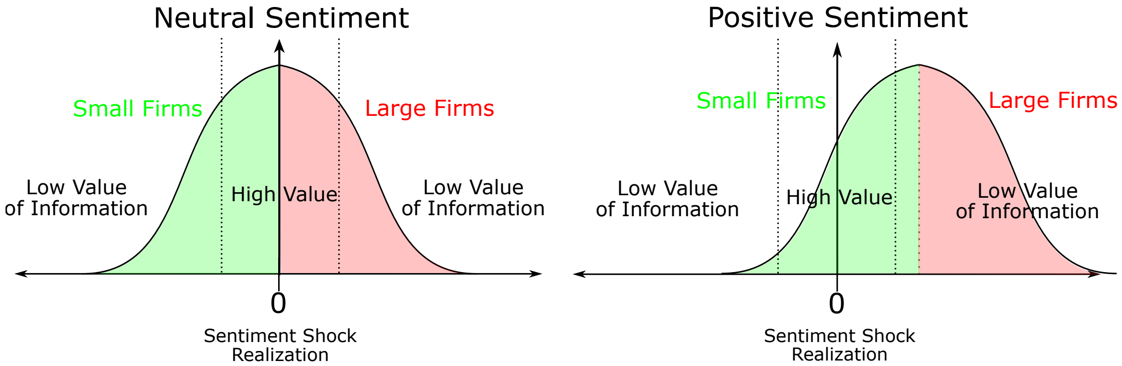

Figure 1.7: Illustration Size Channel.

Notes: Firms that are fairly priced and for which information is valuable are in the center of the firm-size distribution under neutral sentiment ($\varepsilon_t = 0$). In contrast, for positive sentiment shocks, the same firms are in the left part of the firm-size distribution as they appear to be unproductive relative to other firms.

our attention to shocks for which $K_{t+1} < W_t$ leads to

$$K_{t+1}^{\alpha} \propto \exp\left\{\frac{\alpha}{1-\alpha}\omega_{s\varepsilon}\varepsilon_t\right\}. \tag{1.34}$$

As long as traders do not fully invest their wages, the *absolute size* effect can be captured by the term $\exp\left\{\frac{\alpha}{1-\alpha}\omega_{s\varepsilon}\varepsilon_t\right\}$. Intuitively, the effect on investment is stronger when α and, therefore, the returns to scale increase. A further increase in the sentiment shock is ineffectual for the *absolute size* channel once traders fully invest their wages but incentivizes nonetheless more information production through an increase in the value of resources, as captured by $\lambda_t = \max\{1, R_{t+1}\delta\}$ in (P1.2).

Putting all three effects together yields the marginal benefit of information production for a given symmetric information production choice ($\beta_{ijt} = \beta_{jt}$) as

$$\widetilde{MB}\left(\beta_{ijt}, \beta_{jt}\right)\Big|_{\beta_{ijt}=\beta_{jt}} \propto \exp\left\{\underbrace{-\frac{\varepsilon_t^2}{2\left(1+\sigma_\varepsilon^2\right)}}_{\textit{Information-Sensitivity}} \underbrace{-\left(\theta-1\right)\omega_{s\varepsilon}\varepsilon_t}_{\textit{Relative Size}} + \underbrace{\frac{\alpha}{1-\alpha}\omega_{s\varepsilon}\varepsilon_t}_{\textit{Absolute Size}}\right\}. \tag{1.35}$$

For the empirically plausible calibration $\theta - 1 > \frac{\alpha}{1-\alpha}$, positive sentiment shocks always *crowd out* information, as the increase in aggregate investment is dominated by a larger decrease in size of fairly priced firms. Conversely, negative sentiment shocks initially *crowd in* information, as fairly priced firms turn out to be relatively large although aggregate investment goes down. Finally, the *information sensitivity* channel always dominates for large shocks.

Productivity Shocks Productivity shocks have quite different effects on the incentive to produce information. Whereas sentiment shocks affect trading in multiple ways, productivity shocks leave the buying decisions unaffected. The reason is that traders believe that sentiment shocks affect only other traders, whereas productivity shocks affect all traders. The only channel through which productivity shocks change the incentive to produce information is through an increase in aggregate investment (*absolute size* channel) and dividends for all firms. This result is captured in the following proposition.

Proposition 1.5. *Positive (negative) productivity shocks crowd in (out) information.*

The model provides a rationale for the different impact of "good" and "bad" booms as in Gorton and Ordoñez (2020). Whereas productivity-driven "good" booms increase information production and improve allocative efficiency, sentiment-driven "bad" booms *crowd out* information and increase capital misallocation. The results of Propositions 1.4 and 1.5 are pictured in Figure 1.8.

1.5.2 Real Feedback

Financial markets do not only react to aggregate shocks, but also shape the economy's response to aggregate shocks. In the following, aggregate shocks hit an economy that is in steady state. Whether shocks amplify or dampen the effect of shocks on output is determined relative to an economy for which the information choice is fixed at the endogenous steady state information production β^*.

In the economy with fixed information production β^*, the only effect of aggregate shocks is the direct effect on TFP and investment. Positive shocks of both types increase investment, whereas only productivity shocks also have a direct effect on TFP. The opposite is true for negative shocks, which depress investment and TFP in the case of productivity shocks. Whereas the direct effect of aggregate shocks are straightforward, the indirect effects under endogenous information production are more subtle.

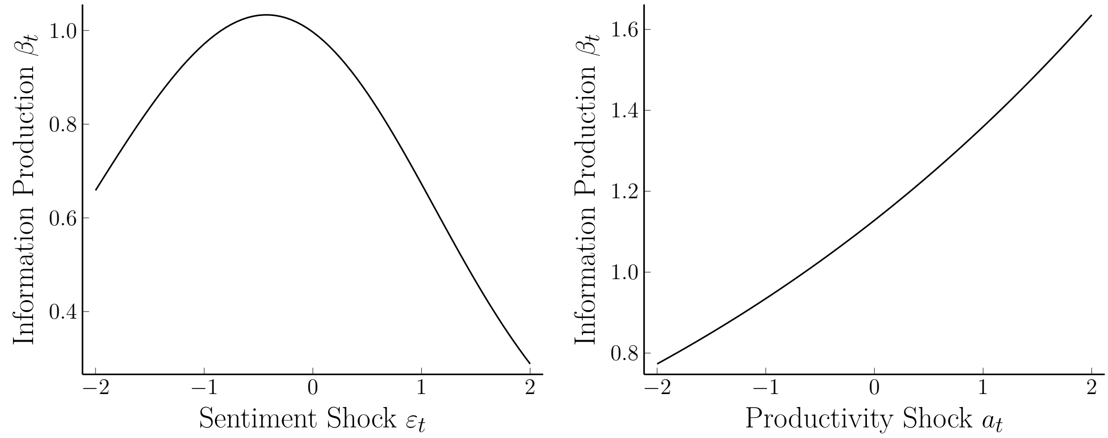

Figure 1.8: Information Production and Aggregate Shocks.

Notes: Information production is non-monotonic in the sentiment shock, with the peak $\bar{\varepsilon}$ being negative for $\theta > \frac{1}{1-\alpha}$. Information production is monotonically increasing in the productivity shock.

There are two indirect effects of sentiment shocks, both of which lead to a non-monotonic response of investment and output. First, sentiment shocks affect the allocative efficiency of financial markets through their effect on information production, which also decreases investment. The cost of misallocation through a decrease in information production depends on the elasticity of substitution between intermediate goods. If the elasticity of substitution is large, misallocation between firms is costly. Moreover, a high elasticity of substitution also leads to a stronger decrease in information production for a positive sentiment shock. In contrast, the costs of misallocating capital are low if the elasticity of substitution is small.

The second effect concerns the pass-through of sentiment shocks. Since traders are unaffected by sentiment if they produce either no or perfect information ($\beta_t \in \{0, \infty\}$), the effect of a given sentiment shock on beliefs must be maximized for an interior value of β_t. Therefore, a change in information production by traders may increase or decrease the effect of a given sentiment shock on their beliefs, which depends on whether steady state information production β^* is above or below the threshold $\frac{\sigma_a^{-2}}{1+\sigma_\varepsilon^{-2}}$. If β^* is above (below) the threshold, then the effect of aggregate sentiment shocks is locally increasing (decreasing) in information production. For example, a positive sentiment shock crowds out information production, leading to an amplification of the shock if the resulting precision choice β^* is still above the threshold $\frac{\sigma_a^{-2}}{1+\sigma_\varepsilon^{-2}}$.

The results for the case with $\beta^* < \frac{\sigma_a^{-2}}{1+\sigma_\varepsilon^{-2}}$ are captured in the following proposition and visualized in Figure 1.9.

Proposition 1.6. *(i) For $\theta > \frac{1}{1-\alpha}$ and $\beta^* < \frac{\sigma_a^{-2}}{1+\sigma_\varepsilon^{-2}}$, information production dampens positive sentiment shocks.*

(ii) Large positive sentiment shocks eventually lead to a decrease in aggregate investment if $\lim_{\varepsilon_t\to\infty}\sqrt{\beta_t(\varepsilon_t)}\varepsilon_t = 0$.

The second result of Proposition 1.6 captures that the costs of misallocation must be eventually so large that they outweigh the investment-stimulating effect of sentiment shocks. Moreover, the direct effect of sentiment shocks vanishes as sentiment shocks grow large, as long as information production declines fast enough. This result is captured in the following corollary.

Corollary 1.2. *If information production declines fast enough as sentiment shocks grow large, then aggregate investment approaches its level without information production* $\beta_t = 0$. *Formally,*

$$\lim_{\varepsilon_t\to\pm\infty}\sqrt{\beta_t(\varepsilon_t)}\varepsilon_t = 0 \Rightarrow \lim_{\varepsilon_t\to\pm\infty} K(\beta_t(\varepsilon_t),\varepsilon_t) = K(0,\varepsilon_t). \tag{1.36}$$

These results may initially seem counterintuitive, since sufficiently large positive sentiment shocks possibly decrease prices and output. However, the decrease in information production must eventually outweigh the expansionary effect of sentiment shocks as the pass-through of sentiment shocks goes to zero. Moreover, this section studies only *anticipated* sentiment shocks. If the same shock was unknown prior to the information production decision, positive sentiment shocks would unambiguously increase investment as in the economy with exogenous information precision.

Similar forces are active for negative shocks with the exception that negative sentiment shocks initially crowd in information production if the elasticity of substitution is large enough ($\theta > \frac{1}{1-\alpha}$). If strong enough, this indirect effect can even lead to negative sentiment shocks being initially expansionary. In contrast, if the elasticity of substitution is relatively small ($\theta < \frac{1}{1-\alpha}$), then negative sentiment shocks always crowd out information production and are, therefore, initially amplified.

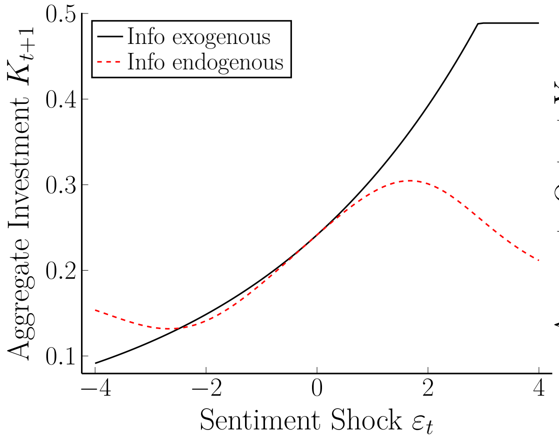

Figure 1.9: Amplification and Dampening for Sentiment Shocks.

Notes: Whether information production dampens or amplifies sentiment shocks depends on the size of the shock and the parameters. As information production affects both allocative efficiency and the pass-through of sentiment shocks, large sentiment shocks eventually drive information so low that investment and output decrease.

Similar to the previous section, the indirect effect of productivity shocks leads generally to amplification. As follows from Proposition 1.5, positive productivity shocks crowd in information production, leading to an improvement in the allocation of capital and incentivizes additional investment. Therefore, compared to the economy with fixed information precision, the reaction of both output and investment to a productivity shock are larger if information precision is allowed to adjust, as can be seen in Figure 1.10. This result is captured in the following proposition.

Proposition 1.7. *Information production amplifies productivity shocks.*

Numerical Illustration

This section provides a numerical illustration of booms driven by productivity and sentiment shocks, focusing on the region of parameters and shocks for which sentiment shocks are expansionary and dampened by information production. To capture the notion of booms, aggregate shocks build up over time according to the auto-regressive

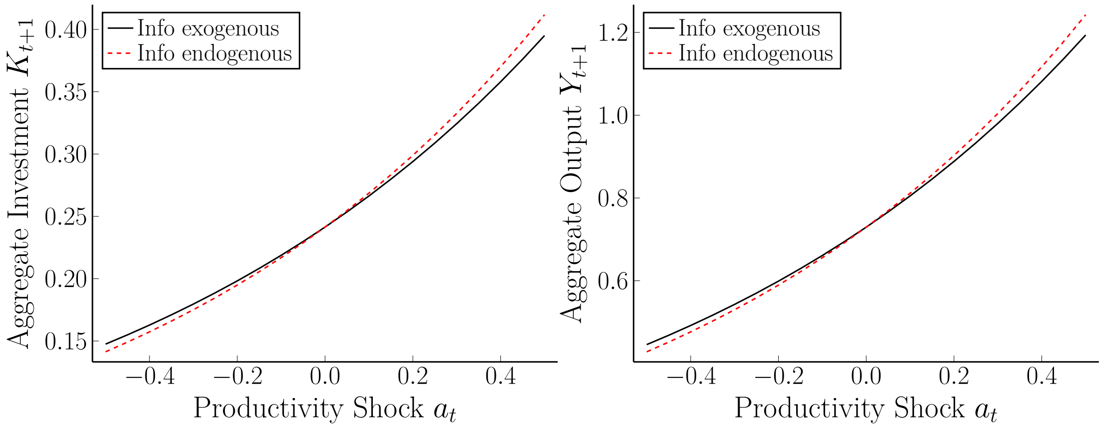

Figure 1.10: Amplification of Productivity Shocks.

Notes: Productivity shocks crowd in information production, leading to an additional increase in TFP and Aggregate Investment K_{t+1}. As a result, the effect of productivity shocks is amplified.

process

$$y_t = \begin{cases} \rho y_{t-1} + \zeta & t \in [0, B] \\ 0 & \text{otherwise} \end{cases}, \tag{1.37}$$

where $y_t \in \{a_t, \varepsilon_t\}$ is the aggregate shock, ζ is a constant innovation, $\rho \in (0, 1)$ is the persistence, and B captures the duration of the boom. After the boom is over, the aggregate shock returns to a neutral stance and remains there.

The expansionary effect of sentiment shocks is dampened, as can be seen in Figure 1.11. Optimistic expectations lead to an increase in investment, but traders decide to cut back on information production, which decreases the allocative efficiency of financial markets. In total, output still increases because the sentiment shock leads to an offsetting increase in investment. In this case, the endogenous response of traders dampens the effect of a positive sentiment shock.

In contrast, productivity-driven booms are generally amplified by an increase in information production, as seen in Figure 1.12, mirroring the result from Figure 1.10 and Proposition 1.7. Expectations of higher productivity tomorrow cause an increase in investment today, which triggers more information production. As a result, the endogenous response of traders amplifies the effect of productivity shocks. Times of high productivity are also times in which financial markets allocate capital efficiently.

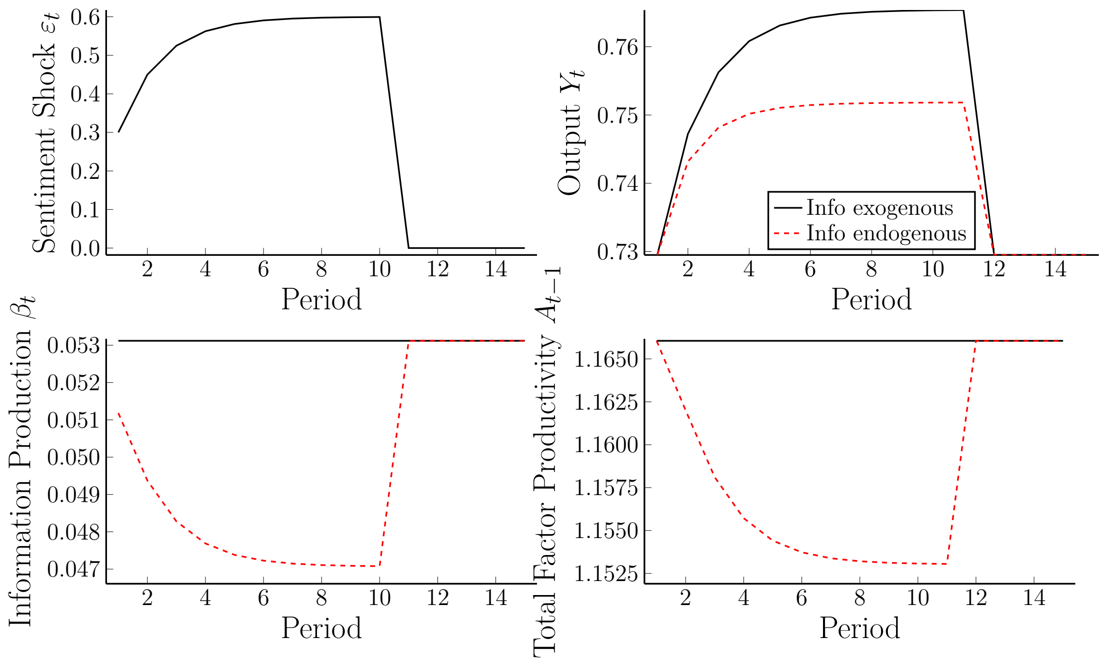

Figure 1.11: Numerical Simulation Sentiment-Driven Boom.

Notes: Sentiment-driven booms are dampened by information production.

1.6 Is there a Role for Policy?

After studying the positive properties of the model, I turn now to the normative implications. There are two sources of inefficiency in my model. First, there are two externalities with respect to the information production decision that work in opposite directions. On the one hand, traders produce information to extract rents from other traders and ignore the negative effects they impose on others. On the other hand, traders do not reap the benefits of an improving capital allocation due to information production and, therefore, do not ignore this positive externality of information production. Whether information production is inefficiently high or low depends on the strength of the rent-stealing motive relative to the usefulness for information in allocating capital.

Second, traders' overconfidence distorts the allocation of capital between firms as described in section 1.4.1, and lets aggregate sentiment shocks drive investment. A state- and price-dependent tax/subsidy on dividends is sufficient to fix this distortion. The formal analysis has been delegated to Appendix 1.E as the focus of this paper is on

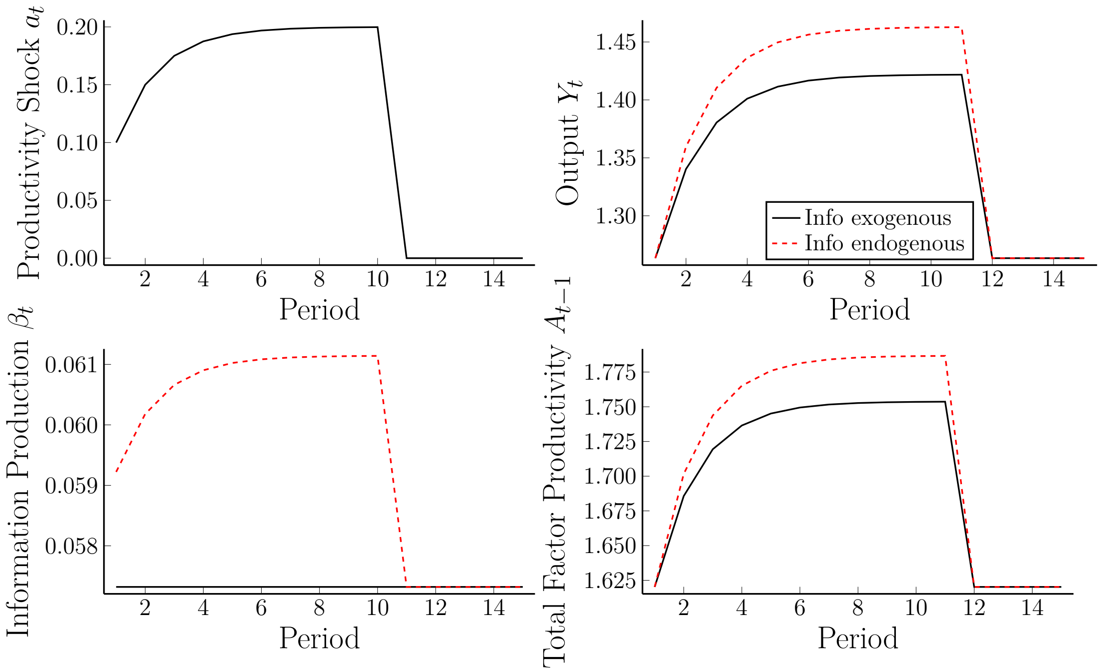

Figure 1.12: Numerical Simulation Productivity-Driven Boom.

Notes: Productivity-driven booms are amplified by information production.

information production.

For the following welfare analysis, I abstract from well-known inter-generational trade-offs using a two-period model. Traders are born with an endowment, produce information, and buy shares. Production takes place in the second period and the final good sector combines intermediate goods into the final good without labor,

$$Y_1 = \left(\int_0^1 Y_{j1}^{\frac{\theta-1}{\theta}} \, dj \right)^{\frac{\alpha\theta}{\theta-1}} . \tag{1.38}$$

The setup is otherwise identical to the main model.

The section proceeds in the following steps. First, I explain in detail why information production is inefficient in the competitive equilibrium and for which parameters information production is likely to be either too high or too low. Next, I consider the optimal intervention if the social planner can only steer the information choice, but cannot decide on aggregate investment. Finally, I propose an implementation for a policy that incentivizes or discourages information production.

1.6.1 Static Information Choice

Endow a social planner with the ability to dictate a level of information production β_{ij0} to each trader, but households autonomously decide on consumption and investment.[25] Moreover, the social planner observes aggregate shocks $\{a_0, \varepsilon_0\}$ before taking her decision. The corresponding maximization problem is

$$\max_{\{\beta_{ij0}\}} \quad C_0 + \delta C_1 - \int_0^1 IA\left(\beta_{ij0}\right) dj \tag{SP1.1}$$

$$s.t. \quad C_1 = A_0\left(\{\beta_{ij0}\}\right) K_1^{\alpha} \tag{1.39}$$

$$C_0 = W_0 - K_1 \tag{1.40}$$

$$K_1 = \min\left\{\left(\alpha\delta A_t^{\alpha_Y}\left(\int_0^1 \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\}^{\theta} dj\right)^{\frac{1}{\theta}}\right)^{\frac{1}{1-\alpha}}, W_t\right\} \tag{1.41}$$

$$\beta_{ij0} \geq 0. \tag{1.42}$$

The social planner maximizes welfare subject to a number of constraints. Old age consumption is equal to aggregate production as in (1.39), for which total factor productivity $A\left(\{\beta_{ij0}\}\right)$ depends on information production. Youth consumption as in (1.40) depends on aggregate investment, which also depends implicitly on information production, as seen in (1.41). Finally, (1.42) is a non-negativity constraint on information production.

Since all traders and firms are ex-ante homogeneous, the social planner chooses the same level of information precision $\beta_0 = \beta_{ij0}$ for all traders and markets. The marginal benefit of increasing β_0 for the social planner is

$$MB^{SP}\left(\beta_0\right) = \delta \underbrace{\frac{\partial A_0\left(\beta_0\right)}{\partial \beta_0}}_{\textit{Change in TFP}} K_1\left(\beta_0\right)^{\alpha} + \left(\delta\alpha A_0\left(\beta_0\right) K_1\left(\beta_0\right)^{\alpha-1} - 1\right) \underbrace{\frac{\partial K_1\left(\beta_0\right)}{\partial \beta_0}}_{\textit{Change in Investment}}. \tag{1.43}$$

The social planner targets both TFP $A_0\left(\beta_0\right)$ and aggregate investment $K_1\left(\beta_0\right)$. Note

[25] The full planner's problem is covered in the Appendix 1.E.

that the latter effect is only relevant if aggregate investment is inefficiently high or low, which is generally the case due to the price distortion described in section 1.4.1 and aggregate sentiment shocks.

The first observation is that (1.43) does not coincide with the marginal benefit in (1.17). Moreover, the difference cannot be expressed in the form of a simple wedge. This finding leads directly to the following proposition.

Proposition 1.8. *Information production is inefficiently high or low in the competitive equilibrium.*

The reason for this result is that the information production decision is subject to two externalities with opposing effects. First, traders produce information to extract rents from other traders, i.e., they ignore a negative externality. In other words, traders seek to get a larger piece of a fixed pie of trading profits. Second, as atomistic traders take prices as given, they do not take into account the allocation-improving effect of *collective* information production, i.e., they do not take into account the positive *spillover* of information production. If all traders produce more precise information, the allocation of capital improves and aggregate productivity increases. Both externalities are explained in more detail in what follows.

Traders think that their information allows them to systematically buy undervalued shares, thus earning a rent. Producing more precise information allows them to better identify profitable trading opportunities. However, if trader ij decides to buy shares, these shares cannot be bought by another trader. Consequently, any rent that accrues to trader ij must be subtracted from rents that are earned by other traders. Although this rent-extracting behavior drives information production in the first place, it can also lead to inefficiently high information production.

In contrast, the social benefit of information production stems from an improvement in the allocation of capital. However, this effect only arises if traders *collectively* produce more precise information . In contrast, individual information production and trading have only infinitesimal effects on prices, which are ignored by price-taking traders in their information production decision. Therefore, information production has a positive *spillover*, which can lead to information production being too low in the competitive equilibrium.

Two simple examples can be constructed to showcase situations in which information production is unambiguously too high or too low in equilibrium. First, assume that the social planner confiscates rents and redistributes them equally. Traders have no incentive to produce information, but the social planner still values information for its effect on the allocation of capital. In this case, information production is inefficiently low. Second, let firm output be given exogenously, such that $Y_{jt} = A_{jt}$. Traders can still make bets on firm revenue by trading shares. However, information production has no social value as production is given exogenously. In this case, information production is inefficiently high.

Cases with too much and too little information can be produced in the clearest way by varying the elasticity of substitution, which captures the importance of capital allocation for aggregate productivity. The comparison between the planner's choice and the competitive outcome is shown in Figure 1.13. First, consider the case of no substitution with ($\theta \to 0$). In that case, every intermediate good is necessary to produce the final good and the necessary mix is pinned down by firm productivities. It follows that an equal distribution of capital becomes optimal and information about firm productivity has no social value since it no longer aids the optimal allocation. In other words, TFP becomes flat in information. Nonetheless, traders find it profitable to produce information as firm revenue still depends on the realization of firm productivity.

Second, if the elasticity of substitution grows arbitrarily large ($\theta \to \infty$), intermediate goods become increasingly substitutable and the allocation of capital more important. In contrast, traders find it at some point unattractive to produce information as most firms will be unable to attract capital, and only the firm with the highest combination of productivity *and* sentiment shock receives the economy's capital stock. As a result, the planner's information precision choice is eventually above the outcome in the competitive equilibrium. The market underproduces information exactly when it is most valuable.

1.6.2 Responding to Aggregate Shocks

The social planner increases information production in response to both negative and positive sentiment shocks for two reasons. First, traders expect that trading becomes less information-sensitive when a sentiment shock hits the economy. However, the value of information for the allocation of capital is only affected insofar as aggregate invest-

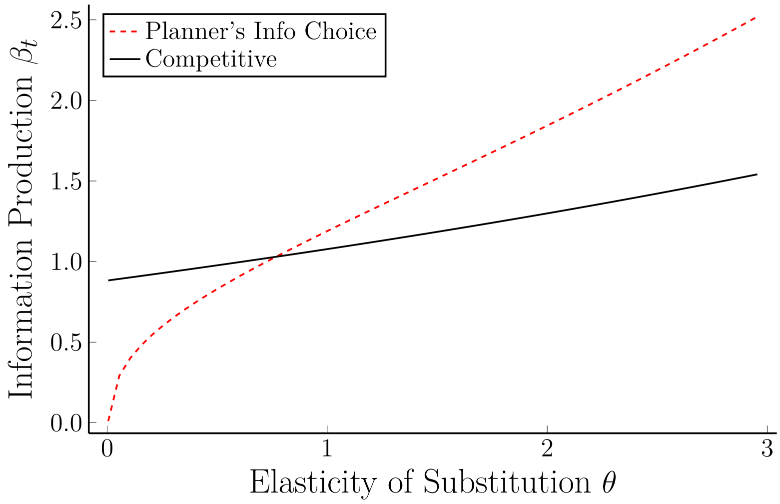

Figure 1.13: Planner's and Market's Information Production and Elasticity of Substitution.

Notes: Low elasticity of substitution: Too much information production. High elasticity of substitution: Too little information production.

ment changes. Second, the social planner also seeks to steer investment through the information production decision. For example, when a positive sentiment shock hits the economy, then producing more precise information eventually dampens the impact of the sentiment shock. The resulting response is asymmetric for positive and negative shocks, as positive shocks increase investment, which makes information more valuable, whereas negative shocks lower investment.

In contrast, the social planner's choice in response to productivity shocks is procyclical. An increase in exogenous productivity a_0 incentivizes information production in two ways. First, note that TFP can be decomposed into two parts, $A_0 = A_0(a_0)\, A_0(\beta_0)$, where the first is exogenously driven by a_0 and the second is related to allocative efficiency through β_0. Therefore, an increase in a_0 amplifies the improvement in the allocative efficiency through an increase in β_0. Second, positive productivity shocks lead to an increase in investment which additionally incentivizes information production.

The social planner's choice is shown in comparison to the competitive equilibrium in Figure 1.14. For the chosen parameters, the social planner chooses generally more pre-

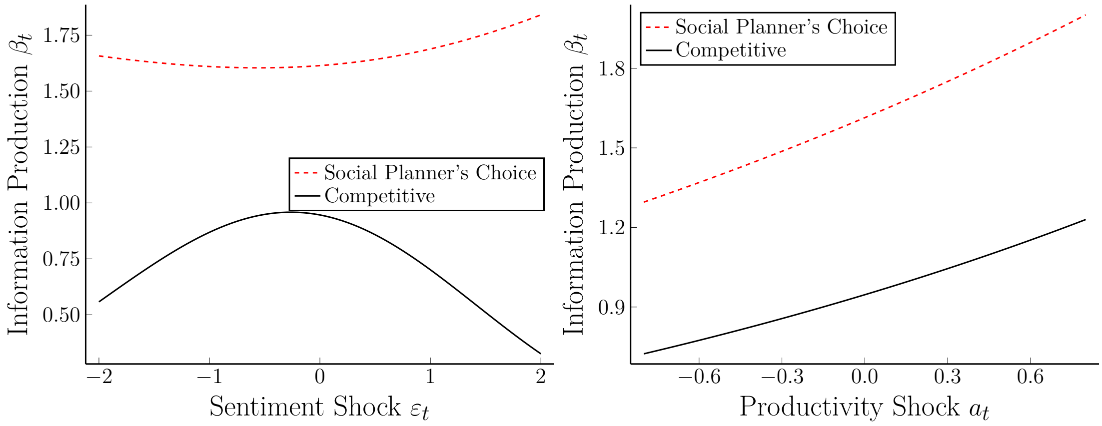

Figure 1.14: Planner's and Market's Information Production and Aggregate Shocks.

cise information than traders choose in the competitive equilibrium. Sentiment shocks widen the difference between the social planner's choice and the competitive outcome, whereas productivity shocks leave the gap largely unchanged. How the social planner can implement this policy is discussed in the following section.

1.6.3 Implementation

Traders are taking a gamble when they decide to buy shares in a given asset. The social planner can incentivize information production by increasing the stakes for each trade. This idea can be implemented through a redistribution of dividends between over- and underperforming firms as shown in the following corollary.

Corollary 1.3. *A state-dependent tax/subsidy $\tau(a_{jt}, z_{jt})$ on dividends with the properties,*

(i) No price distortions: $\tilde{\mathbb{E}}\left\{\tau(a_{j0}, z_{j0})\,\Pi_{j1} | s_{ij0} = z_{j0}, z_{j0}\right\} = \tilde{\mathbb{E}}\left\{\Pi_{j1} | s_{ij0} = z_{j0}, z_{j0}\right\}$

(ii) Monotonicity of beliefs: $\frac{\partial \tilde{\mathbb{E}}\{\tau(a_{j0}, z_{j0})\Pi_{j1}|s_{ij0}, z_{j0}\}}{\partial s_{ij0}} > 0$

(iii) Monotonicity of prices: $\frac{\partial \tilde{\mathbb{E}}\{\tau(a_{j0}, z_{j0})\Pi_{j1}|s_{ij0}=z_{j0}, z_{j0}\}}{\partial z_{j0}} > 0$

encourages (discourages) information productionwhen

$$\tau\left(a_{j0},z_{j0}\right)\geq(\leq)\,1\iff\Pi_{j1}\geq\tilde{\mathbb{E}}\left\{\Pi_{j1}|s_{ij0}=z_{j0},z_{j0}\right\}\ \textit{and}\ \frac{\partial\mathcal{P}\left\{x_{ij0}=2\right\}}{\partial\beta_{ij0}}\geq 0$$
$$\tau\left(a_{j0},z_{j0}\right)\leq(\geq)\,1\iff\Pi_{j1}\leq\tilde{\mathbb{E}}\left\{\Pi_{j1}|s_{ij0}=z_{j0},z_{j0}\right\}\ \textit{and}\ \frac{\partial\mathcal{P}\left\{x_{ij0}=2\right\}}{\partial\beta_{ij0}}\leq 0$$

and the inequalities are strict for at least some realizations of $\{a_{j0},z_{j0}\}$.

Intuitively, the social planner can make assets more or less risky by taxing/subsidizing dividends depending on realized productivity and market expectations. For example, subsidizing dividend payments of over-performing firms and taxing under-performing firms makes any investment riskier and information production more attractive. To avoid distorting prices, subsidies and taxes must offset each other in expectations.

As an illustration, the following combination of a tax $\tau\left(a_{j0},z_{j0}\right)$ and a lump-sum transfer $T\left(a_{j0},z_{j0}\right)$ encourage information production, where I assume $a_0=-\frac{\sigma_a^2}{2}$ as a normalization,

$$\tau\left(a_{j0},z_{j0}\right)=\begin{cases}0 & a_{j0}<\omega_a z_{j0}\\ 1 & a_{j0}\geq\omega_a z_{j0}\end{cases}\tag{1.44}$$

$$T\left(a_{j0},z_{j0}\right)=\begin{cases}0 & a_{j0}<\omega_a z_{j0}\\ \tilde{\mathbb{E}}\left\{\Pi_{j1}|a_{j0}<z_{j0},s_{ij0}=z_{j0},z_{j0}\right\} & a_{j0}\geq\omega_a z_{j0}\end{cases},\tag{1.45}$$

where $\omega_a=\frac{\beta_0(1+\sigma_\varepsilon^{-2})}{\sigma_a^{-2}+\beta_0(1+\sigma_\varepsilon^{-2})}$ and the post-tax dividend payment is

$$\hat{\Pi}(a_{j0},z_{j0})=\tau\left(a_{j0},z_{j0}\right)\Pi_{j1}+T\left(a_{j0},z_{j0}\right).\tag{1.46}$$

The tax is confiscatory if the realization of the productivity shock a_{j0} is below the mean expectation of the marginal trader $\omega_a z_{j0}$, i.e., the firm disappoints market expectations. The expected tax revenue from the perspective of the marginal trader is transferred to buyers if the realization of a_{j0} is above $\omega_a z_{j0}$, i.e., the firm exceeds market expectations. A tax schedule that incentivizes information production, therefore, increases both the potential downsides and upsides of any trade. The before- and after-tax dividend schedule is shown in Figure 1.15 for the case with $z_{j0}=0$.

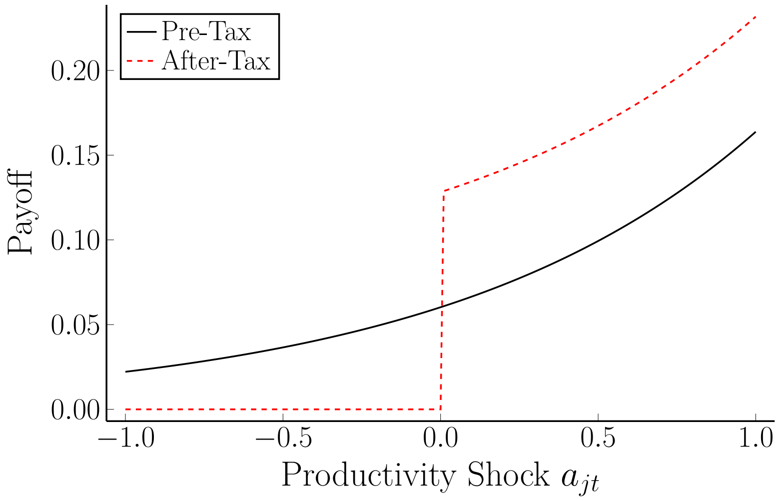

Figure 1.15: Information Production Incentivizing Tax Schedule.

Information production can be discouraged by flattening the dividend function instead. A straightforward and common implementation is through a progressive dividend tax in combination with the deduction of losses from realized gains, effectively offsetting part of the incurred losses by reducing the tax owed. In the model, the social planner can completely crowd out information production by buying all shares and selling shares that are claims on aggregate output. As there is no aggregate uncertainty, such shares pay a deterministic dividend and traders do not produce information.

To recapitulate, the social planner generally chooses a level of information production that deviates from the competitive equilibrium. If the efficient allocation of capital is sufficiently important, e.g., due to a high elasticity of substitution, then the social planner chooses a higher level of information production than would arise in the competitive equilibrium. Moreover, the social planner increases information production in response to both negative and positive sentiment shocks. In contrast, information production increases with the productivity shock. Finally, taxes and subsidies that increase the exposure to risk stemming from firm productivity increase the incentive to produce information.

1.7 Discussion

1.7.1 Asset Purchases

During the last decade, central banks have repeatedly used asset purchases to stabilize financial markets and accelerate economic growth and price inflation (for a brief overview, see Gagnon and Sack 2018). These interventions were accompanied by concerns that asset purchases might harm market efficiency and lead to an increase in capital misallocation.[26] Although my model is too stylized to give a full assessment of asset purchases, it can be used to shed light on the effect of asset purchases on information production in financial markets.[27]

In my model, asset purchases have real effects by exploiting that information is dispersed between traders. The mechanism works as follows: Asset purchases reduce the number of shares in the hands of traders, leading to an upward shift in the identity of the marginal trader. The marginal traders turn out to be more optimistic than in absence of asset purchases, and consequently, asset prices increase. Additionally, announced asset purchases affect information production. Traders anticipate the reduction in asset supply distorts prices upward, discouraging information production similar to a positive sentiment shock. Therefore, my model can provide a rationale for the concerns about asset purchases and declines in market efficiency.

However, asset purchases can also be used to reduce distortions in financial markets, for example, through negative sentiment shocks. When a sufficiently large negative sentiment shock hits the economy, traders anticipate that prices will be depressed, which discourages information production as trading becomes less information-sensitive. The central bank can offset the downward bias on asset prices by purchasing assets. This counter-measure can lead to unbiased prices, which restore the incentive to produce information for traders at the same time as increasing asset prices. This logic is captured

[26]See da Silva and Rungcharoenkitkul (2017), DNB (2017), Fernandez, Bortz, and Zeolla (2018), Borio and Zabai (2018), Acharya et al. (2019), and Kurtzman and Zeke (2020). In particular, the Dutch central bank argues in their 2016 annual report (DNB 2017): *"The large-scale purchase programmes and the flood of liquid assets has set the risk compass in financial markets spinning, with misallocations as a result."*

[27]Although most central banks focused on buying government bonds as a form of quantitative easing, also interventions in corporate bond markets were common, which can be interpreted through the lens of my model (Gagnon and Sack 2018). Moreover, the Bank of Japan bought directly shares in stock ETFs (Okimoto 2019).

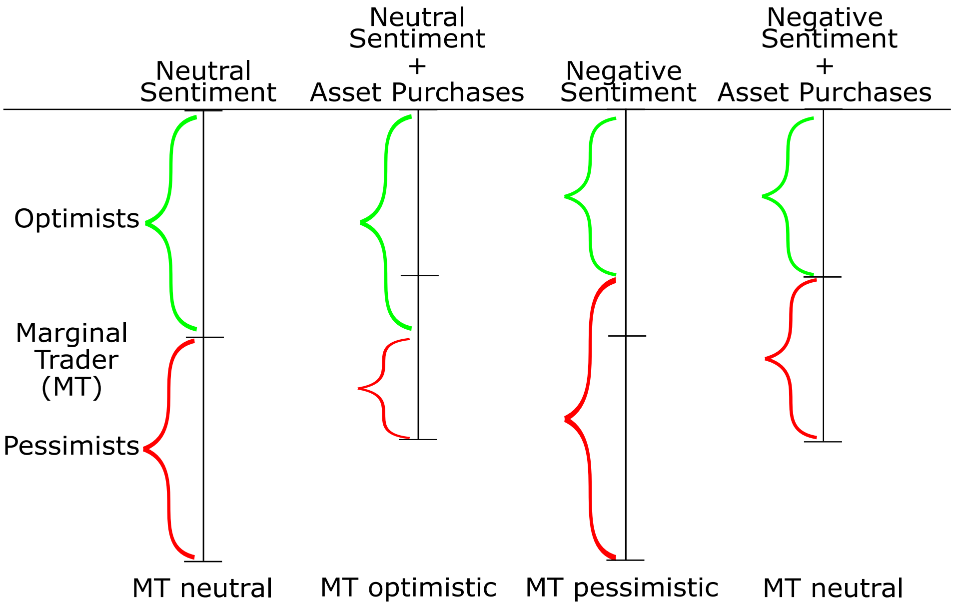

Figure 1.16: Asset Purchases Counter Negative Sentiment Shocks.

in the following proposition and is visualized in Figure 1.16.

Proposition 1.9. *Let the social planner acquire* $d^{SP} \in (-1, 1)$ *units of assets, such that* $1 - d^{SP}$ *shares are left for traders. Then,*

(i) asset purchases ($d^{SP} > 0$) *undo negative sentiment shocks both in terms of investment and information production.*

(ii) asset sales ($d^{SP} < 0$) *undo positive sentiment shocks both in terms of investment and information production.*

In other words, asset purchases and sales can *increase* market efficiency by countering sentiment shocks. This finding is relevant for central banks in deciding when to start shrinking the size of their balance sheets. Central banks can avoid the adverse effects of asset sales by waiting until sentiment has reached a more neutral level. A reduction in asset holdings can then even increase information production and market efficiency.

1.7.2 Uncertainty

Traders have Imperfect Information about Aggregate Shocks

The analysis so far assumed that traders observed aggregate states perfectly before deciding on information precision. This assumption is not crucial for the results, which also hold when traders have only imperfect information about aggregate states before they make their information production decision. Nonetheless, in reality, traders or policymakers do not have perfect knowledge about the current aggregate state.

The simplest setting to think about the effects of uncertainty is to reveal aggregate shocks after the information production decision but before trading. Furthermore, assume that aggregate productivity and sentiment shocks are auto-correlated.[28] Then, the laws of motion for aggregate shocks are given by

$$a_t = \rho_a a_{t-1} + \xi_t^a \tag{1.47}$$

$$\varepsilon_t = \rho_\varepsilon \varepsilon_{t-1} + \xi_t^\varepsilon, \tag{1.48}$$

where $\rho_a \in (0,1)$ and $\rho_\varepsilon \in (0,1)$ capture the persistence of aggregate shocks and $\xi_t^a \sim \mathcal{N}\left(0, \sigma_{\xi a}^2\right)$ and $\xi_t^\varepsilon \sim \mathcal{N}\left(0, \sigma_{\xi\varepsilon}^2\right)$ are the corresponding innovations. Traders can learn about past aggregate states by observing past aggregate investment K_t and output Y_t. Whereas K_t is moved by both productivity and sentiment, output Y_t reacts only to productivity after controlling for K_t^α. For example, if investment was high, but output was disappointing, investment must have been driven by a positive sentiment shock. The prior for traders about aggregate states is then given by

$$a_t | a_{t-1} \sim \mathcal{N}\left(\rho_a a_{t-1}, \sigma_{\xi a}^2\right) \tag{1.49}$$

$$\varepsilon_t | \varepsilon_{t-1} \sim \mathcal{N}\left(\rho_\varepsilon \varepsilon_{t-1}, \sigma_{\xi\varepsilon}^2\right). \tag{1.50}$$

In this setting, past sentiment shocks generate expectations about future sentiment shocks. The analysis of Proposition 1.4 still applies, as traders evaluate the value of information for different realizations of the sentiment shock ε_t.

[28] An alternative would be not to reveal aggregate shocks before trading takes place. In this setting traders learn from private and public signals also about aggregate states. Similarly, the social planner can use publicly available information to guide her interventions. The insights are broadly the same as in the case when aggregate shocks are revealed after the information production decision.

Policy under Uncertainty

The policy analysis is not substantially changed under aggregate uncertainty if the social planner has to take her decision before aggregate shocks are revealed. Indeed, negative effects of sentiment shocks on information production can be offset without knowing the exact realization, as only *expectations* of sentiment shocks affect information production. The social planner can collect information about traders' expectations of sentiment through surveys. This information can be used to implement a policy that offsets the effect of anticipated sentiment shocks on information production.

The effect of uncertainty is more subtle when the social planner also tries to steer investment. In this case, realizations of sentiment shocks matter. Therefore, any intervention that does not explicitly condition on the realizations of sentiment shocks has to weigh the costs and benefits of taxes or subsidies on investment in different states. Increasing information production can diminish the impact of sentiment shocks for all realizations.

A special case arises when traders are informed about aggregate shocks, but the social planner is not. In this case, multiple indicators can be used by the social planner to identify whether a boom is driven by sentiment or productivity. A sentiment-driven boom crowds out information production and decreases the variance of prices, leading all firms to look more alike. In contrast, a productivity-driven boom crowds in information, leading to more dispersion in asset prices and firm capital. For example, if asset prices increase across the board and the dispersion in asset prices or returns between firms shrinks, the social planner wants to lower investment and increase information production. Instead, if there are winners and losers even as asset prices are booming, price discovery still occurs, and traders are producing information. Using dispersion in asset prices and returns is more attractive than measuring information production directly, as to whether asset prices reflect fundamentals can only be backed out after production happened. However, asset prices are available continuously. This result is captured in the following proposition.

Proposition 1.10. *For* $\sigma_\varepsilon^2 \leq 1$,

(i) the cross-sectional variance in asset prices is increasing in β_t.

(ii) the cross-sectional variance in asset price returns is increasing in β_t.

Finally, if policymakers need to commit to interventions before prices form and aggregate shocks are persistent, past realizations of price-earnings ratios can also be informative regarding future aggregate shocks. For example, if investment was high, but output was relatively low, then investment must have been driven by sentiment, and future investment is also likely to be driven by sentiment.

1.7.3 Empirical Evidence

Many measures seek to capture a notion of information in financial markets. However, the literature so far has not converged on any single measure. Roll (1988) suggested a measure that attributes movements in asset prices that are uncorrelated with the market or industry portfolio with new firm-specific information. However, firm-specific variance can also stem from firm-specific noise (for an overview, see Ningning and Hongquan 2014). Chousakos, Gorton, and Ordoñez (2020) employed a measure that follows a similar idea.

In contrast, Bai, Philippon, and Savov (2016) and Farboodi et al. (2020) suggested a measure that uses asset prices to forecast earnings. According to this measure, financial markets are informative if firms with higher earnings also have a higher market capitalization. The downside of their approach is that it implicitly assumes that the data generating processes for earnings and prices are identical between firms, as they run regressions for cross-sections of firms.

Dávila and Parlatore (2021) avoided these objections by providing a micro-founded procedure to estimate (relative) price informativeness at the firm level, allowing for different data generating processes for each firm. Relative price informativeness captures a notion of how precise the price signal is relative to prior uncertainty, which corresponds exactly to $\frac{\beta_{jt}\sigma_\varepsilon^{-2}}{\sigma_a^{-2}+\beta_{jt}\sigma_\varepsilon^{-2}}$ in my model. Their measure can be used to provide suggestive evidence that information precision indeed depends on the cycle. I use an estimate of utilization-adjusted TFP following Basu, Fernald, and Kimball (2006) from the San Francisco Fed to verify the connection between information and aggregate productivity and as an indicator for the type of shock that drives the cycle.

Using data from the US between 1995 and 2017, Figure 1.1 provides suggestive evidence that information in financial markets varies depending on what type of shock drives the cycle. Because the model focuses on cycles instead of long-run developments, both time series are detrended using a cubic time trend between 1995 and 2017 and smoothed with a two-year moving average. The resulting time series is shown in the left graph, whereas the original can be seen on the right. Both graphs have gray bars that indicate recessions following the methodology of the NBER for dating recessions. The first striking observation is that the cyclical components of price informativeness and TFP growth are positively related. As so far as cyclical movements in TFP growth capture changes to allocational efficiency, this provides evidence that information in financial markets indeed impacts TFP.

A second exercise allows us to back out which type of shocks drove the expansions up to 2001 and from 2002 to 2008. The period between 1995 and 2001 was marked by an acceleration in TFP growth, accompanied by an increase in price informativeness. This increase suggests that the dot-com boom was driven by technological innovations, for example, the introduction of advanced information technologies. In contrast, the expansion between 2002 and 2008 was marked by a sharp decline in TFP growth into negative territory and a fall in price informativeness. Through the lens of the model, an expansion accompanied by a decline in TFP signals a sentiment boom (see also Borio et al. 2015; Doerr 2018). The finding that price informativeness was also declining verifies the model's prediction that information production declines during sentiment booms.

This narrative is also supported when using return non-synchronicity as a measure of price informativeness as Roll (1988). I use the database *CRSP* to compute the standard deviation in monthly stock returns for the US. Following standard practices, I drop the financial sector with four digit SICC codes 6xxx and firms with market caps in the bottom 20 percent, for which I take breaking points from Kenneth R. French's website. Similarly, I only include ordinary common shares (share code ten and eleven) which are traded on the NYSE, NYSE American and NASDAQ (exchange code one, two, and three).

Similarly to price informativeness as in Figure 1.1, the dispersion of monthly returns increases during the dot-com boom and decreases during the subsequent housing boom, as can be seen in Figure 1.17. Viewed through the lens of the model, this suggests that the dot-com boom has been driven by productivity, whereas the housing boom has been driven by sentiment. Different to Figure 1.2, the dispersion in monthly returns stays relatively low after the Great Financial Crisis, whereas price informativeness following Dávila and Parlatore (2021) decreases but stays historically high. A possible explanation is that return dispersion also captures changes in the variance of fundamentals and

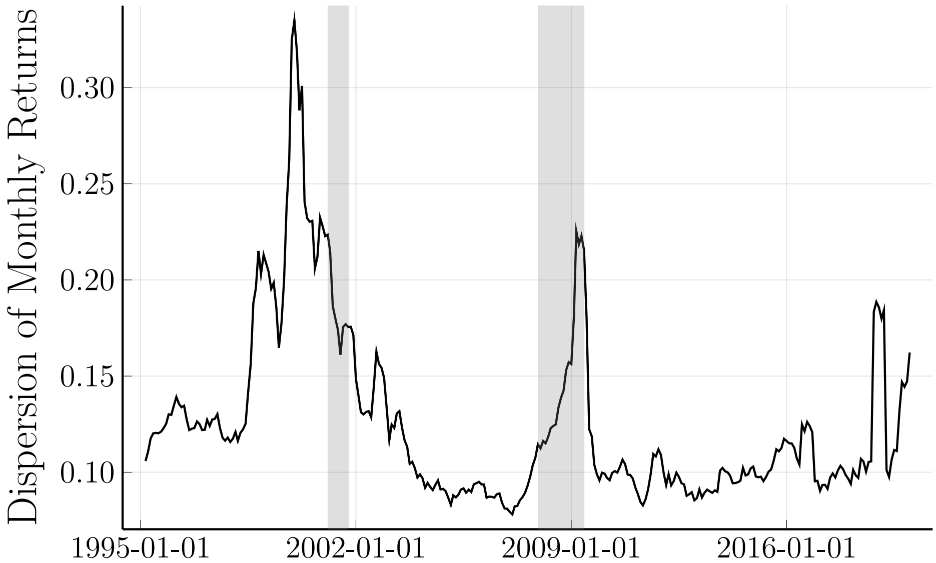

Figure 1.17: Return Dispersion.

Notes: Return dispersion was high during the dot-com boom leading up to 2001, but decreased substantially during the housing boom leading up to the Great Financial Crisis.

noise, whereas the measure of Dávila and Parlatore (2021) aims to correct for changes in variances. Therefore, volatility may have remained low due to a decrease in variances, whereas informativeness remained high as information production decreased by less.

1.8 Conclusion

I develop a tractable framework to study information production in financial markets embedded in a standard macroeconomic model. In such a model, total factor productivity has an endogenous component that depends on the traders' decentralized information production. When asset prices are more informative, more capital is allocated to the most productive firms, and total factor productivity increases. I add to the literature by studying the effect of aggregate shocks on information production.

I prove that sentiment shocks, defined as waves of non-fundamental optimism or pessimism, crowd out information production as trading becomes less information-sensitive. Although such optimism increases investment, it also worsens the allocation of capital. This result rationalizes the empirical finding that credit booms often worsen aggregate productivity (Borio et al. 2015; Gopinath et al. 2017; Doerr 2018; Gorton and Ordoñez 2020) through a novel information mechanism. In contrast, expectations of heightened productivity crowd in information, thereby improving capital allocation and aggregate productivity beyond the initial shock. This dichotomy mirrors the “good” and “bad” booms of Gorton and Ordoñez (2020). My model suggests that “good” booms are driven by productivity, whereas “bad” booms are driven by sentiment.

From a normative perspective, I show that information production is too high or too low in the competitive equilibrium. There are two externalities with opposing effects. On the one hand, traders produce information to increase trading rents at the expense of other traders. This rent-extracting behavior can lead to excessive information production. On the other hand, traders do not reap the benefits of improving the capital allocation through *collective* information production. This information spillover can lead to information production being too low. Generally, information production is too low in the competitive equilibrium exactly when the allocation of capital matters the most and, hence, information is most valuable.

Finally, I apply the model to evaluate the effect of large-scale asset purchase programs. I show that asset purchases can discourage information production. This finding confirms the concerns of policymakers about such programs (e.g., DNB 2017). However, asset purchases can also improve capital allocation if they effectively reduce aggregate mispricing of assets. Therefore, policymakers need to know which force is currently driving the cycle to react appropriately. The model suggests that dispersion in asset prices and returns identify the source of fluctuations in real-time. For example, sentiment booms decrease information production, which lowers the dispersion in asset prices and returns and lets firms appear more alike.

Appendix

1.A Trading

Every household i consists of many traders indexed by $ij \in [0,1]$. The information set of each trader consists of $\{s_{ijt}, \{z_{jt}\}, a_t, \varepsilon_t\}$, i.e., traders observe their private signal, all public signals and the aggregate states. This setting allows that traders have rational expectations about aggregates, but still disagree about firm-specific variables, which motivates trade. I impose that $\kappa_H = 2$ and $\kappa_L = 0$ to avoid distortions in asset prices that stem from the choice of position limits.

The beliefs of traders about firm productivity A_{jt} are relevant for their trading decision. Trader ij's beliefs are given by

$$\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}, z_{jt}\right\} = \exp\left\{\omega_{p,ijt}a_t + \omega_{s,ijt}s_{ijt} + \omega_{z,ijt}\left(z_{jt} - \frac{\varepsilon_t}{\sqrt{\beta_{jt}}}\right) + \frac{1}{2}\mathbb{V}_{ijt}\right\}. \tag{1.51}$$

Similarly, the beliefs of the marginal trader are

$$\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\} = \exp\left\{\omega_{p,jt}a_t + \omega_{s,jt}z_{jt} + \omega_{z,jt}\left(z_{jt} - \frac{\varepsilon_t}{\sqrt{\beta_{jt}}}\right) + \frac{1}{2}\mathbb{V}_{jt}\right\}, \tag{1.52}$$

where $a_{jt} \overset{iid}{\sim} \mathcal{N}\left(a_t, \sigma_a^2\right)$, $\varepsilon_{jt} \overset{iid}{\sim} \mathcal{N}\left(\varepsilon_t, \sigma_\varepsilon^2\right)$ and ω-terms are the corresponding Bayesian weights,

$$\omega_{z,ijt} = \frac{\beta_{jt}\sigma_\varepsilon^{-2}}{\sigma_a^{-2} + \beta_{ijt} + \beta_{jt}\sigma_\varepsilon^{-2}}, \quad \omega_{z,jt} = \frac{\beta_{jt}\sigma_\varepsilon^{-2}}{\sigma_a^{-2} + \beta_{jt} + \beta_{jt}\sigma_\varepsilon^{-2}} \tag{1.53}$$

$$\omega_{s,ijt} = \frac{\beta_{ijt}}{\sigma_a^{-2} + \beta_{ijt} + \beta_{jt}\sigma_\varepsilon^{-2}}, \quad \omega_{s,jt} = \frac{\beta_{jt}}{\sigma_a^{-2} + \beta_{jt} + \beta_{jt}\sigma_\varepsilon^{-2}} \tag{1.54}$$

$$\omega_{p,ijt} = \frac{\sigma_a^{-2}}{\sigma_a^{-2} + \beta_{ijt} + \beta_{jt}\sigma_\varepsilon^{-2}}, \quad \omega_{p,jt} = \frac{\sigma_a^{-2}}{\sigma_a^{-2} + \beta_{jt} + \beta_{jt}\sigma_\varepsilon^{-2}}, \tag{1.55}$$

and $\{\mathbb{V}_{jt}, \mathbb{V}_{ijt}\}$ stand for posterior uncertainty

$$\mathbb{V}_{ijt} = \frac{1}{\sigma_a^{-2} + \beta_{ijt} + \beta_{jt}\sigma_\varepsilon^{-2}}, \quad \mathbb{V}_{jt} = \frac{1}{\sigma_a^{-2} + \beta_{jt} + \beta_{jt}\sigma_\varepsilon^{-2}}. \tag{1.56}$$

The private information precision β_{ijt} is highlighted in blue and is part of the information acquisition decision. Alternatively, the beliefs of the marginal trader who observed $s_{ijt} = z_{jt}$ can be expressed as a function of shocks,

$$\ln \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\} = \omega_{p,jt}a_t + \omega_{s\varepsilon,jt}\varepsilon_t + \omega_{a,jt}a_{jt} + \frac{\omega_{a,jt}}{\sqrt{\beta_{jt}}}\left(\varepsilon_{jt} - \varepsilon_t\right) + \frac{1}{2}\mathbb{V}_{jt}, \tag{1.57}$$

where the corresponding Bayesian weights are

$$\omega_{a,jt} = \omega_{z,jt} + \omega_{s,jt} \tag{1.58}$$

$$\omega_{\varepsilon,jt} = \omega_{a,jt}/\sqrt{\beta_{jt}} \tag{1.59}$$

$$\omega_{s\varepsilon,jt} = \omega_{s,jt}/\sqrt{\beta_{jt}}. \tag{1.60}$$

Trader ij buys shares of firm j whenever

$$\tilde{\mathbb{E}}\left\{\Pi_{jt+1}|s_{ijt}, z_{jt}\right\} \geq \tilde{\mathbb{E}}\left\{\Pi_{jt+1}|s_{ijt} = z_{jt}, z_{jt}\right\} \tag{1.61}$$

$$\iff \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}, z_{jt}\right\} \geq \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\}. \tag{1.62}$$

The above inequality is equivalent to

$$\omega_{p,ijt}a_t + \omega_{s,ijt}s_{ijt} + \omega_{z,ijt}\left(z_{jt} - \frac{1}{\sqrt{\beta_{jt}}}\varepsilon_t\right) + \frac{\mathbb{V}_{ijt}}{2} \geq \tag{1.63}$$

$$\omega_{p,jt}a_t + \omega_{s,jt}z_{jt} + \omega_{z,jt}\left(z_{jt} - \frac{1}{\sqrt{\beta_{jt}}}\varepsilon_t\right) + \frac{\mathbb{V}_{jt}}{2}. \tag{1.64}$$

The inequality can be expressed as a cutoff for the idiosyncratic noise,

$$\eta_{ijt} \geq \frac{\sqrt{\beta_{ijt}}}{\omega_{s,ijt}}\left(\omega_{p,ijt}a_t + \omega_{z,ijt}\left(z_{jt} - \frac{1}{\sqrt{\beta_{jt}}}\varepsilon_t\right) + \frac{\mathbb{V}_{ijt}}{2}\right) + \sqrt{\beta_{ijt}}a_{jt} \tag{1.65}$$
$$- \frac{\sqrt{\beta_{ijt}}}{\omega_{s,ijt}}\left(\omega_{p,jt}a_t + \omega_{s,jt}z_{jt} + \omega_{z,jt}\left(z_{jt} - \frac{1}{\sqrt{\beta_{jt}}}\varepsilon_t\right) + \frac{\mathbb{V}_{jt}}{2}\right).$$

Since η_{ijt} is standard-normally distributed, the perceived probability of buying can be written in closed form

$$\begin{aligned}
&\mathcal{P}\left\{x_{ijt} = 2|a_{jt}, \varepsilon_{jt}, \beta_{ijt}, \beta_{jt}\right\} \\
&= \Phi\left(-\frac{\sqrt{\beta_{ijt}}}{\omega_{s,ijt}}\left(\omega_{p,ijt}a_t + \omega_{z,ijt}\left(z_{jt} - \frac{1}{\sqrt{\beta_{jt}}}\varepsilon_t\right) + \frac{1}{2}\mathbb{V}_{ijt}\right) + \sqrt{\beta_{ijt}}a_{jt}\right. \\
&\left. - \frac{\sqrt{\beta_{ijt}}}{\omega_{s,ijt}}\left(\omega_{p,jt}a_t + \omega_{s,jt}z_{jt} + \omega_{z,jt}\left(z_{jt} - \frac{1}{\sqrt{\beta_{jt}}}\varepsilon_t\right) + \frac{1}{2}\mathbb{V}_{jt}\right)\right),
\end{aligned} \tag{1.66}$$

where $\Phi(\cdot)$ is the standard-normal cdf. For a symmetric information choice ($\beta_{ijt} = \beta_{jt}$), the buying probability can be simplified to

$$\mathcal{P}\left\{x_{ijt} = 2|a_{jt}, \varepsilon_{jt}, \beta_{ijt}, \beta_{jt}\right\}\big|_{\beta_{ijt}=\beta_{jt}} = \Phi\left(-\varepsilon_{jt}\right). \tag{1.67}$$

Traders think that they are more likely to buy shares when the realization of the sentiment shock is relatively low and shares are therefore cheap relative to their fundamental value.

Finally, traders choose their information precision taking the symmetric choice of all other traders as given. The derivative of the probability of buying with respect to β_{ijt}

is

$$\frac{\partial \mathcal{P}\left\{x_{ijt}=2|a_{jt},\varepsilon_{jt},\beta_{ijt},\beta_{jt}\right\}}{\partial \beta_{ijt}} \tag{1.68}$$

$$\begin{aligned}
&= \phi\left(-\frac{\sqrt{\beta_{ijt}}}{\omega_{s,ijt}}\left(\omega_{p,ijt}a_t+\omega_{z,ijt}\left(z_{jt}-\frac{1}{\sqrt{\beta_{jt}}}\varepsilon_t\right)+\frac{1}{2}\mathbb{V}_{ijt}\right)+\sqrt{\beta_{ijt}}a_{jt}\right.\\
&\left.-\frac{\sqrt{\beta_{ijt}}}{\omega_{s,ijt}}\left(\omega_{p,jt}a_t+\omega_{s,jt}z_{jt}+\omega_{z,jt}\left(z_{jt}-\frac{1}{\sqrt{\beta_{jt}}}\varepsilon_t\right)+\frac{1}{2}\mathbb{V}_{jt}\right)\right)\\
&*\left(-\frac{1}{2\beta_{ijt}^{3/2}}\left(\sigma_a^{-2}a_t+\beta_{jt}\sigma_\varepsilon^{-2}\left(z_{jt}-\frac{\varepsilon_t}{\sqrt{\beta_{jt}}}\right)+\frac{1}{2}\right)+\frac{a_{jt}}{2\sqrt{\beta_{ijt}}}\right.\\
&-\left(\frac{1}{\sqrt{\beta_{ijt}}}-\frac{1}{2\beta_{ijt}^{3/2}}\left(\mathbb{V}_{ijt}\right)^{-1}\right)\\
&\left.*\left(\omega_{p,jt}a_t+\omega_{s,jt}z_{jt}+\omega_{z,jt}\left(z_{jt}-\frac{\varepsilon_t}{\sqrt{\beta_{jt}}}\right)+\frac{1}{2}\mathbb{V}_{jt}\right)\right)
\end{aligned} \tag{1.69}$$

where $\phi(\cdot)$ is the standard normal pdf. For a symmetric information choice ($\beta_{ijt}=\beta_{jt}$) this expression can be simplified to

$$\left.\frac{\partial \mathcal{P}\left\{x_{ijt}=2|a_{jt},\varepsilon_{jt},\beta_{ijt},\beta_{jt}\right\}}{\partial \beta_{ijt}}\right|_{\beta_{ijt}=\beta_{jt}} =$$
$$\phi\left(\varepsilon_{jt}\right)\left[\frac{1}{2\sqrt{\beta_{jt}}}\left(a_{jt}+z_{jt}\right)-\frac{1}{\sqrt{\beta_{jt}}}\left(\omega_{p,jt}a_t+\omega_{s,jt}z_{jt}+\omega_{z,jt}\left(z_{jt}-\frac{\varepsilon_t}{\sqrt{\beta_{jt}}}\right)+\frac{1}{2}\mathbb{V}_{jt}\right)\right]. \tag{1.70}$$

1.B Position Limits

1.B.1 Exogenous Position Limits

In the main text, I have assumed that traders can buy up to two units of each stock and normalized the total asset supply to one. Assume now that traders' position limits are given by $x_{ijt} \in [0, \kappa_H]$. Consider first some special cases.

Let $\kappa_H \in [0, 1)$. In this case, demand by traders is insufficient to buy the total asset supply. The result is that the stock price collapses to zero, all traders buy κ_H units of firm j's stock, and the price is uninformative because it does not vary according to firm productivity. Similarly, if $\kappa_H = 1$, traders can clear the market, but the same outcome arises.

In contrast, if there are no upper limits to how much traders can buy ($\kappa_H = \infty$), the most optimistic trader alone can clear the whole market. Expectations about dividends and the interest rate R_{t+1} go to infinity, but prices are remain finite. Information becomes useless for traders because the probability of buying in any given market is zero.

To avoid these edge cases, I focus on position limits for which the market clearing condition gives an interior solution for the threshold, i.e., $\kappa_H \in (1, \infty)$. The market-clearing condition

$$\kappa_H \left(1 - \Phi\left(\sqrt{\beta_{jt}}\left(\hat{s}\left(P_{jt}\right) - a_{jt}\right) - \varepsilon_{jt}\right)\right) = 1 \tag{1.71}$$

leads to the threshold

$$\hat{s}\left(P_{jt}\right) = a_{jt} + \frac{\varepsilon_{jt} + \Phi^{-1}\left(1 - \frac{1}{\kappa_H}\right)}{\sqrt{\beta_{jt}}}. \tag{1.72}$$

The resulting expectations of dividends can be expressed by multiplying the price when traders can buy up to two units with a factor related to κ_H,

$$\tilde{\mathbb{E}}\left\{\Pi_{jt+1} | s_{ijt} = \hat{s}\left(P_{jt}\right), z_{jt}\right\} = \underbrace{\exp\left\{\Phi^{-1}\left(1 - \frac{1}{\kappa_H}\right) / \sqrt{\beta_{jt}}\right\}^{\theta}}_{\textit{Bias through Choice of Position Limits}} \tilde{\mathbb{E}}\left\{\Pi_{jt+1} | s_{ijt} = z_{jt}, z_{jt}\right\}. \tag{1.73}$$

As is now evident, the bias is equal to zero for $\kappa_H = 2$. In other words, the marginal trader is neither an optimist nor pessimist. Consequently, the interest rate is also distorted,

$$R_{t+1} = \exp\left\{\Phi^{-1}\left(1 - \frac{1}{\kappa_H}\right) / \sqrt{\beta_{jt}}\right\}^{\theta} \frac{\int_0^1 \tilde{\mathbb{E}}\left\{\Pi_{jt+1} | s_{ijt} = z_{jt}, z_{jt}\right\} dj}{K_{t+1}}. \tag{1.74}$$

Holding K_{t+1} constant leads to an allocation of capital that is independent of the position limit κ_H,

$$K_{jt+1} = \frac{\tilde{\mathbb{E}}\left\{\Pi_{jt+1}|s_{ijt} = \hat{s}\left(P_{jt}\right), z_{jt}\right\}}{R_{t+1}}$$
$$= \frac{\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\}^{\theta}}{\int_0^1 \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\}^{\theta} dj} K_{t+1}. \quad (1.75)$$

Still, a different interest rate will affect aggregate investment through (1.30). If buyers are relatively optimistic ($\kappa_H > 2$), then the interest rate and aggregate investment increase. Setting $\kappa_H = 2$ is for the model with exogenously given information inconsequential and only avoids introducing a multiplicative factor for expectations of dividends.

For the information production decision, the choice of position limits has similar effects as aggregate sentiment shocks or reductions in asset supply. The main idea is the same: when the aggregate sentiment shock is positive, traders expect the trading decision to become less information-sensitive, making information less valuable. The same effect is present when setting a higher position limit $\kappa_H > 2$. In this case, however, it is counteracted by an increase in the traders' buying capacity. Depending on which effect dominates, the highest information choice is achieved for $\kappa_H < 2$ or $\kappa_H > 2$.

Position limits affect the analysis for sentiment shocks and revert the logic outlined in the main text. For example, assume that $\kappa_H = 1 + \eta$, where $\eta > 0$ is a small number. Then almost all traders need to buy shares to clear the market. It follows that trading is fully information-insensitive because all traders expect to buy κ_H units of nearly all firms irrespective of the private signal. Different from the intuition before, a positive sentiment shock makes traders think that the trading decision becomes *more information-sensitive*. Recall that the trading decision is most information-sensitive if the ex-ante probability of buying is 50%. As the increase in the sentiment shock pushes the ex-ante probability of trading towards 50% from below, a sentiment shock can make the trading decision more information-sensitive and encourage information production.

The choice of $\kappa_H = 2$ in the main text guarantees that the marginal trader is, on average, neither optimistic nor pessimistic absent aggregate sentiment shocks. Moreover, considering aspects outside of the model, excess or lack of demand can lead to the entry or exit of traders because firms are predictably under- or overpriced. It can also lead to the additional entry or exit of firms for the same reason. Both forces tend to undo the effects of too much or too little demand. Finally, denoting position limits in units of shares is mainly an analytical simplification when including risk-neutral traders in the financial markets. A formulation where position limits depend on the trader's wealth is investigated in Chapter 2.

1.B.2 Short-Sales

Short-sales were ruled out in the main text for analytical convenience, but its presence would not affect the main results of the model. Assume that traders can take also negative positions, such that $x_{ijt} \in [-2, 2]$. The market-clearing condition becomes

$$\underbrace{2\left(1-\Phi\left(\sqrt{\beta_{jt}}\left(\hat{s}\left(P_{jt}\right)-a_{jt}\right)-\varepsilon_{jt}\right)\right)}_{\textit{buying}}-\underbrace{2\Phi\left(\sqrt{\beta_{jt}}\left(\hat{s}\left(P_{jt}\right)-a_{jt}\right)-\varepsilon_{jt}\right)}_{\textit{selling}}=1, \tag{1.76}$$

with the threshold

$$\hat{s}\left(P_{jt}\right)=a_{jt}+\frac{\varepsilon_{jt}+\Phi^{-1}\left(\frac{1}{4}\right)}{\sqrt{\beta_{jt}}}. \tag{1.77}$$

In contrast to before, more traders need to buy to clear the market, because previously inactive traders now short shares and thereby increase the asset supply. Therefore, allowing short-sales leads to a lower price because the marginal trader will be more pessimistic than before. The bias can be avoided by imposing asymmetric position limits, e.g., $x_{ijt} \in [-2, 4]$, in which case the marginal trader still is identified by the signal $\hat{s}\left(P_{jt}\right)=a_{jt}+\varepsilon_{jt}/\sqrt{\beta_{jt}}$.

Similar to the previous section, the choice of position limits can change the effect of sentiment shocks on the incentive to produce information by introducing a bias to the asset price. This issue is avoided by ruling out short-sales and setting $\kappa_H = 2$.

1.B.3 Endogenous Position Limits

Finally, let traders choose position limits $x_{ijt} \in [\kappa_L, \kappa_H]$ subject to cost $c^L\left(\kappa_L\right)$ and $c^H\left(\kappa_H\right)$ before trading takes place. One interpretation is that funds and credit lines have to be allocated between markets, which can be costly. Such shifting of funds can be valuable if traders expect that some markets are under- or overpriced. For example, if market j is hit by a positive sentiment shock, traders may want to extend their ability to short shares in this market while reducing their ability to buy, by shifting collateral towards this market and cash towards other markets. Generally, endogenously choosing position limits will tend to counteract the effects of sentiment shocks.

The effect on the information production decision is more subtle. Consider as a partial equilibrium example that trader ij received private information that shares of firm j will be underpriced. In anticipation of a depressed market, trader ij extends her ability to buy but completely forgoes short-sales. The opportunity cost of buying when prices are too high is proportional to $\kappa_H - \kappa_L$. Therefore, the value of information is increasing in $\kappa_H - \kappa_L$. Whether the adjustment of position limits increases information production depends, therefore, on whether $\kappa_H - \kappa_L$ is increased as a result.

More formally, the expected trading rents can be written as

$$\widetilde{EU}\left(\beta_{ijt}, \beta_{jt}\right) = \tilde{\mathbb{E}}\left\{\left(\kappa_H \mathcal{P}\left\{x_{ijt} = \kappa_H\right\} + \kappa_L \mathcal{P}\left\{x_{ijt} = \kappa_L\right\}\right)\left(\frac{1}{R_{t+1}}\Pi_{jt+1} - P_j t\right)\right\}, \tag{1.78}$$

where profits are equal to the dividend earned minus the opportunity cost of the price paid. Because there are no trading costs, it must be that $\mathcal{P}\left\{x_{ijt} = \kappa_L\right\} = 1 - \mathcal{P}\left\{x_{ijt} = \kappa_H\right\}$:

$$\begin{aligned}\widetilde{EU}\left(\beta_{ijt}, \beta_{jt}\right) = \tilde{\mathbb{E}}&\left\{\left(\kappa_H - \kappa_L\right)\mathcal{P}\left\{x_{ijt} = \kappa_H\right\}\left(\frac{1}{R_{t+1}}\Pi_{jt+1} - P_j t\right)\right\} \\ &- \kappa_L \tilde{\mathbb{E}}\left\{\frac{1}{R_{t+1}}\Pi_{jt+1} - P_{jt}\right\}.\end{aligned} \tag{1.79}$$

Taking the derivative with respect to β_{ijt} yields

$$\widetilde{MB}\left(\beta_{ijt}, \beta_{jt}\right) = \tilde{\mathbb{E}}\left\{\left(\kappa_H - \kappa_L\right)\frac{\partial \mathcal{P}\left\{x_{ijt} = \kappa_H\right\}}{\partial \beta_{ijt}}\left(\frac{1}{R_{t+1}}\Pi_{jt+1} - P_{jt}\right)\right\}. \tag{1.80}$$

The marginal benefit of producing information is proportional to $\kappa_H - \kappa_L$. Therefore, if traders decide to expand $\kappa_H - \kappa_L$ in response to a shock, the endogenous choice of position limits will tend to increase information production.

1.C Intermediate Good Firms

1.C.1 Micro-Foundation

Intermediate good firms sell their whole revenue stream to traders to focus the analysis on information frictions. This assumption can be micro-founded by assuming that there are at least two entrepreneurs without private wealth for each variety j. Entrepreneurs need to turn to financial markets to fund their projects, but the market is competitive in the sense that, at most, one entrepreneur for each variety j can sell her shares to traders. A mechanism chooses the entrepreneur who promises the highest rate of return on her shares. If there is a tie, the successful entrepreneur is chosen randomly among the entrepreneurs who offer the highest return.

Formally, the entrepreneur's problem is

$$\max_{K_{jkt+1}, D_{jkt+1}\left(A_{jt}, K_{jkt+1}\right)} C_{jkt} + \delta \mathbb{E}\left\{C_{jkt+1}\right\} \tag{1.81}$$

$$s.t. \quad K_{jkt+1} + C_{jkt} \leq P_{jkt}. \tag{1.82}$$

$$C_{jkt+1} \leq \Pi_{jkt+1}\left(A_{jt}, K_{jkt+1}\right) - D_{jkt+1}\left(A_{jt}, K_{jkt+1}\right) \tag{1.83}$$

$$C_{jkt}, C_{jkt+1}, K_{jkt+1}, D_{jkt+1}\left(A_{jt}, K_{jkt+1}\right) \geq 0 \tag{1.84}$$

where

$$P_{jkt} = \begin{cases} 0 & \text{if } \exists k' \neq k : R_{jkt+1} < R_{jk't+1} \\ \begin{cases} 0 & \text{w.p. } 1 - \frac{1}{|k''|} \\ \frac{1}{R_{t+1}} \tilde{\mathbb{E}}\left\{D_{jkt+1}\left(A_{jt}, K_{jkt+1}\right) | s_{ijt} = z_{jt}, z_{jt}\right\} & \text{w.p. } \frac{1}{|k''|} \end{cases} & \text{if } \substack{\exists k'' \neq k : R_{jkt+1} = R_{jk''t+1} \\ \forall k' \neq k : R_{jkt+1} \geq R_{jk't+1}} \\ \frac{1}{R_{t+1}} \tilde{\mathbb{E}}\left\{D_{jkt+1}\left(A_{jt}, K_{jkt+1}\right) | s_{ijt} = z_{jt}, z_{jt}\right\} & \text{if } \forall k' \neq k : R_{jkt+1} > R_{jk't+1} \end{cases} . \tag{1.85}$$

The entrepreneur maximizes her utility over consumption today and tomorrow using the same utility function as households.[29] When young, entrepreneurs can either consume or invest in their firm. When old, entrepreneurs pay out a dividend D_{jkt+1} and consume what remains of revenue Π_{jkt+1}.

[29] Entrepreneurs can only raise funds by selling claims to revenue and cannot otherwise borrow or lend. This setting guarantees that asset prices are an invertible function of z_{jt}, a noisy signal of firm productivity, without which the equilibrium in the financial market does not exist. See Albagli, Hellwig, and Tsyvinski (2011b, 2017) for a discussion of this issue.

The entrepreneur is only able to sell her shares at a positive price if she offers the highest return in market j. If the entrepreneur promises a lower rate of return R_{jkt+1} than some other entrepreneur k', she will not be able to sell her shares and will raise nothing. If she promises the highest rate of return in the economy, but another entrepreneur promises the same return, she will be able to sell her shares with probability $1/|k''|$ where $|k''|$ is the number of entrepreneurs who also promise the highest return. If only she promises the highest return, she will be able to sell her shares with probability one. Finally, the expected rate of return is given by

$$R_{jkt+1} = \frac{\mathbb{E}\left\{D_{jkt+1}\left(A_{jt}, K_{jkt+1}\right)\right\}}{P_{jkt+1}}. \tag{1.86}$$

There is perfect competition between entrepreneurs because productivity A_{jt} is attached to variety j. Therefore, all entrepreneurs sell their goods at the same price ρ_{jt+1} . The only equilibrium is one in which at least two entrepreneurs choose

$$D_{jkt+1}\left(A_{jt}, K_{jkt+1}\right) = \Pi_{jkt+1}\left(A_{jt}, K_{jkt+1}\right) \tag{1.87}$$

$$K_{jkt+1} = P_{jkt}. \tag{1.88}$$

It show first that the choice above is an equilibrium and then I show that it is the only equilibrium. Any entrepreneur k who chooses (1.87) and (1.88) can only deviate by either investing less or paying a lower dividend. In either case, another entrepreneur exists who promises a higher return on investment and entrepreneur k is unable to sell her shares. Similarly, any entrepreneur who does not choose (1.87) and (1.88) does not have a profitable deviation either. Choosing to invest less or promising a lower dividend leads to no change, as the rate of return is only further depressed. Investing more or promising a higher dividend is similarly inconsequential as long as the entrepreneur does not choose at least (1.87) and (1.88). If she chooses to deviate to (1.87) and (1.88), she still earns zero profits but gets to produce with positive probability. Therefore, (1.87) and (1.88) are an equilibrium.

To show that at least two entrepreneurs choosing (1.87) and (1.88) is the only equilibrium, it is necessary to show that profitable deviations exist for all other choices of investment and dividends. First, consider that only one entrepreneur k chooses (1.87) and (1.88) and that all others either invest strictly less or pay a lower dividend in some states. Then entrepreneur k can raise her profits by either investing less or paying a lower dividend while still promising the highest rate of return. Second, assume that all entrepreneurs choose an investment and dividend policy that leads to positive profits in at least some states. In this case, there is a profitable deviation for any entrepreneur. Entrepreneur k can invest more or pay a larger dividend to promise the highest rate of return while still keeping positive profits. Therefore, the only equilibrium is given by at least two entrepreneurs choosing (1.87) and (1.88).

1.C.2 Entrepreneurs with Market Power: Equity

Alternatively, assume that there is only one entrepreneur per variety. Furthermore, assume that entrepreneurs are patient and restricted to selling equity contracts as captured in the following Assumption.

Assumption 1.3 (*Equity Contracts*). *Entrepreneurs can only sell claims to a fraction* $\lambda\left(P_{jt}, P_t\right) \in [0, 1]$ *of firm revenue.*

The share of revenue that is sold to the market is allowed to depend on the price P_{jt} and on the aggregate value of the stock market P_t. The entrepreneur's maximization problem is

$$\max_{\lambda(P_{jt},P_t),K_{jt+1}} \quad \mathbb{E}\left\{C_{jt+1}|P_{jt}\right\} \tag{1.89}$$

$$C_{jt+1} \leq \Pi\left(A_{jt}, K_{jt+1}\right) - D\left(A_{jt}, K_{jt+1}\right) \tag{1.90}$$

$$D\left(A_{jt}, K_{jt+1}\right) = \lambda\left(P_{jt}, P_t\right)\Pi\left(A_{jt}, K_{jt+1}\right) \tag{1.91}$$

$$0 \leq K_{jt+1} \leq P_{jt}. \tag{1.92}$$

The entrepreneur maximizes her old age consumption, consisting of firm revenue $\Pi\left(A_{jt}, K_{jt+1}\right)$ after paying dividends $D\left(A_{jt}, K_{jt+1}\right)$ subject to constraints (1.90), (1.91) and (1.92). The first constraint states that consumption cannot be negative. The second constraint follows from Assumption 1.3. The final constraint imposes non-negativity on investment and states that entrepreneurs cannot borrow additional funds from other sources.

Plugging in the constraints into the objective yields the simplified problem

$$\max_{\lambda(P_{jt},P_t),K_{jt+1}} \quad \mathbb{E}\left\{\Pi_{jt+1} - \lambda\left(P_{jt},P_t\right)\Pi_{jt+1}|z_{jt}\right\} \tag{1.93}$$

$$0 \leq K_{jt+1} \leq P_{jt}. \tag{1.94}$$

Following the same steps as in the main text, the asset price P_{jt} can be expressed as

$$P_{jt} = \alpha\frac{\lambda\left(P_{jt},P_t\right)}{R_{t+1}}Y_t^{\alpha_Y}\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}K_{jt+1}^{\frac{\theta-1}{\theta}}. \tag{1.95}$$

It is optimal for the entrepreneur to invest everything she raises, which allows firm capital to be written as

$$P_{jt} = K_{jt+1} = \left(\alpha\frac{\lambda\left(P_{jt},P_t\right)}{R_{t+1}}Y_t^{\alpha_Y}\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}\right)^{\theta}. \tag{1.96}$$

Plugging the expression for capital back into the entrepreneur's problem leads to the final problem

$$\max_{\lambda(P_{jt},P_t)} \quad \left(1-\lambda\left(P_{jt},P_t\right)\right)\lambda\left(P_{jt},P_t\right)^{\theta-1}c \tag{1.97}$$

$$c = \left(1-\alpha^{\theta-1}\right)\mathbb{E}\left\{A_{jt}|z_{jt}\right\}\left(\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}\frac{Y_t^{\alpha_Y}}{R_{t+1}}\right)^{\theta-1}. \tag{1.98}$$

The first-order condition to the simplified problem is

$$\begin{aligned}
&\frac{\partial}{\partial\lambda_{jt}}\left(1-\lambda\left(P_{jt},P_t\right)\right)\lambda\left(P_{jt},P_t\right)^{\theta-1} = 0\\
\iff &\left(\theta-1\right)\lambda\left(P_{jt},P_t\right)^{\theta-2} - \theta\lambda\left(P_{jt},P_t\right)^{\theta-1} = 0\\
&\iff \forall j,t: \quad \lambda_{jt} = \frac{\theta-1}{\theta}
\end{aligned} \tag{1.99}$$

Therefore, all entrepreneurs irrespective of $\left(P_{jt},P_t\right)$ sell a constant fraction $\lambda_{jt} = \frac{\theta-1}{\theta}$ of revenue to the financial market. The resulting dividend per share is

$$D_{jt+1} = \frac{\theta-1}{\theta}\alpha Y_{t+1}^{\alpha_Y}A_{jt}K_{jt+1}^{\frac{\theta-1}{\theta}}. \tag{1.100}$$

Assigning market power to entrepreneurs, therefore, effectively leads to a markup on the price of the intermediate good as traders only receive a fraction $\frac{\theta-1}{\theta}$ of firm revenue for completely funding firm investment. The effect is to depress investment, which can be undone through an ad-valorem subsidy of $\tau = \frac{\theta}{\theta-1}$ in the market for intermediate goods.

1.C.3 Entrepreneurs with Market Power: Credit Markets

The main focus of this paper is to study booms that are caused by productivity or sentiment. The literature extensively studies such booms in credit markets. The model can be extended to cover debt securities that are centrally traded instead of stock markets. Assume that the entrepreneur's technology is given by

$$Y_{jt} = \begin{cases} A^{\frac{\theta-1}{\theta}} K_{jt} & \text{w.p. } \pi_{jt} \\ 0 & \text{w.p. } 1-\pi_{jt} \end{cases}. \tag{1.101}$$

In the main text, entrepreneurs were sure to succeed , but their productivity was uncertain. Now, assume that entrepreneurs run projects that are either successful and give a certain payoff or fail and produce nothing. Success or failure is determined by the realization of a normally distributed variable,

$$\mathcal{P}\left(Y_{jt} > 0\right) = P\left(a_{jt} > \bar{a}\right) = \Phi\left(\frac{a_t - \bar{a}}{\sqrt{\sigma_a^2}}\right) = \pi_t. \tag{1.102}$$

The entrepreneur's project succeeds whenever $a_{jt} \sim \mathcal{N}\left(a_t, \sigma_a^2\right)$ has a realization above the threshold $\bar{a}$. Households have dispersed information about the firm-specific shock $s_{ijt} = a_{jt} + \frac{\eta_{ijt}+\varepsilon_{jt}}{\sqrt{\beta_{ijt}}}$ where η_{ijt} is idiosyncratic noise and ε_{jt} is correlated noise. Same as before, traders suffer from correlation neglect and perceive only their own signal to be $s_{ijt} = a_{jt} + \eta_{ijt}/\sqrt{\beta_{ijt}}$. The household's problem is the same as in the main model.

To finance their projects, entrepreneurs issue a unit mass of debt securities with the payoff

$$X_{jt} = \begin{cases} \lambda_{jt} & \text{if } Y_{jt} > 0 \\ 0 & \text{otherwise} \end{cases}. \tag{1.103}$$

The security pays an amount λ_{jt} when the project succeeds and pays zero otherwise.[30] The entrepreneur maximizes the revenue that she can keep in case of success after repaying her debt obligations

$$\max_{\lambda_{jt},K_{jt+1}} \quad \rho_{jt+1}Y_{jt+1}\left(a_{jt},K_{jt+1}\right)-X_{jt}\left(a_{jt},\lambda_{jt}\right) \tag{1.104}$$

$$0\leq K_{jt+1}\leq P_{jt}. \tag{1.105}$$

The entrepreneur invests all raised funds, $K_{jt+1}=P_{jt}$. According to the beliefs of the marginal trader, the price of debt and firm capital can then be written as

$$K_{jt+1}=P_{jt}=\frac{\lambda_{jt}}{R_{t+1}}\Phi\left(\frac{\tilde{\mathbb{E}}\left\{a_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}-\bar{a}}{\sqrt{\mathbb{V}}}\right), \tag{1.106}$$

where $\mathbb{V}=\left(\sigma_a^{-2}+\beta_{jt}\left(1+\sigma_\varepsilon^{-2}\right)\right)^{-1}$is the posterior uncertainty. The solution to the entrepreneur's problem is

$$\lambda_{jt}=\left(\frac{\theta-1}{\theta}\alpha Y_{t+1}^{\alpha_Y}A\right)^{\theta}\frac{\Phi\left(\frac{\tilde{\mathbb{E}}\left\{a_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}-\bar{a}}{\sqrt{\mathbb{V}}}\right)^{\theta-1}}{\left(R_{t+1}\right)^{\theta-1}}, \tag{1.107}$$

which depends on the market valuation of debt or equivalently the interest rate that entrepreneur j faces. Using (1.106) and (1.107) in the expression for firm-revenue allows the entrepreneur's decision to be expressed as a fraction of output in the case of success,

$$\frac{\lambda_{jt}}{\rho_{jt+1}Y_{jt+1}\left(a_{jt},K_{jt+1}\right)}=\frac{\theta-1}{\theta}. \tag{1.108}$$

This result recovers the optimal equity contract from section 1.C.2.

In contrast to the model with equity, there is an additional channel through which shocks affect information production. The binary payoff function introduces changes in the variance of outcomes for firms driven by productivity and sentiment shocks. The variance of outcomes is captured by $\pi_{jt}\left(1-\pi_{jt}\right)$, whereas riskiness normally would only be captured by the probability of failure, $1-\pi_{jt}$. Intuitively, a project is entirely

[30] Quantity and payoffs can be interchanged by denoting the mass of securities by λ_{jt} so the payoff in the good state is normalized to one. Instead, the quantity is normalized to one and the payoff is allowed to vary.

safe whenever the probability of success, π_{jt}, is equal to one. In this case, learning about the firm-specific shock, a_{jt}, is inconsequential. The same reasoning applies if the project is sure to fail ($\pi_{jt} = 0$). Therefore, the effect of changes to a_t is ambiguous. Positive shocks to a_t trigger additional information production only when π_{jt} was low before, but they crowd out information when debt becomes safe as a consequence. Therefore, aggregate shocks affect the (perceived) riskiness of debt.[31]

Although my model abstracts from banks and credit intermediation, it replicates the main stylized facts of credit booms before financial crises. First, credit booms are episodes of sharp increases in lending and economic activity (Jordà, Schularick, and Taylor 2011). This is the case in the model presented here, as the volume of credit increases in response to a positive aggregate shock. As a result, investment and economic activity increase. Second, credit becomes riskier as lending standards are relaxed, i.e., riskier firms get access to credit (Keys et al. 2010). In response to a sentiment shock, all firms are considered to be safer than they actually are. Because there is more scope for a change in beliefs for relatively risky firms, the sentiment shock leads to a disproportionate increase in funding for risky firms (low realization of a_{jt}). Third, credit spreads decrease in the boom phase before a financial crisis (Krishnamurthy and Muir 2017). In the case of sentiment-driven booms, all firms are perceived to be safer than they actually are and, therefore, spreads are low.

1.D Multi-Sector /-Country Model and Sector-/ Country-Specific Shocks

Let the economy consist of $N \in \mathbb{N}$ sectors or countries. Each consists of a unit mass of firms indexed by $nj \in N \times [0, 1]$. Similarly, each household now has one trader for each firm in each sector or country. The aggregate production function becomes

$$Y_t = L^{1-\alpha} \left[\sum_{n \in N} \left(\int_0^1 Y_{njt}^{\frac{\theta_n - 1}{\theta_n}} dnj \right)^{\frac{\theta_n}{\theta_n - 1} \frac{\theta - 1}{\theta}} \right]^{\frac{\alpha\theta}{\theta - 1}}. \tag{1.109}$$

[31] An alternative interpretation is that productivity shocks affect productivity conditional on success, A. In this case, productivity shocks would have no effect on the riskiness of debt and would behave similarly to a productivity shock in the model with equity markets.

where θ_n is the elasticity substitution inside sector or country $n \in N$. Productivity and sentiment shocks can now also be sector-specific, such that $a_{njt} \sim \mathcal{N}\left(a_{nt}, \sigma^2_{a,n}\right)$, $a_{nt} \sim \mathcal{N}\left(a_t, \sigma^2_a\right)$ and $\varepsilon_{njt} \sim \mathcal{N}\left(\varepsilon_{nt}, \sigma^2_{\varepsilon,n}\right)$, $\varepsilon_{nt} \sim \mathcal{N}\left(\varepsilon_t, \sigma^2_a\right)$. Sector-specific and aggregate shocks are observable.

In reality, booms rarely affect the whole economy equally. For example, the dot-com boom of the late 1990s was mainly about information technology and the emerging internet. Similarly, the housing boom in the 2000s concentrated in the financial and construction sector. In contrast to the economy-wide sentiment shock studied in the paper's main body, sector-specific sentiment shocks lead directly to an increase in capital misallocation as the marginal product of capital declines in the shocked sector.

Nonetheless, the main finding leads then to an additional insight: sector-specific sentiment shocks lead to an increase in capital misallocation inside the shocked sector. Not only is there too much investment in a specific sector, but this investment also flows to increasingly unproductive firms, thus amplifying the welfare costs of a sentiment boom. This result is captured in the following corollary analogously to Proposition 1.4.

Corollary 1.4. *A sector / country-specific positive sentiment shock can lead to an increase in capital misallocation inside the sector or country.*

At the same time, the redirection of capital investment towards the positively shocked sector can hurt non-shocked sectors, leading to a negative spillover of positive shocks across sectors. This is the case whenever aggregate investment is fixed ($\delta \to \infty$) or goods from different sectors are close substitutes ($\theta \to \infty$). This analysis also extends to a multi-country setting with free capital flows, in which a sentiment boom in one country leads to an increase in capital misallocation in both countries.

Corollary 1.5. *If aggregate investment is fixed ($\delta \to \infty$) or sector-/country-specific goods are close substitutes ($\theta \to \infty$), a sector- / country-specific positive shock leads to an increase in capital misallocation also in all other sectors / countries.*

1.E Full Social Planner Problem

In the main text, the social planner could only intervene by choosing information production. Now, the social planner can also choose consumption and investment in the aggregate and cross-section to maximize social welfare, defined as the sum of the utilities of all traders. Therefore, the social planner is able to achieve the second best by fixing all inefficiencies. The maximization problem is

$$\max_{K_{j1},C_0,C_1,\beta_{j0}} \quad C_0 + \delta\mathbb{E}_0\left\{C_1\right\} - \int_0^1 IA\left(\beta_{j0}\right)dj \tag{SPFull}$$

$$s.t. \quad K_1 = W_0 - C_0 \tag{1.110}$$

$$C_1 \leq Y_1\left(\left\{K_{j1}\right\}, \left\{\beta_{j0}\right\}\right) \tag{1.111}$$

$$C_0 \leq W_0 \tag{1.112}$$

$$K_{j1}, C_0, C_1, \beta_{j0} \geq 0. \tag{1.113}$$

Constraint (1.110) states that aggregate capital in period 1 is equal to endowments W_0 minus youth consumption C_0. Resource constraints for consumption are given in (1.111) and (1.112). Finally, non-negativity constraints on consumption, information production and capital are given in (1.113). The solution to the social planner's problem is given in the following proposition.

Proposition 1.11. *The social planner's allocation under perfect information about aggregate shocks* $\{a_0, \varepsilon_0\}$ *is given by* $\left\{C_0^{SP}, K_{j1}^{SP}, K_1^{SP}, \beta_0^{SP}\right\}$, *where*

$$K_{j1}^{SP} = \frac{\mathbb{E}\left\{A_{j0}|z_{j0}\right\}^\theta}{\int_0^1 \mathbb{E}\left\{A_{j0}|z_{j0}\right\}^\theta dj} K_1^{SP} \tag{1.114}$$

leading to aggregate output

$$Y_1^{SP} = A_0^{SP}\left(K_1^{SP}\right)^\alpha \quad \textit{with } A_0^{SP} = \left(\int_0^1 \mathbb{E}\left\{A_{j0}|z_{j0}\right\}^\theta dj\right)^{\frac{\alpha}{\theta-1}}. \tag{1.115}$$

The interest rate is

$$R_1^{SP} = \alpha A_0^{SP}\left(K_1^{SP}\right)^{\alpha-1}, \tag{1.116}$$

leading to aggregate investment

$$K_1^{SP} = \min\left\{\left(\alpha\delta A_0^{SP}\right)^{\frac{1}{1-\alpha}}, W_0\right\}. \tag{1.117}$$

The symmetric information production choice is

$$\textit{for all } \beta_{j0} = \beta_0^{SP} : \delta \left. \frac{\partial A_0^{SP}}{\partial \beta_0} \right|_{\beta_0 = \beta_0^{SP}} \left(K_1^{SP}\right)^{\alpha} = \left. \frac{\partial I A_0}{\partial \beta_0} \right|_{\beta_0 = \beta_0^{SP}} . \tag{1.118}$$

The social planner fixes the two aforementioned inefficiencies. First, the social planner distributes capital optimally by attributing the correct precision to the price signal z_{jt} as in (1.28) and (1.114). As a result, ex-ante marginal products of capital are equalized between firms. This reallocation of capital leads to an increase in TFP as in Proposition 1.3 compared to the competitive allocation. Second, the social planner chooses information production β_0^{SP} to increase TFP instead of trading rents. Given that the social planner optimally distributes capital between firms as in (1.114), an increase in β_0^{SP} always benefits aggregate productivity A_0^{SP}.

1.E.1 Implementation

In this section, I investigate how the social planner can implement the centralized allocation through the use of taxes and subsidies. Net proceeds and costs of taxes and subsidies are distributed lump-sum between old traders.

The social planner can apply a tax/subsidy on dividend income to achieve the constrained efficient allocation of capital. Under this state-dependent tax/subsidy, traders receive

$$\Pi_{j1}^{DE} = \tau^{Bias}\left(z_{j0}\right) \Pi_{j1}, \quad \text{where } \tau^{Bias}\left(z_{j0}\right) = \frac{\mathbb{E}\left\{A_{j0} | z_{j0}\right\}}{\tilde{\mathbb{E}}\left\{A_{j0} | s_{ij0} = z_{j0}, z_{j0}\right\}}. \tag{1.119}$$

As seen in Figure 1.5, $\tau^{Bias}\left(z_{j0}\right)$ is a subsidy on dividends whenever $K_{j1}^{eff} < K_{j1}$. If the social planner has information about aggregate shocks, the tax/subsidy corrects also for aggregate sentiment shocks through the marginal trader's expectations $\tilde{\mathbb{E}}\left\{A_{j0} | s_{ij0} = z_{j0}, z_{j0}\right\}$. A tax (subsidy) can lower (increase) investment in response to a positive (negative) sentiment shock.

Moreover, a tax/subsidy $\tau^{Info}\left(\beta_{ij0}\right)$ on information production is sufficient to induce the socially optimal level,

$$\frac{\partial IA^{DE}\left(\beta_{ij0}\right)}{\partial\beta_{ij0}} = \tau^{Info}\left(\beta_{ij0}\right)\frac{\partial IA\left(\beta_{ij0}\right)}{\partial\beta_{ij0}}, \quad \tau^{Info}\left(\beta_{ij0}\right) = \frac{\widetilde{MB}\left(\beta_{ij0},\beta_{j0}\right)\Big|_{\beta_{ij0}=\beta_{j0}}}{\delta\frac{\partial A_0}{\partial\beta_0}K_1^{\alpha}}. \tag{1.120}$$

Applying the after-tax marginal cost leads directly to the first-order condition as in (1.118). The results are summarized in the following proposition:

Proposition 1.12. *The social planner's allocation* $\left\{K_1^{SP}, K_{j1}^{SP}, \beta_0^{SP}\right\}$ *can be implemented through taxes/subsidies (1.119) and (1.120).*

Alternatively, the social planner can use transaction taxes to implement the optimal capital allocation. Since Tobin (1972), financial transaction taxes have been discussed with the objective of reducing volatility by making short-term speculation less profitable. This perspective cannot be studied here as assets are short-lived and only traded once. Nonetheless, a transaction tax can be used to drive a wedge between how much traders pay for shares and how much is invested in capital. The following proposition shows how such a transaction tax can be used to stabilize investment against sentiment shocks and reallocate capital across firms.

Corollary 1.6. *(i) Aggregate investment can be stabilized with respect to sentiment shocks through a transaction tax,*

$$K_{j1}^{DE} = \tau^{Trans}\left(\varepsilon_0\right)P_{j0}, \quad \tau^{Trans}\left(\varepsilon_0\right) = \exp\left\{-\omega_{s\varepsilon}\varepsilon_0\right\}, \quad \omega_{s\varepsilon} = \frac{\sqrt{\beta_0}}{\sigma_a^{-2}+\beta_0\left(1+\sigma_\varepsilon^{-2}\right)}. \tag{1.121}$$

(ii) The dividend tax/subsidy $\tau^{Bias}\left(z_{j0}\right)$ *(1.119) can be substituted by a state-dependent transaction tax,*

$$K_{j1}^{DE} = \tau^{Trans}\left(P_{j0}\right)P_{j0}, \quad \tau^{Trans}\left(P_{j0}\right) = \frac{\mathbb{E}\left\{A_{j0}|z_{j0}\right\}}{\tilde{\mathbb{E}}\left\{A_{j0}|s_{ij0}=z_{j0},z_{j0}\right\}}. \tag{1.122}$$

1.F Information Structure

I assume that traders are overconfident in that they wrongly believe that sentiment drives the beliefs of all other traders but does not drive their own beliefs. Whereas it is empirically reasonable to assume that behavioral biases play a role in financial markets, I chose this approach for tractability. Avoiding the introduction of non-optimizing agents greatly simplifies embedding a model of informative financial markets in a macro setting and facilitates the welfare analysis. Moreover, overconfidence is sufficient to motivate costly information production and to avoid the Grossman-Stiglitz paradox. This assumption is not necessary for deriving the main result that sentiment shocks crowd out information and can identically be derived with noise trades in partial equilibrium.

In the following, I walk through different assumptions for the information structure and their relationships to information aggregation and production.

Exogenous Public Signals

The simplest case is one in which traders do not have private signals but instead observe public signals of the form $z_{jt} = a_{jt} + \varepsilon_{jt}/\sqrt{\beta}$. This setting mirrors the allocation in Proposition 1.3. However, it has nothing to say about the origin of the signal. How does it come about, and what determines its precision?

Heterogeneous Private Signals

To say something about the aggregation of information, endow traders with heterogeneous private signals as in (1.5). Following the same steps as in section 1.3.3 leads to the market equilibrium.

Under rationality, the constrained efficient capital allocation as in Proposition 1.3 is achieved, but information production is ruled out. As in the model with overconfidence, observing the asset price is isomorphic to observing $z_{jt} = \int s_{ijt} dj$. Rational traders realize that they have nothing to learn from their private signal after observing the public signal z_{jt}. However, setting up this equilibrium requires that traders use their private signals to make the buying decision. In this setting, traders are indifferent between buying and not buying, as they all share the same information set. Therefore, the indifference can be broken in favor of buying whenever $s_{ijt} \geq z_{jt}$.[32]

The main drawback of this approach is that it rules out costly information production. The private signal s_{ijt} becomes fully uninformative after observing the public signal z_{jt} in equilibrium ($\forall i : \beta_{ijt} = \beta_{jt}$). In this case, the trader finds it optimal to devi-

[32] This class of equilibria is referred to as "fully revealing rational expectations equilibria" and is studied in Grossman (1976).

ate to $\beta_{ijt} < \beta_{jt}$, which guarantees an informative private signal. However, since traders are ex-ante homogeneous, asymmetric equilibria cannot exist. There is no equilibrium with costly information production and rationality, similar to the result in Grossman and Stiglitz (1980).

To overcome this problem, I assume that traders think that their private signal does not contain sentiment. Therefore, they do not discard their private signal s_{ijt} after observing the price signal z_{jt}. The posterior of trader ij becomes

$$a_{jt}|s_{ijt}, z_{jt} \sim \mathcal{N}\left(\frac{\beta_{ijt}s_{ijt} + \beta_{jt}\sigma_{\varepsilon}^{-2}z_{jt}}{\sigma_a^{-2} + \beta_{ijt} + \beta_{jt}\sigma_{\varepsilon}^{-2}}, \frac{1}{\sigma_a^{-2} + \beta_{ijt} + \beta_{jt}\sigma_{\varepsilon}^{-2}}\right), \tag{1.123}$$

where I have marked in blue terms that follow from the overconfidence assumption. It follows that traders have posteriors that are too precise, as they think that their private signals remain informative after observing z_{jt}. This misperception motivates traders to invest in their private signal with the anticipation of trading rents. Finally, this bias leads to an overreaction of prices to the price signal z_{jt} as described in section 1.4.1. This price distortion appears also in models with rational and noise traders, if rational traders learn from prices and have heterogeneous private signals.[33] The main focus of this paper, however, is not on the price distortion, but rather on time-varying price informativeness and the allocational efficiency of financial markets.

1.G Proofs

Proof of Proposition 1.1. This proof follows the same steps as the proof for Proposition 1 in Albagli, Hellwig, and Tsyvinski (2017), since the financial market in my model is isomorphic to their model. Their proof is repeated here for completeness. The only difference is that K_{jt+1} depends on the price signal z_{jt}, whereas k in Albagli, Hellwig, and Tsyvinski (2017) is determined before trading takes place. Therefore, it is necessary to assume that $K_{jt+1}(z_{jt})$ is non-decreasing in z_{jt} as the price might otherwise be not invertible, which is confirmed ex-post. The proof begins in the following.

There must be a threshold $\hat{s}(P_{jt})$ such that all households with $s_{ijt} \geq \hat{s}(P_{jt})$ find it profitable to buy two units of share j and otherwise abstain from trading. It follows that

[33]For a more detailed discussion, see Albagli, Hellwig, and Tsyvinski (2011a, 2015, 2021).

the price must be equal to the valuation of the trader who is merely indifferent between buying and not buying,

$$P_{jt} = \frac{1}{R_{t+1}} \tilde{\mathbb{E}} \left\{ D\left(A_{jt}, K_{jt+1}\right) | s_{ijt} = \hat{s}\left(P_{jt}\right), P_{jt} \right\}. \tag{1.124}$$

This monotone demand schedule leads to total demand

$$D(\theta, \varepsilon, P) = 2\left(1 - \Phi\left(\sqrt{\beta_{jt}}\left(\hat{s}\left(P_{jt}\right) - a_{jt}\right) - \varepsilon_{jt}\right)\right). \tag{1.125}$$

Equalizing total demand with a normalized supply of one leads to the market-clearing condition

$$2\left(1 - \Phi\left(\sqrt{\beta_{jt}}\left(\hat{s}\left(P_{jt}\right) - a_{jt}\right) - \varepsilon_{jt}\right)\right) = 1, \tag{1.126}$$

with the unique solution $\hat{s}\left(P_{jt}\right) = z_{jt} = a_{jt} + \frac{\varepsilon_{jt}}{\sqrt{\beta_{jt}}}$. If P_{jt} is pinned down by z_{jt}, then P_{jt} is invertible, given that K_{jt+1} is non-decreasing in z_{jt}. It follows, then, that observing P_{jt} is equivalent to observing $z_{jt} \sim \mathcal{N}\left(a_{jt}, \beta_{jt}^{-1}\sigma_{\varepsilon}^{2}\right)$. Traders treat the signal z_{jt} and their private signal $s_{ijt} \sim \mathcal{N}\left(a_{jt}, \beta_{ijt}^{-1}\right)$ as mutually independent. Using this result, the price can be restated as

$$P\left(z_{jt}, K_{jt+1}\right) = \frac{1}{R_{t+1}} \tilde{\mathbb{E}} \left\{ \Pi\left(A_{jt}, K_{jt+1}\right) | s_{ijt} = z_{jt}, z_{jt} \right\}, \tag{1.127}$$

where posterior expectations of trader ij are given by

$$a_{jt} | s_{ijt}, z_{jt} \sim \mathcal{N}\left(\frac{\sigma_a^{-2} a_t + \beta_{ijt} s_{ijt} + \beta_{jt}\sigma_{\varepsilon}^{-2} z_{jt}}{\sigma_a^{-2} + \beta_{ijt} + \beta_{jt}\sigma_{\varepsilon}^{-2}}, \frac{1}{\sigma_a^{-2} + \beta_{ijt} + \beta_{jt}\sigma_{\varepsilon}^{-2}}\right). \tag{1.128}$$

Using the result that the firm invests all proceeds into capital ($K_{jt+1} = P_{jt}$), it follows indeed that K_{jt+1} is non-decreasing in z_{jt} and P_{jt} is an invertible function of the price signal z_{jt}. It remains to show the uniqueness of the above solution. Begin with the assumption that demand $x\left(s_{ijt}, P_{jt}\right)$ is non-increasing in P_{jt}. It follows that $\hat{s}\left(P_{jt}\right)$ is non-decreasing in P_{jt}. There are two cases to differentiate. First, if $\hat{s}\left(P_{jt}\right)$ is strictly increasing in P_{jt}, then the price is indeed uniquely pinned-down by z_{jt} and invertible; it can be expressed like above. Secondly, assume that the threshold is flat over some interval, such that $\hat{s}\left(P_{jt}\right) = \hat{s}$ over some interval $P_{jt} \in \left(P', P''\right)$ for $P' \neq P''$. Furthermore, choose $\epsilon > 0$ small enough such that $\hat{s}\left(P_{jt}\right)$ is increasing to the left and the right of the interval, i.e., over $P_{jt} \in (P' - \epsilon, P')$ and $P_{jt} \in (P'', P'' + \epsilon)$. In these regions, $\hat{s}(P_{jt})$ is monotonically increasing in P_{jt}, which is uniquely pinned down by z_{jt} and invertible; observing the price is equivalent to observing the signal z_{jt}. In this case the price can be expressed as before for $z_{jt} \in (\hat{s}(P' - \varepsilon), \hat{s})$ and $z_{jt} \in (\hat{s}, \hat{s}(P'' + \epsilon))$. This leads to a contradiction in the assumption that $P' \neq P''$, because $P(z_{jt}, K_{jt+1})$ is both continuous and monotonically increasing in z_{jt}. Therefore, $\hat{s}(P_{jt})$ cannot be flat and the above solution is indeed unique. □

Proof of Proposition 1.2. (i) Using (1.4) in (1.21) leads to the expression for firm capital

$$K_{jt+1} = \left(\frac{\alpha Y_{t+1}^{\alpha_Y}}{R_{t+1}} \tilde{\mathbb{E}} \left\{ A_{jt} | s_{ijt} = z_{jt}, z_{jt} \right\} \right)^{\theta}. \tag{1.129}$$

Plugging R_{t+1} from (1.22) into (1.21) using (1.129) leads to

$$K_{jt+1} = \frac{\tilde{\mathbb{E}} \left\{ A_{jt} | s_{ijt} = z_{jt}, z_{jt} \right\}^{\theta}}{\int_0^1 \tilde{\mathbb{E}} \left\{ A_{jt} | s_{ijt} = z_{jt}, z_{jt} \right\}^{\theta} dj} K_{t+1}, \tag{1.130}$$

which finishes the proof.

(ii) Plugging the above expression for firm capital (1.25) into the aggregate production function (1.2) leads to

$$\begin{aligned}
Y_t &= \left(\int_0^1 Y_{jt}^{\frac{\theta-1}{\theta}} dj \right)^{\frac{\alpha\theta}{\theta-1}} \\
&= \left(\int_0^1 A_{jt-1} K_{jt}^{\frac{\theta-1}{\theta}} dj \right)^{\frac{\alpha\theta}{\theta-1}} \\
&= \frac{\left(\int_0^1 A_{jt-1} \tilde{\mathbb{E}} \left\{ A_{jt-1} | s_{ijt-1} = z_{jt-1}, z_{jt-1} \right\}^{\theta-1} dj \right)^{\frac{\alpha\theta}{\theta-1}}}{\left(\int_0^1 \tilde{\mathbb{E}} \left\{ A_{jt-1} | s_{ijt-1} = z_{jt-1}, z_{jt-1} \right\}^{\theta} dj \right)^{\alpha}} K_t^{\alpha} \\
&= A_{t-1} L^{1-\alpha} K_t^{\alpha}
\end{aligned} \tag{1.131}$$

where total factor productivity is

$$
\begin{aligned}
A_{t-1} &= \frac{\left(\int_0^1 A_{jt-1}\tilde{\mathbb{E}}\left\{A_{jt-1}|s_{ijt-1}=z_{jt-1},z_{jt-1}\right\}^{\theta-1}dj\right)^{\frac{\alpha\theta}{\theta-1}}}{\left(\int_0^1 \tilde{\mathbb{E}}\left\{A_{jt-1}|s_{ijt-1}=z_{jt-1},z_{jt-1}\right\}^{\theta}dj\right)^{\alpha}} \\
&= \exp\left\{\theta a_{t-1} + \left(\left(\theta-1\right)\omega_a+1\right)^2\frac{\sigma_a^2}{2} + \left(\theta-1\right)^2\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{2} + \left(\theta-1\right)\omega_{s\varepsilon}\varepsilon_{t-1}\right. \\
&\left. + \frac{\left(\theta-1\right)}{2}\mathbb{V}\right\}^{\frac{\alpha\theta}{\theta-1}} \\
&: \exp\left\{\theta a_{t-1} + \theta^2\omega_a^2\frac{\sigma_a^2}{2} + \theta^2\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{2} + \theta\omega_{s\varepsilon}\varepsilon_{t-1} + \frac{\theta}{2}\mathbb{V}\right\}^{\alpha} \\
&= \exp\left(\frac{\alpha\theta}{\theta-1}a_{t-1} + \left(\frac{\alpha\theta}{\theta-1}\left(\left(\theta-1\right)\omega_a+1\right)^2 - \alpha\theta^2\omega_a^2\right)\frac{\sigma_a^2}{2}\right. \\
&\left. + \left(\alpha\theta\left(\theta-1\right)-\alpha\theta^2\right)\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{2}\right) \\
&= \exp\left(\frac{\alpha\theta}{\theta-1}a_{t-1} + \alpha\theta\left(\left(\theta-1\right)\omega_a^2 + 2\omega_a + \frac{1}{\theta-1} - \theta\omega_a^2\right)\frac{\sigma_a^2}{2} - \alpha\theta\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{2}\right) \\
&= \exp\left(\frac{1}{\theta-1}\left(a_{t-1}+\frac{\sigma_a^2}{2}\right) + \omega_a\left(2-\omega_a\right)\frac{\sigma_a^2}{2} - \omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{2}\right)^{\alpha\theta}. \qquad (1.132)
\end{aligned}
$$

The weights $\{\omega_a, \omega_\varepsilon, \omega_{s\varepsilon}\}$ and $\mathbb{V}$ are derived in Appendix 1.A. Finally, total factor productivity can be expressed as

$$
\ln A_{t-1}\left(a_{t-1},\beta_{t-1}\right) = \underbrace{\frac{\alpha\theta}{\theta-1}\left(a_{t-1}+\frac{\sigma_a^2}{2}\right)}_{\textit{exogenous}} + \underbrace{\kappa^a\left(\beta_{t-1}\right)\sigma_a^2 - \kappa^\varepsilon\left(\beta_{t-1}\right)\sigma_\varepsilon^2}_{\textit{allocative efficiency}} \qquad (1.133)
$$

where $\kappa^a\left(\beta_{t-1}\right) = \frac{\alpha\theta}{2}\omega_a\left(2-\omega_a\right)$ and $\kappa^\varepsilon\left(\beta_{t-1}\right) = \frac{\alpha\theta}{2}\omega_\varepsilon^2$.

(iii) I will show that the allocative efficiency component of TFP takes its minimum for $\beta_{t-1} > 0$ if $\sigma_\varepsilon^2 > 1$. The allocational efficiency component is proportional to

$$
\begin{aligned}
&\omega_a (2-\omega_a)\,\sigma_a^2 - \omega_\varepsilon^2 \sigma_\varepsilon^2 \\
&= \frac{2\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)}{\sigma_a^{-2}+\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)}\sigma_a^2 - \frac{\beta_{t-1}^2\left(1+\sigma_\varepsilon^{-2}\right)^2}{\left(\sigma_a^{-2}+\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)\right)^2}\sigma_a^2 - \frac{\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)^2}{\left(\sigma_a^{-2}+\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)\right)^2}\sigma_\varepsilon^2 \\
&= \frac{2\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right) + 2\beta_{t-1}^2\left(1+\sigma_\varepsilon^{-2}\right)^2\sigma_a^2 - \beta_{t-1}^2\left(1+\sigma_\varepsilon^{-2}\right)^2\sigma_a^2 - \beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)^2\sigma_\varepsilon^2}{\left(\sigma_a^{-2}+\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)\right)^2} \\
&= \frac{2\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right) + \beta_{t-1}^2\left(1+\sigma_\varepsilon^{-2}\right)^2\sigma_a^2 - \beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)^2\sigma_\varepsilon^2}{\left(\sigma_a^{-2}+\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)\right)^2} \\
&= \frac{\beta_{t-1}^2\sigma_a^2\left(1+\sigma_\varepsilon^{-2}\right)^2 + \beta_{t-1}\left(\sigma_\varepsilon^{-2}-\sigma_\varepsilon^2\right)}{\left(\sigma_a^{-2}+\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)\right)^2}, \qquad (1.134)
\end{aligned}
$$

which is weakly positive for all values of β_{t-1} if $\sigma_\varepsilon^2 < 1$.

(iv) Using the previous result, it remains to take the derivative of (1.134) with respect to β_{t-1}. Denote

$$
a = \beta_{t-1}^2\sigma_a^2\left(1+\sigma_\varepsilon^{-2}\right)^2 + \beta_{t-1}\left(\sigma_\varepsilon^{-2}-\sigma_\varepsilon^2\right) \qquad (1.135)
$$

$$
b = \left(\sigma_a^{-2}+\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)\right)^2 \qquad (1.136)
$$

After some algebra,

$$
\begin{aligned}
\frac{\partial a}{\partial \beta_{t-1}} b = {} & 2\beta_{t-1}^3\sigma_a^2\left(1+\sigma_\varepsilon^{-2}\right)^4 + 4\beta_{t-1}^2\left(1+\sigma_\varepsilon^{-2}\right)^3 + 2\beta_{t-1}\sigma_a^{-2}\left(1+\sigma_\varepsilon^{-2}\right)^2 \\
& + \beta_{t-1}^2\left(1+\sigma_\varepsilon^{-2}\right)^2\left(\sigma_\varepsilon^{-2}-\sigma_\varepsilon^2\right) + 2\beta_{t-1}\sigma_a^{-2}\left(1+\sigma_\varepsilon^{-2}\right)\left(\sigma_\varepsilon^{-2}-\sigma_\varepsilon^2\right) \\
& + \sigma_a^{-4}\left(\sigma_\varepsilon^{-2}-\sigma_\varepsilon^2\right) \qquad (1.137)
\end{aligned}
$$

and

$$
\begin{aligned}
\frac{\partial b}{\partial \beta_{t-1}} a = {} & 2\beta_{t-1}^3\sigma_a^2\left(1+\sigma_\varepsilon^{-2}\right)^4 + 2\beta_{t-1}^2\left(1+\sigma_\varepsilon^{-2}\right)^3 \\
& + 2\beta_{t-1}^2\left(\sigma_\varepsilon^{-2}-\sigma_\varepsilon^2\right)\left(1+\sigma_\varepsilon^{-2}\right)^2 + 2\beta_{t-1}\sigma_a^{-2}\left(\sigma_\varepsilon^{-2}-\sigma_\varepsilon^2\right)\left(1+\sigma_\varepsilon^{-2}\right) \qquad (1.138)
\end{aligned}
$$

Dropping the positive denominator, the derivative of (1.134) is $\frac{\partial a}{\partial \beta_{t-1}} b - \frac{\partial b}{\partial \beta_{t-1}} a$, which is after dividing through $\left(1+\sigma_\varepsilon^{-2}\right)^2$,

$$
2\beta_{t-1}^2 + \beta_{t-1}^2\sigma_\varepsilon^{-2} + 2\beta_{t-1}\sigma_a^{-2} + \sigma_a^{-4}\frac{\sigma_\varepsilon^{-2}-\sigma_\varepsilon^2}{\left(1+\sigma_\varepsilon^{-2}\right)^2} + \beta_{t-1}^2\sigma_\varepsilon^2, \qquad (1.139)
$$

which is positive for all values of β_{t-1} if $\sigma_\varepsilon^2 < 1$.

$\square$

Lemma 1.1. *Denote the efficient Bayesian weights*

$$\omega_a^{eff} = \frac{\beta_t \sigma_\varepsilon^{-2}}{\sigma_a^{-2} + \beta_t \sigma_\varepsilon^{-2}} \tag{1.140}$$

$$\omega_\varepsilon^{eff} = \frac{\sqrt{\beta_t} \sigma_\varepsilon^{-2}}{\sigma_a^{-2} + \beta_t \sigma_\varepsilon^{-2}} \tag{1.141}$$

and posterior uncertainty

$$\mathbb{V}^{eff} = \frac{1}{\sigma_a^{-2} + \beta_t \sigma_\varepsilon^{-2}}, \tag{1.142}$$

These are the weights that a rational uninformed observer would use after observing the price signal z_{jt}, in contrast to the Bayesian weights that overconfident traders use as introduced in 1.A. Then the following relationships hold:

(i)

$$\underbrace{\left(\omega_a^{eff}\right)^2 \sigma_a^2 + \left(\omega_\varepsilon^{eff}\right)^2 \sigma_\varepsilon^2}_{=Var(\mathbb{E}\{a_{jt}|z_{jt}\})} + \underbrace{\mathbb{V}^{eff}}_{Var(a_{jt}|z_{jt})} = \underbrace{\sigma_a^2}_{Var(a_{jt})} . \tag{1.143}$$

(ii)

$$\sigma_a^2 - \mathbb{V}^{eff} = \omega_a^{eff} \sigma_a^2 \tag{1.144}$$

(iii)

$$\theta \omega_a^{eff} \sigma_a^2 + \mathbb{V}^{eff} = (\theta - 1)\, \omega_a^{eff} \sigma_a^2 + \sigma_a^2 \tag{1.145}$$

Proof. (i):

$$\begin{aligned}
\left(\omega_a^{eff}\right)^2 \sigma_a^2 + \left(\omega_\varepsilon^{eff}\right)^2 \sigma_\varepsilon^2 + \mathbb{V}^{eff} &= \frac{\beta_t^2 \sigma_\varepsilon^{-4}}{\left(\sigma_a^{-2} + \beta_t \sigma_\varepsilon^{-2}\right)^2} \sigma_a^2 + \frac{\beta_t \sigma_\varepsilon^{-4}}{\left(\sigma_a^{-2} + \beta_t \sigma_\varepsilon^{-2}\right)^2} \sigma_\varepsilon^2 \\
&\quad + \frac{1}{\sigma_a^{-2} + \beta_t \sigma_\varepsilon^{-2}} \\
&= \left(\mathbb{V}^{eff}\right)^2 \left(\sigma_a^{-2} + 2\beta_t \sigma_\varepsilon^{-2} + \beta_t^2 \sigma_\varepsilon^{-4} \sigma_a^2\right) \\
&= \left(\mathbb{V}^{eff}\right)^2 \left(\sigma_a^{-4} + 2\beta_t \sigma_\varepsilon^{-2} \sigma_a^{-2} + \beta_t^2 \sigma_\varepsilon^{-4}\right) \sigma_a^2 \\
&= \left(\mathbb{V}^{eff}\right)^2 \left(\sigma_a^{-2} + \beta_t \sigma_\varepsilon^{-2}\right)^2 \sigma_a^2 \\
&= \sigma_a^2.
\end{aligned} \tag{1.146}$$

(ii):

$$\begin{aligned}
\sigma_a^2 - \mathbb{V}^{eff} &= \sigma_a^2 - \frac{1}{\sigma_a^2 + \beta\sigma_\varepsilon^{-2}} \\
&= \frac{\sigma_a^2\sigma_a^{-2} + \sigma_a^2\beta\sigma_\varepsilon^{-2} - 1}{\sigma_a^{-2} + \beta\sigma_\varepsilon^{-2}} \\
&= \frac{\beta\sigma_\varepsilon^{-2}}{\sigma_a^{-2} + \beta\sigma_\varepsilon^{-2}}\sigma_a^2 \\
&= \omega_a^{eff}\sigma_a^2 \qquad (1.147)
\end{aligned}$$

(iii):

$$\begin{aligned}
\theta\omega_a^{eff}\sigma_a^2 + \mathbb{V}^{eff} &= \frac{\theta\beta_{t-1}\sigma_\varepsilon^{-2}\sigma_a^2}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}} + \frac{1}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}} \\
&= \frac{\theta\beta_{t-1}\sigma_\varepsilon^{-2} + \sigma_a^{-2}}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}}\sigma_a^2 \\
\text{Add and subtract} &= \frac{\theta\beta_{t-1}\sigma_\varepsilon^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2} - \beta_{t-1}\sigma_\varepsilon^{-2} + \sigma_a^{-2}}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}}\sigma_a^2 \\
\text{Split} &= (\theta - 1)\frac{\beta_{t-1}\sigma_\varepsilon^{-2}}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}}\sigma_a^2 + \frac{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}}\sigma_a^2 \\
&= (\theta - 1)\,\omega_a^{eff}\sigma_a^2 + \sigma_a^2 \qquad (1.148)
\end{aligned}$$

□

Proof of Proposition 1.3. (i) An efficient allocation of capital equalizes marginal products between firms. Demand for firm capital follows from the following maximization problem

$$\max_{K_{jt+1}} \alpha Y_{t+1}^{\alpha_Y}\mathbb{E}\left\{A_{jt}|z_{jt}\right\} K_{jt+1}^{\frac{\theta-1}{\theta}} - R_{t+1}K_{jt+1} \qquad (1.149)$$

with the first-order condition

$$K_{jt+1} = \left(\frac{\theta-1}{\theta}\frac{\mathbb{E}\left\{A_{jt}|z_{jt}\right\}}{R_{t+1}}\alpha Y_{t+1}^{\alpha_Y}\right)^{\theta}. \qquad (1.150)$$

Integrating over all firms on both sides yields

$$R_{t+1} = \left(\int_0^1 \mathbb{E}\left\{ A_{jt} | z_{jt} \right\}^{\theta} dj \right)^{\frac{1}{\theta}} \alpha \frac{\theta - 1}{\theta} Y_{t+1}^{\alpha_Y} K_{t+1}^{-\frac{1}{\theta}}. \tag{1.151}$$

Plugging this expression back into the first-order condition leads to the constrained efficient allocation

$$K_{jt+1} = \frac{\mathbb{E}\left\{ A_{jt} | z_{jt} \right\}^{\theta}}{\int_0^1 \mathbb{E}\left\{ A_{jt} | z_{jt} \right\}^{\theta} dj} K_{t+1}. \tag{1.152}$$

(ii) Plugging (1.28) into (1.2) leads to

$$Y_t = A_{t-1}^{eff} K_t^{\alpha}, \tag{1.153}$$

where the constrained efficient level of total factor productivity is

$$A_{t-1}^{eff} = \frac{\left(\int_0^1 A_{jt-1} \mathbb{E}\left\{ A_{jt-1} | z_{jt-1} \right\}^{\theta - 1} dj \right)^{\frac{\alpha\theta}{\theta-1}}}{\left(\int_0^1 \mathbb{E}\left\{ A_{jt-1} | z_{jt-1} \right\}^{\theta} dj \right)^{\alpha}} \overset{\text{L.I.E.}}{=} \left(\int_0^1 \mathbb{E}\left\{ A_{jt-1} | z_{jt-1} \right\}^{\theta} dj \right)^{\frac{\alpha}{\theta-1}}. \tag{1.154}$$

The analytical expression can be obtained by evaluating the conditional expectations and using the constrained efficient Bayesian weights and posterior uncertainty,

$$\omega_p^{eff} = \frac{\sigma_a^{-2}}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}}, \quad \omega_a^{eff} = \frac{\beta_{t-1}\sigma_\varepsilon^{-2}}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}} \tag{1.155}$$

$$\omega_\varepsilon^{eff} = \frac{\sqrt{\beta_{t-1}}\sigma_\varepsilon^{-2}}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}}, \quad \mathbb{V}^{eff} = \frac{1}{\sigma_a^{-2} + \beta_{t-1}\sigma_\varepsilon^{-2}} \tag{1.156}$$

leading to

$$
\begin{aligned}
A_{t-1}^{eff} &= \left(\int_0^1 \mathbb{E}\left\{A_{j0}|z_{j0}\right\}^{\theta} dj\right)^{\frac{\alpha}{\theta-1}} \\
&= \left(\int_0^1 \left\{\exp\left\{\theta\omega_p^{eff} a_t + \theta\omega_a^{eff} a_{jt} + \theta\omega_\varepsilon^{eff}(\varepsilon_{jt}-\varepsilon_t) + \theta\frac{\mathbb{V}^{eff}}{2}\right\}\right\} dj\right)^{\frac{\alpha}{\theta-1}} \\
&= \exp\left\{\theta\omega_p^{eff} a_t + \theta\omega_a^{eff} a_t + \frac{\theta^2}{2}\left(\omega_a^{eff}\right)^2\sigma_a^2 + \frac{\theta^2}{2}\left(\omega_\varepsilon^{eff}\right)^2\sigma_\varepsilon^2 + \theta\frac{\mathbb{V}^{eff}}{2}\right\}^{\frac{\alpha}{\theta-1}} \\
&= \exp\left\{a_t + \frac{\theta}{2}\left(\omega_a^{eff}\right)^2\sigma_a^2 + \frac{\theta}{2}\left(\omega_\varepsilon^{eff}\right)^2\sigma_\varepsilon^2 + \frac{\mathbb{V}^{eff}}{2}\right\}^{\frac{\alpha\theta}{\theta-1}} \\
&\overset{\text{Lemma 1.1 (i) and (ii)}}{=} \exp\left\{a_t + \frac{1}{2}\left(\theta\omega_a^{eff}\sigma_a^2 + \mathbb{V}^{eff}\right)\right\}^{\frac{\alpha\theta}{\theta-1}} \\
&\overset{\text{Lemma 1.1 (iii)}}{=} \exp\left\{a_t + \frac{1}{2}\left(\sigma_a^2 + (\theta-1)\,\omega_a^{eff}\sigma_a^2\right)\right\}^{\frac{\alpha\theta}{\theta-1}} \\
&= \exp\left\{\frac{1}{1-\theta}\left(a_t + \frac{\sigma_a^2}{2}\right) + \omega_a^{eff}\frac{\sigma_a^2}{2}\right\}^{\alpha\theta}
\end{aligned}
\tag{1.157}
$$

TFP under the efficient allocation of capital can be similarly decomposed into two expressions,

$$
\ln A_{t-1}^{eff} = \underbrace{\frac{\alpha\theta}{\theta-1}\left(a_{t-1} + \frac{\sigma_a^2}{2}\right)}_{\textit{exogenous}} + \underbrace{\alpha\theta\omega_a^{eff}\frac{\sigma_a^2}{2}}_{\textit{allocative efficiency}} . \tag{1.158}
$$

It follows that

$$
\frac{\partial\omega_a^{eff}}{\partial\beta_{t-1}} > 0 \Rightarrow \frac{\partial A_{t-1}^{eff}}{\partial\beta_{t-1}} > 0, \tag{1.159}
$$

which completes the proof.

(iii) As under both allocations capital is distributed equally between firms for $\beta_{t-1} = 0$, total factor productivity also coincides,

$$
A^{eff}(a_{t-1}, 0) = A(a_{t-1}, 0) = \exp\left(\frac{\alpha\theta}{\theta-1}\left(a_{t-1} + \frac{\sigma_a^2}{2}\right)\right). \tag{1.160}
$$

In the perfect information case ($\beta_{t-1} = \infty$) the efficient and market allocation also coincide,

$$\lim_{\beta_{t-1}\to\infty} A^{eff}(a_{t-1},\beta_{t-1}) = \lim_{\beta_{t-1}\to\infty} A(a_{t-1},\beta_{t-1}) = \exp\left(\frac{1}{\theta-1}\left(a_{t-1}+\frac{\sigma_a^2}{2}\right)+\frac{\sigma_a^2}{2}\right)^{\alpha\theta}. \tag{1.161}$$

For $\beta_t \in (0,\infty)$, note that the exogenous TFP component coincides under the efficient and market allocation. Therefore, it suffices to show that the allocative efficiency component is higher under the efficient allocation than under the market allocation,

$$\omega_a^{eff}\sigma_a^2 > \omega_a\left(2-\omega_a\right)\sigma_a^2 - \omega_\varepsilon^2\sigma_\varepsilon^2. \tag{1.162}$$

Using the simplification of the RHS from (1.134) leads to

$$\frac{\sigma_\varepsilon^{-2}\sigma_a^2}{\sigma_a^{-2}+\beta\sigma_\varepsilon^{-2}} > \frac{\sigma_\varepsilon^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)^2\sigma_a^2-\sigma_\varepsilon^2}{\left(\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)\right)^2}. \tag{1.163}$$

Multiplying on both sides by $\left(\sigma_a^{-2}+\beta\sigma_\varepsilon^{-2}\right)\left(\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)\right)^2$ and simplifying leads to

$$\begin{aligned}
&2\sigma_\varepsilon^{-2}\beta+2\sigma_\varepsilon^{-4}\beta > \beta\sigma_\varepsilon^{-4}+\beta\left(1+\sigma_\varepsilon^{-2}\right)^2-\sigma_a^{-2}\sigma_\varepsilon^2-\beta\\
\iff\ &2\sigma_\varepsilon^{-2}\beta+\sigma_\varepsilon^{-4}\beta > \beta+2\beta\sigma_\varepsilon^{-2}+\beta\sigma_\varepsilon^{-4}-\sigma_a^{-2}\sigma_\varepsilon^2-\beta\\
\iff\ &0 > -\sigma_a^{-2}\sigma_\varepsilon^2
\end{aligned} \tag{1.164}$$

Since $\sigma_a^2, \sigma_\varepsilon^2 > 0$ and all conversion were equivalent, the initial inequality holds and total factor productivity under the constrained efficient allocation is larger than under the market allocation for $\beta_{t-1} \in (0,\infty)$. □

Proof of Corollary 1.1. The distortion due to overconfidence vanishes if the expectations of the marginal trader $\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}$ and $\mathbb{E}\left\{A_{jt}|z_{jt}\right\}$ coincide under the following conditions:

(i) when the private signal is infinitely noisy, both expectations converge to the unconditional mean,

$$\lim_{\beta\to 0}\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\} = \lim_{\beta\to 0}\mathbb{E}\left\{A_{jt}|z_{jt}\right\} = \mathbb{E}\left\{A_{jt}\right\}. \tag{1.165}$$

When the private signal is infinitely precise, both expectations converge to the actual realization of A_{jt},

$$\lim_{\beta\to\infty} \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\} = \lim_{\beta\to\infty} \mathbb{E}\left\{A_{jt}|z_{jt}\right\} = A_{jt} \tag{1.166}$$

(ii) When the variance of firm-specific productivity shocks goes to zero, i.e., the prior becomes arbitrarily precise, both expectations converge to the mean of the distribution,

$$\lim_{\sigma_a^2\to 0} \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\} = \lim_{\sigma_a\to 0} \mathbb{E}\left\{A_{jt}|z_{jt}\right\} = \exp\left\{a_t\right\}. \tag{1.167}$$

When the variance of firm-specific productivity shocks goes to infinity, i.e., the prior becomes arbitrarily noisy, both allocations coincide because they put full weight on the price signal z_{jt},

$$\lim_{\sigma_a^{-2}\to 0} \omega_z = \lim_{\sigma_a^{-2}\to 0} \omega_z^{eff} = 1 \tag{1.168}$$

where

$$\omega_z = \frac{\beta_t\left(1+\sigma_\varepsilon^{-2}\right)}{\sigma_a^{-2}+\beta_t\left(1+\sigma_\varepsilon^{-2}\right)}, \quad \omega_z^{eff} = \frac{\beta_t\sigma_\varepsilon^{-2}}{\sigma_a^{-2}+\beta_t\sigma_\varepsilon^{-2}}. \tag{1.169}$$

(iii) When the variance of firm-specific sentiment shocks goes to zero, financial markets perfectly aggregate dispersed information as the precision of the price signal goes to infinity. In this case, both expectations converge to the true realization:

$$\lim_{\sigma_\varepsilon^2\to 0} \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\} = \lim_{\sigma_\varepsilon\to 0} \mathbb{E}\left\{A_{jt}|z_{jt}\right\} = \exp\left\{a_{jt}\right\}. \tag{1.170}$$

□

Lemma 1.2 (Joining two Normal PDFs). *Let $f(\varepsilon_{jt})$ be the pdf of $\varepsilon_{jt} \sim \mathcal{N}(\varepsilon_t, \sigma_\varepsilon^2)$ and $\phi(\cdot)$ the standard-normal pdf. Then*

$$f(\varepsilon_{jt})\,\phi(\varepsilon_{jt}) = \exp\left\{-\frac{\varepsilon_t^2}{2\left(1+\sigma_\varepsilon^2\right)}\right\}\sqrt{\frac{1}{2\pi\left(1+\sigma_\varepsilon^2\right)}}\tilde{f}(\varepsilon_{jt}) \tag{1.171}$$

where $\tilde{f}(\varepsilon_{jt})$ is the transformed pdf of $\varepsilon_{jt} \sim \mathcal{N}\left(\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}, \frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\right)$.

Proof. Write out the pdfs explicitly,

$$\phi\left(\varepsilon_{jt}\right) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{\varepsilon_{jt}^2}{2}\right\} \tag{1.172}$$

$$f\left(\varepsilon_{jt}\right) = \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}} \exp\left\{-\frac{\left(\varepsilon_{jt}-\varepsilon_t\right)^2}{2\sigma_\varepsilon^2}\right\} \tag{1.173}$$

$$f\left(\varepsilon_{jt}\right)\phi\left(\varepsilon_{jt}\right) = \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}}\frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{\varepsilon_{jt}^2}{2} - \frac{\left(\varepsilon_{jt}-\varepsilon_t\right)^2}{2\sigma_\varepsilon^2}\right\}. \tag{1.174}$$

Rearranging the term inside the exponential function,

$$\begin{aligned}
\frac{\left(\varepsilon_{jt}-\varepsilon_t\right)^2}{\sigma_\varepsilon^2} + \varepsilon_{jt} &= \frac{\varepsilon_{jt} - 2\varepsilon_{jt}\varepsilon_t + \varepsilon_t^2}{\sigma_\varepsilon^2} + \varepsilon_{jt} \\
\text{join fractions} &= \frac{\left(1+\sigma_\varepsilon^2\right)\varepsilon_{jt} - 2\varepsilon_t\varepsilon_{jt} + \varepsilon_t^2}{\sigma_\varepsilon^2} \\
\text{divide by } \left(1+\sigma_\varepsilon^2\right) &= \frac{\varepsilon_{jt} - 2\varepsilon_{jt}\frac{\varepsilon_t}{1+\sigma_\varepsilon^2} + \frac{\varepsilon_t^2}{1+\sigma_\varepsilon^2}}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} \\
\text{add and subtract} &= \frac{\varepsilon_{jt} - 2\varepsilon_{jt}\frac{\varepsilon_t}{1+\sigma_\varepsilon^2} + \frac{\varepsilon_t^2}{1+\sigma_\varepsilon^2}}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} + \frac{\left(\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2 - \left(\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} \\
\text{exchange terms} &= \frac{\varepsilon_{jt} - 2\varepsilon_{jt}\frac{\varepsilon_t}{1+\sigma_\varepsilon^2} + \left(\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} + \frac{\frac{\varepsilon_t^2}{1+\sigma_\varepsilon^2} - \left(\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} \\
\text{binomial} &= \frac{\left(\varepsilon_{jt} - \frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} + \frac{\frac{\varepsilon_t^2}{1+\sigma_\varepsilon^2} - \left(\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} \\
&= \frac{\left(\varepsilon_{jt} - \frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} + \frac{1 - \frac{1}{1+\sigma_\varepsilon^2}}{\sigma_\varepsilon^2}\varepsilon_t^2 \\
&= \frac{\left(\varepsilon_{jt} - \frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} + \frac{\frac{1+\sigma_\varepsilon^2-1}{1+\sigma_\varepsilon^2}}{\sigma_\varepsilon^2}\varepsilon_t^2 \\
&= \frac{\left(\varepsilon_{jt} - \frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}} + \frac{\varepsilon_t^2}{1+\sigma_\varepsilon^2}
\end{aligned} \tag{1.175}$$

This allows to write

$$
\begin{aligned}
f(\varepsilon_{jt})\,\phi(\varepsilon_{jt}) &= \frac{1}{\sqrt{2\pi\sigma_\varepsilon^2}}\frac{1}{\sqrt{2\pi}}\exp\left\{-\frac{\varepsilon_{jt}}{2}-\frac{(\varepsilon_{jt}-\varepsilon_t)^2}{2\sigma_\varepsilon^2}\right\}\\
&= \frac{1}{\sqrt{\sigma_\varepsilon^2}}\frac{1}{\sqrt{2\pi}}\frac{1}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}\frac{\left(\varepsilon_{jt}-\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}\right\}\exp\left\{-\frac{1}{2}\left(\frac{\varepsilon_t^2}{1+\sigma_\varepsilon^2}\right)\right\}\\
&= \frac{1}{\sqrt{\sigma_\varepsilon^2}}\sqrt{\frac{\sigma_\varepsilon^2}{2\pi(1+\sigma_\varepsilon^2)}}\frac{1}{\sqrt{2\pi\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}}\exp\left\{-\frac{1}{2}\frac{\left(\varepsilon_{jt}-\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}\right)^2}{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}\right\}\\
&* \exp\left\{-\frac{1}{2}\left(\frac{\varepsilon_t^2}{1+\sigma_\varepsilon^2}\right)\right\}\\
&= \exp\left\{-\frac{1}{2}\left(\frac{\varepsilon_t^2}{1+\sigma_\varepsilon^2}\right)\right\}\sqrt{\frac{1}{2\pi(1+\sigma_\varepsilon^2)}}\tilde{f}(\varepsilon_{jt}),
\end{aligned}
\tag{1.176}
$$

where $\tilde{f}(\varepsilon_{jt})$ is the transformed pdf of $\varepsilon_{jt}\sim\mathcal{N}\left(\frac{\varepsilon_t}{1+\sigma_\varepsilon^2},\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\right)$. □

Lemma 1.3 (Auxiliary Results Market Allocation)**.** *Denote the Bayesian weights*

$$\omega_a = \frac{\beta_{jt}(1+\sigma_\varepsilon^{-2})}{\sigma_a^{-2}+\beta_{jt}(1+\sigma_\varepsilon^{-2})}, \tag{1.177}$$

$$\omega_\varepsilon = \frac{\sqrt{\beta_{jt}}(1+\sigma_\varepsilon^{-2})}{\sigma_a^{-2}+\beta_{jt}(1+\sigma_\varepsilon^{-2})} \tag{1.178}$$

$$\omega_{z\varepsilon} = \frac{\sqrt{\beta_{jt}}\sigma_\varepsilon^{-2}}{\sigma_a^{-2}+\sqrt{\beta_{jt}}(1+\sigma_\varepsilon^{-2})}, \tag{1.179}$$

and posterior uncertainty

$$\mathbb{V} = \frac{1}{\sigma_a^{-2}+\beta_t(1+\sigma_\varepsilon^{-2})} \tag{1.180}$$

then

(i) $\omega_a^2\sigma_a^2+\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}+\mathbb{V}=\sigma_a^2$,

(ii) $\omega_a^2\sigma_a^2+\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}=\sigma_a^2-\mathbb{V}=\omega_a\sigma_a^2$,

(iii) $\frac{\omega_\varepsilon}{1+\sigma_\varepsilon^2}=\omega_{z\varepsilon}$.

Proof. (i)

$$\begin{aligned}
\omega_a^2\sigma_a^2 + \omega_\varepsilon^2 \frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2} + \mathbb{V} &= \frac{\sigma_a^2\beta^2\left(1+\sigma_\varepsilon^{-2}\right)^2}{\left(\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)\right)^2} + \frac{\beta\left(1+\sigma_\varepsilon^{-2}\right)}{\left(\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)\right)^2} \\
&\quad + \frac{\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)}{\left(\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)\right)^2} \\
&= \frac{\sigma_a^{-2}+2\beta\left(1+\sigma_\varepsilon^{-2}\right)+\sigma_a^2\beta^2\left(1+\sigma_\varepsilon^{-2}\right)^2}{\left(\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)\right)^2} \\
&= \frac{\sigma_a^{-4}+2\sigma_a^{-2}\beta\left(1+\sigma_\varepsilon^{-2}\right)+\beta^2\left(1+\sigma_\varepsilon^{-2}\right)^2}{\left(\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)\right)^2}\sigma_a^2 \\
&= \frac{\left(\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)\right)^2}{\left(\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)\right)^2}\sigma_a^2 \\
&= \sigma_a^2 \qquad (1.181)
\end{aligned}$$

(ii) The first equality follows from (i). Then

$$\begin{aligned}
\sigma_a^2 - \mathbb{V} &= \sigma_a^2 - \frac{1}{\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)} \\
&= \frac{\sigma_a^2\sigma_a^{-2}+\sigma_a^2\beta\left(1+\sigma_\varepsilon^{-2}\right)-1}{\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)} \\
&= \frac{\beta\left(1+\sigma_\varepsilon^{-2}\right)}{\sigma_a^{-2}+\beta\left(1+\sigma_\varepsilon^{-2}\right)}\sigma_a^2 \\
&= \omega_a\sigma_a^2. \qquad (1.182)
\end{aligned}$$

(iii)

$$\frac{\omega_\varepsilon}{1+\sigma_\varepsilon^2} = \frac{\omega_\varepsilon}{\sigma_\varepsilon^2\left(1+\sigma_\varepsilon^{-2}\right)} = \frac{\omega_\varepsilon\sigma_\varepsilon^{-2}}{\left(1+\sigma_\varepsilon^{-2}\right)} = \omega_{z\varepsilon}. \qquad (1.183)$$

□

Lemma 1.4. *In the symmetric equilibrium with $\beta_{ijt} = \beta_{jt}$ with $K_{t+1} < W_t$,*

(i) Sentiment shocks ε_t affect the marginal benefit of information production through

three channels,

$$\widetilde{MB}\left(\beta_{ijt},\beta_{jt}\right)\Big|_{\beta_{ijt}=\beta_{jt}}\propto\exp\left\{\underbrace{-\frac{\varepsilon_t^2}{2\left(1+\sigma_\varepsilon^2\right)}}_{\text{Information-Sensitivity}}\underbrace{-\left(\theta-1\right)\omega_{s\varepsilon}\varepsilon_t}_{\text{Relative Size}}+\underbrace{\frac{\alpha}{1-\alpha}\omega_{s\varepsilon}\varepsilon_t}_{\text{Absolute Size}}\right\}. \tag{1.184}$$

(ii) Productivity shocks a_t increase the marginal benefit of information production,

$$\widetilde{MB}\left(\beta_{ijt},\beta_{jt}\right)\Big|_{\beta_{ijt}=\beta_{jt}}\propto\exp\left\{\left(\frac{\alpha\theta-\theta+1}{\left(1-\alpha\right)\alpha\theta}+\frac{\alpha}{1-\alpha}+1\right)a_t\right\}. \tag{1.185}$$

***Proof of Lemma 1.4**.* *(i)* Assume $a_t=0$ without loss of generality. The marginal benefit to increasing β_{ijt} is

$$\widetilde{MB}\left(\beta_{ijt},\beta_{jt}\right)\Big|_{\beta_{ijt}=\beta_{jt}}=\int_{-\infty}^{\infty}g\left(a_{jt}\right)\int_{-\infty}^{\infty}f\left(\varepsilon_{jt}\right)\frac{\partial\mathcal{P}\left\{x_{ijt}=2\right\}}{\partial\beta_{ijt}}\Bigg|_{\beta_{ijt}=\beta_{jt}}\alpha A_t^{\alpha_Y} \tag{1.186}$$

$$*\frac{\left(A_{jt}-\mathbb{E}\left\{A_{jt}|s_{ijt}=z_{jt}z_{jt}\right\}\right)\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta-1}}{\left(\int_0^1\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta}dj\right)^{\frac{\theta-1}{\theta}}}K_{t+1}^{\alpha}d\varepsilon_{jt}da_{jt},$$

where $g\left(a_{jt}\right)$ is the pdf of $a_{jt}\sim\mathcal{N}\left(0,\sigma_a^2\right)$ and $f\left(\varepsilon_{jt}\right)$ is the pdf of $\varepsilon_{jt}\sim\mathcal{N}\left(\varepsilon_t,\sigma_\varepsilon^2\right)$. The most immediate effect comes from changes to aggregate investment K_{t+1}^{α}. For $\delta R_{t+1}=1$,

$$K_{t+1}^{\alpha}=\left(\alpha\delta A_t^{\alpha_Y}\left(\int_0^1\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta}dj\right)^{\frac{1}{\theta}}\right)^{\frac{\alpha}{1-\alpha}}\propto\exp\left\{\frac{\alpha}{1-\alpha}\omega_{s\varepsilon}\varepsilon_t\right\}. \tag{1.187}$$

The *Absolute Size* channel is summarized by $\exp\left\{\frac{\alpha}{1-\alpha}\omega_{s\varepsilon}\varepsilon_t\right\}$. Next, the derivative of the probability of buying at $\beta_{ijt}=\beta_{jt}$ is

$$\frac{\partial\mathcal{P}\left\{x_{ijt}=2\right\}}{\partial\beta_{ijt}}\Bigg|_{\beta_{ijt}=\beta_{jt}}=\phi\left(\varepsilon_{jt}\right)\left(\frac{\omega_{p,jt}}{\sqrt{\beta_{jt}}}a_{jt}+\frac{\varepsilon_{jt}}{2\beta_{jt}}-\frac{\omega_{\varepsilon,jt}\varepsilon_{jt}-\omega_{z\varepsilon,jt}\varepsilon_t}{\sqrt{\beta_{jt}}}+\frac{1}{2}\frac{\mathbb{V}_{jt}}{\sqrt{\beta_{jt}}}\right), \tag{1.188}$$

where $\phi(\cdot)$ is the standard-normal pdf. Combine $f(\varepsilon_{jt})$ with $\phi(\varepsilon_{jt})$ using Lemma 1.2,

$$\phi(\varepsilon_{jt})\, f(\varepsilon_{jt}) = \exp\left\{-\frac{\varepsilon_t^2}{2(1+\sigma_\varepsilon^2)}\right\}\sqrt{\frac{1}{2\pi(1+\sigma_\varepsilon^2)}}\tilde{f}(\varepsilon_{jt}), \tag{1.189}$$

where $\tilde{f}(\varepsilon_{jt})$ is the pdf of a fictional variable $\varepsilon_{jt} \sim \mathcal{N}\left(\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}, \frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\right)$. The *Information-Sensitivity* channel is summarized by $\exp\left\{-\frac{\varepsilon_t^2}{2(1+\sigma_\varepsilon^2)}\right\}$. For the rest of the proof, use

$$\varepsilon_{jt} = \sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x + \frac{\varepsilon_t}{1+\sigma_\varepsilon^2} \tag{1.190}$$

$$d\varepsilon_{jt} = \sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}dx. \tag{1.191}$$

Substitute ε_{jt} out of the terms in parenthesis for $\left.\frac{\partial \mathcal{P}\{x_{ijt}=2\}}{\partial \beta_{ijt}}\right|_{\beta_{ijt}=\beta_{jt}}$ leads to

$$\frac{\omega_{p,jt}}{\sqrt{\beta_{jt}}}a_{jt} + \left(\frac{1}{2\beta_{jt}} - \frac{\omega_{\varepsilon,jt}}{\sqrt{\beta_{jt}}}\right)\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x + \frac{\varepsilon_t}{2\beta_{jt}(1+\sigma_\varepsilon^2)} + \frac{1}{2}\frac{\mathbb{V}_{jt}}{\sqrt{\beta_{jt}}}. \tag{1.192}$$

Substituting ε_{jt} out of $\tilde{\mathbb{E}}\{A_{jt}|s_{ijt}=z_{jt}, z_{jt}\}$,

$$\begin{aligned}\tilde{\mathbb{E}}\{A_{jt}|s_{ijt}=z_{jt}, z_{jt}\} &= \exp\left\{\omega_a a_{jt} + \omega_\varepsilon \varepsilon_{jt} - \omega_{z\varepsilon}\varepsilon_t + \frac{1}{2}\mathbb{V}\right\} \\ &= \exp\left\{\omega_a a_{jt} + \omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x + \cancel{\omega_\varepsilon\frac{\varepsilon_t}{1+\sigma_\varepsilon^2} - \omega_{z\varepsilon}\varepsilon_t} + \frac{1}{2}\mathbb{V}\right\} \\ &\overset{\text{Lemma 1.3 (iii)}}{=} \exp\left\{\omega_a a_{jt} + \omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x + \frac{1}{2}\mathbb{V}\right\}.\end{aligned} \tag{1.193}$$

Substitute ε_{jt} out of the firm-specific multiplier for firm capital,

$$\begin{aligned}
&\frac{\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta-1}}{\left(\int_0^1\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta}dj\right)^{\frac{\theta-1}{\theta}}}\\
&=\exp\left\{(\theta-1)\,\omega_a a_{jt}+(\theta-1)\,\omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x+(\theta-1)\,\omega_\varepsilon\frac{\varepsilon_t}{1+\sigma_\varepsilon^2}-(\theta-1)\,\omega_\varepsilon\varepsilon_t-\frac{(\theta-1)\,\theta}{2}\left(\omega_a^2\sigma_a^2+\omega_\varepsilon^2\sigma_\varepsilon^2\right)\right\}\\
&\propto\exp\left\{(\theta-1)\,\omega_a a_{jt}+(\theta-1)\,\omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x-(\theta-1)\,\omega_\varepsilon\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\varepsilon_t\right\}\\
&=\exp\left\{(\theta-1)\,\omega_a a_{jt}+(\theta-1)\,\omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x-(\theta-1)\,\omega_{s\varepsilon}\varepsilon_t\right\}\\
&=\exp\left\{(\theta-1)\,\omega_a a_{jt}+(\theta-1)\,\omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x\right\}\exp\left\{-(\theta-1)\,\omega_{s\varepsilon}\varepsilon_t\right\},
\end{aligned}\tag{1.194}$$

where I used Lemma 1.3 (iii) repeatedly. The *Relative Size* channel is summarized through $\exp\{-(\theta-1)\,\omega_{s\varepsilon}\varepsilon_t\}$. It remains to show that there are no other terms in $\widetilde{MB}\,(\beta_{ijt},\beta_{jt})$ that depend on ε_t. It is sufficient to show that

$$\int_{-\infty}^{\infty}g\,(a_{jt})\int_{-\infty}^{\infty}\tilde{f}\,(\varepsilon_{jt})\,\frac{\left(A_{jt}-\mathbb{E}\left\{A_{jt}|s_{ijt}=z_{jt}z_{jt}\right\}\right)\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta-1}}{\left(\int_0^1\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta}dj\right)^{\frac{\theta-1}{\theta}}}d\varepsilon_{jt}da_{jt}\overset{!}{=}0.\tag{1.195}$$

Substituting ε_{jt} out leads to

$$\begin{aligned}
&\int_{-\infty}^{\infty}g\,(a_{jt})\int_{-\infty}^{\infty}\cancel{\frac{1}{\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}}}\phi\,(x)\,\frac{\left(A_{jt}-\mathbb{E}\left\{A_{jt}|s_{ijt}=z_{jt}z_{jt}\right\}\right)\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta-1}}{\left(\int_0^1\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta}dj\right)^{\frac{\theta-1}{\theta}}}\cancel{\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}}dxda_{jt}\\
&\propto\int_{-\infty}^{\infty}g\,(a_{jt})\int_{-\infty}^{\infty}\phi\,(x)\left(A_{jt}\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta-1}-\mathbb{E}\left\{A_{jt}|s_{ijt}=z_{jt}z_{jt}\right\}^{\theta}\right)dxda_{jt}\\
&=\int_{-\infty}^{\infty}g\,(a_{jt})\int_{-\infty}^{\infty}\phi\,(x)\left(\exp\left\{((\theta-1)\,\omega_a+1)\,a_{jt}+(\theta-1)\,\omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x+\frac{(\theta-1)}{2}\mathbb{V}\right\}\right.\\
&\quad\left.-\exp\left\{\theta\omega_a a_{jt}+\theta\omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x+\frac{\theta}{2}\mathbb{V}\right\}\right)dxda_{jt}\\
&\propto\int_{-\infty}^{\infty}g\,(a_{jt})\int_{-\infty}^{\infty}\phi\,(x)\left(\exp\left\{((\theta-1)\,\omega_a+1)\,a_{jt}+(\theta-1)\,\omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x-\frac{\mathbb{V}}{2}\right\}\right.\\
&\quad\left.-\exp\left\{\theta\omega_a a_{jt}+\theta\omega_\varepsilon\sqrt{\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}}x\right\}\right)dxda_{jt}\\
&=\int_{-\infty}^{\infty}g\,(a_{jt})\left(\exp\left\{((\theta-1)\,\omega_a+1)\,a_{jt}+\frac{(\theta-1)^2\omega_\varepsilon^2}{2}\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}-\frac{\mathbb{V}}{2}\right\}-\exp\left\{\theta\omega_a a_{jt}+\frac{\theta^2\omega_\varepsilon^2}{2}\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\right\}\right)da_{jt}\\
&=\exp\left\{\frac{((\theta-1)\,\omega_a+1)^2}{2}\sigma_a^2+\frac{(\theta-1)^2\omega_\varepsilon^2}{2}\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}-\frac{\mathbb{V}}{2}\right\}-\exp\left\{\frac{\theta^2\omega_a^2}{2}\sigma_a^2+\frac{\theta^2\omega_\varepsilon^2}{2}\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\right\}.
\end{aligned}\tag{1.196}$$

It remains to show that

$$\left(\left(\theta-1\right)\omega_a+1\right)^2\sigma_a^2+\left(\theta-1\right)^2\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}-\mathbb{V}\overset{!}{=}\theta^2\omega_a^2\sigma_a^2+\theta^2\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}. \tag{1.197}$$

Using Lemma 1.3 (i), the LHS is equal to

$$\begin{aligned}
&\left(\left(\theta^2-2\theta+1\right)\omega_a^2+2\left(\theta-1\right)\omega_a\cancel{+1}\right)\sigma_a^2+\left(\theta^2-2\theta+1\right)\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\cancel{-\sigma_a^2}+\omega_a^2\sigma_a^2+\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\\
&=\left(\left(\theta^2-2\theta+2\right)\omega_a^2+2\left(\theta-1\right)\omega_a\right)\sigma_a^2+\left(\theta^2-2\theta+2\right)\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\\
&=2\left(\theta-1\right)\omega_a\sigma_a^2+\left(\theta^2+2\left(1-\theta\right)\right)\left(\omega_a^2\sigma_a^2+\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\right)\\
&\overset{\text{Lemma 1.3 (ii)}}{=}2\left(\theta-1\right)\omega_a\sigma_a^2+\left(\theta^2+2\left(1-\theta\right)\right)\left(\omega_a\sigma_a^2\right)\\
&=\theta^2\omega_a\sigma_a^2.
\end{aligned} \tag{1.198}$$

Using Lemma 1.3 (ii), the RHS is equal to

$$\theta^2\left(\omega_a^2\sigma_a^2+\omega_\varepsilon^2\frac{\sigma_\varepsilon^2}{1+\sigma_\varepsilon^2}\right)=\theta^2\omega_a\sigma_a^2. \tag{1.199}$$

Combining both confirms the conjecture. The marginal benefit of information production depends on ε_t only through the multiplicative effects in (1.187), (1.189) and (1.194), such that

$$\left.\widetilde{MB}\left(\beta_{ijt},\beta_{jt}\right)\right|_{\beta_{ijt}=\beta_{jt}}\propto\exp\left\{\underbrace{\frac{\alpha}{1-\alpha}\omega_{s\varepsilon}\varepsilon_t}_{\textit{Absolute Size}}\ \underbrace{-\frac{\varepsilon_t^2}{2\left(1+\sigma_\varepsilon^2\right)}}_{\textit{Information Sensitivity}}\ \underbrace{-\left(\theta-1\right)\omega_{s\varepsilon}\varepsilon_t}_{\textit{Relative Size}}\right\}. \tag{1.200}$$

(ii) Follow the same strategy as in *(i)* and use the same expression for the marginal benefit of information production (1.186). Start with the expressions for aggregate investment, K_{t+1}^{α}, and productivity $A_t^{\alpha_Y}$ in (1.186). For $\delta R_{t+1}=1$, they are equal to

$$A_t^{\alpha_Y}K_{t+1}^{\alpha}=A_t^{\alpha_Y}\left(\alpha\delta A_t^{\alpha_Y}\left(\int_0^1\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta}dj\right)^{\frac{1}{\theta}}\right)^{\frac{\alpha}{1-\alpha}}, \tag{1.201}$$

Recall that $\alpha_Y = \frac{\alpha\theta-\theta+1}{\alpha\theta}$. Then

$$A_t^{\alpha_Y} \propto \exp\left\{\frac{\alpha\theta-\theta+1}{\alpha\theta}a_t\right\} \tag{1.202}$$

$$\left(\int_0^1 \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta} dj\right)^{\frac{1}{\theta}} \propto \exp\left\{a_t\right\}. \tag{1.203}$$

Putting both together yields

$$\begin{aligned} A_t^{\alpha_Y}K_{t+1}^{\alpha} &\propto \exp\left\{\frac{\alpha\theta-\theta+1}{\alpha\theta}\left(1+\frac{\alpha}{1-\alpha}\right)a_t+\frac{\alpha}{1-\alpha}a_t\right\} \\ &= \exp\left\{\frac{\alpha\theta-\theta+1}{(1-\alpha)\,\alpha\theta}a_t+\frac{\alpha}{1-\alpha}a_t\right\} \\ &= \exp\left\{\left(\frac{\alpha\theta-\theta+1}{(1-\alpha)\,\alpha\theta}+\frac{\alpha}{1-\alpha}\right)a_t\right\}. \end{aligned} \tag{1.204}$$

Again, using substitution with

$$a_{jt} = \sqrt{\sigma_a^2}y + a_t \tag{1.205}$$

$$da_{jt} = \sqrt{\sigma_a^2}dy \tag{1.206}$$

it follows that

$$A_{jt} \propto \exp\left\{a_t\right\} \tag{1.207}$$

$$\mathbb{E}\left\{A_{jt}|s_{ijt}=z_{jt}z_{jt}\right\} \propto \exp\left\{a_t\right\}, \tag{1.208}$$

which yields

$$\frac{\left(A_{jt}-\mathbb{E}\left\{A_{jt}|s_{ijt}=z_{jt}z_{jt}\right\}\right)\tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta-1}}{\left(\int_0^1 \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt}=z_{jt},z_{jt}\right\}^{\theta} dj\right)^{\frac{\theta-1}{\theta}}} \propto \exp\left\{a_t\right\}. \tag{1.209}$$

The change in the buying probability does *not* depend on a_t

$$\left.\frac{\partial\mathcal{P}\left\{x_{ijt}=2\right\}}{\partial\beta_{ijt}}\right|_{\beta_{ijt}=\beta_{jt}} = \phi\left(\varepsilon_{jt}\right)\left(\frac{\omega_p}{\sqrt{\beta_{jt}}}\sqrt{\sigma_a^2}y+\frac{\varepsilon_{jt}}{2\beta_{jt}}-\frac{\omega_{\varepsilon}\varepsilon_{jt}-\omega_{z\varepsilon}\varepsilon_t}{\sqrt{\beta_{jt}}}+\frac{\mathbb{V}_{jt}}{2\sqrt{\beta_{jt}}}\right). \tag{1.210}$$

Since no other terms depend on a_t and I substituted a_{jt} out, it follows that

$$\widetilde{MB}\left(\beta_{ijt},\beta_{jt}\right)\Big|_{\beta_{ijt}=\beta_{jt}} \propto \exp\left\{\left(\frac{\alpha\theta-\theta+1}{(1-\alpha)\,\alpha\theta}+\frac{\alpha}{1-\alpha}+1\right)a_t\right\} \tag{1.211}$$

□

Proof of Proposition 1.4*.* The cutoff can be derived by using the result from Lemma 1.4 (i) and taking the derivative with respect to ε_t to the following expression,

$$\frac{\partial}{\partial\varepsilon_t}\left(-\frac{\varepsilon_t^2}{2\left(1+\sigma_\varepsilon^2\right)}-(\theta-1)\,\omega_{s\varepsilon}\varepsilon_t+\frac{\alpha}{1-\alpha}\omega_{s\varepsilon}\varepsilon_t\right)\overset{!}{=}0. \tag{1.212}$$

Denote $\bar{\varepsilon}$ as the value of ε_t for which the above expression is maximized. Then,

$$\begin{aligned}
-\frac{\bar{\varepsilon}}{1+\sigma_\varepsilon^2}-(\theta-1)\,\omega_{s\varepsilon}+\frac{\alpha}{1-\alpha}\omega_{s\varepsilon}=0&\\
\iff\ \frac{\bar{\varepsilon}}{1+\sigma_\varepsilon^2}=\left(\frac{1}{1-\alpha}-\theta\right)\omega_{s\varepsilon}&\\
\iff\ \bar{\varepsilon}=\left(1+\sigma_\varepsilon^2\right)\left(\frac{1}{1-\alpha}-\theta\right)\omega_{s\varepsilon},&
\end{aligned} \tag{1.213}$$

where $\omega_{s\varepsilon}=\frac{\sqrt{\beta_t}}{\sigma_a^{-2}+\beta_t\left(1+\sigma_\varepsilon^{-2}\right)}$. For $\varepsilon_t<\bar{\varepsilon}$, information production β_t is increasing in ε_t. For $\varepsilon_t>\bar{\varepsilon}$, information production β_t is decreasing in ε_t. □

Proof of Proposition 1.5*.* Follows from Lemma 1.4 (ii). □

Proof of Proposition 1.6*.* *(i)* Using the result from Proposition 1.4 (ii) and the assumption that $\theta>\frac{1}{1-\alpha}$, it must be that positive sentiment shocks crowd out information production. Moreover, as $\beta^*<\frac{\sigma_a^{-2}}{1+\sigma_\varepsilon^{-2}}$, it must be that the pass-through of aggregate sentiment shocks $\omega_{s\varepsilon}$ decreases in information production. Therefore, the pass-through is smaller when information is endogenous than if information production is fixed at β^*. As a result, sentiment shocks are dampened by information production in financial markets, as less precise information by itself leads to less investment and lowers the pass-through of sentiment shocks.

(ii) $\lim_{\varepsilon_t\to\infty}\sqrt{\beta_t(\varepsilon_t)}\varepsilon_t=0$ guarantees that the pass-through of sentiment shocks goes faster to zero than the sentiment shock goes to infinity, i.e., the direct effect of sentiment shocks on investment disappears as shocks become arbitrarily large. Moreover,

Lemma 1.4 (i) shows that through the information-sensitivity effect $\lim_{\varepsilon_t \to \infty} \beta_t(\varepsilon_t) = 0$. □

Proof of Corollary 1.2. Follows directly from Proposition 1.6 (ii). □

Proof of Proposition 1.7. Follows from Proposition 1.5 and Assumption 1.2. An increase in aggregate productivity encourages more information production, which also leads to more investment. As a result, productivity shocks are amplified. □

Proof of Proposition 1.8. Since $\widetilde{MB}(\beta_{ijt}, \beta_{jt}) \propto \kappa_H - \kappa_L$ whereas $MB^{SP}(\beta_t)$ is not a function of position limits $\{\kappa_H, \kappa_L\}$, it must that $\widetilde{MB}(\beta_{ijt}, \beta_{jt})\Big|_{\beta_{ijt}=\beta_{jt}=\beta_t} \neq MB^{SP}(\beta_t)$ for almost all values of β_t. Therefore, the information production in the competitive economy and social planner allocation do not coincide almost everywhere. □

Proof of Corollary 1.3. The marginal benefit of information production after applying the tax/subsidy $\tau(a_{j0}, z_{j0})$ is

$$\widetilde{MB}(\beta_{ij0}, \beta_{j0}) \propto \tilde{\mathbb{E}}\left\{\frac{\partial \mathcal{P}\{x_{ij0} = 2\}}{\partial \beta_{ij0}}\Big(\tau(a_{j0}, z_{j0})\Pi_{j1} - \tilde{\mathbb{E}}\{\tau(a_{j0}, z_{j0})\Pi_{j1} | s_{ij0} = z_{j0}, z_{j0}\}\Big)\right\}. \quad (1.214)$$

Assume that a tax fulfills the following conditions

$$\tau(a_{j0}, z_{j0}) \geq (\leq) 1 \iff \Pi_{j1} \geq \tilde{\mathbb{E}}\{\Pi_{j1} | s_{ij0} = z_{j0}, z_{j0}\} \text{ and } \frac{\partial \mathcal{P}\{x_{ij0} = 2\}}{\partial \beta_{ij0}} \geq 0 \quad (1.215)$$

$$\tau(a_{j0}, z_{j0}) \leq (\geq) 1 \iff \Pi_{j1} \leq \tilde{\mathbb{E}}\{\Pi_{j1} | s_{ij0} = z_{j0}, z_{j0}\} \text{ and } \frac{\partial \mathcal{P}\{x_{ij0} = 2\}}{\partial \beta_{ij0}} \leq 0 \quad (1.216)$$

and the inequalities are strict for at least some $\{a_{j0}, z_{j0}\}$. In the first case, $\tau(a_{j0}, z_{j0}) \geq 1$ whenever the trading rents $\Pi_{j1} - \tilde{\mathbb{E}}\{\tau(a_{j0}, z_{j0})\Pi_{j1} | s_{ij0} = z_{j0}, z_{j0}\}$ are positive and producing additional information leads to an increase of the probability of trading in that state. Therefore, the gains of buying when it is profitable are increased, which encourages information production. The same reasoning applies for $\tau(a_{j0}, z_{j0}) \leq 1$, as losses become more painful, increasing the incentive for information production. This set of taxes increases $\widetilde{MB}(\beta_{ij0}, \beta_{j0})$ and encourages information production.

The reverse reasoning applies when $\tau(a_{j0}, z_{j0}) \geq 1$ for losses and $\tau(a_{j0}, z_{j0}) \leq 1$ for gains, which leads to a decrease in $\widetilde{MB}(\beta_{ij0}, \beta_{j0})$. As a result, gains and losses are reduced, which discourages information production. □

Proof of Proposition 1.9. Let the social planner buy $d^{SP} \in (-1, 1)$ units of shares in all markets. The market-clearing condition for market j becomes

$$2\left(1 - \Phi\left(\sqrt{\beta_{j0}}\left(\hat{s}\left(P_{j0}\right) - a_{j0}\right) - \varepsilon_{j0}\right)\right) = 1 - d^{SP}, \tag{1.217}$$

Keeping position limits fixed, the social planner's demand d^{SP} changes the identity of the marginal trader. If the social planner purchases more assets, the marginal trader becomes more optimistic on average. The threshold signal becomes,

$$\hat{s}\left(P_{j0}, d^{SP}\right) = a_{j0} + \frac{\varepsilon_{j0} + \Phi^{-1}\left(\frac{1+d^{SP}}{2}\right)}{\sqrt{\beta_{j0}}}. \tag{1.218}$$

It follows immediately that asset purchases or sales with $d^{SP} = 2\Phi(-\varepsilon_0) - 1$ ensure that the marginal trader holds unbiased beliefs,

$$\hat{s}\left(P_{j0}, d^{SP}\right) = a_{j0} + \frac{\varepsilon_{j0} - \varepsilon_0}{\sqrt{\beta_{j0}}}. \tag{1.219}$$

It follows that prices are unbiased and aggregate investment is at a level as if the sentiment shock was absent.

Traders expect to buy in equilibrium whenever $s_{ijt} > \hat{s}\left(P_{j0}, d^{SP}\right)$. Asset purchases/sales set the threshold $\hat{s}\left(P_{j0}, d^{SP}\right)$ at a level as if the aggregate sentiment shock was $\varepsilon_0 = 0$, effectively undoing any change to the incentive to produce information. Because the trader thinks that she is unaffected by the sentiment shock and the markets become unaffected by the sentiment shock due to asset purchases / sales d^{SP}, also the information production decision reverts to the level without the aggregate sentiment shock. □

Proof of Proposition 1.10. (i) Denote $k_{jt+1} = \ln K_{jt+1}$. Using (1.25) allows to write the variance of the log of firm capital stocks as

$$Var\left(k_{jt+1}\right) = \theta^2 Var\left(\ln \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\}\right)$$
$$= \frac{\theta^2}{2}\left(\omega_a^2\sigma_a^2 + \omega_\varepsilon^2\sigma_\varepsilon^2\right). \tag{1.220}$$

Which can be expressed as

$$\omega_a^2\sigma_a^2 + \omega_\varepsilon^2\sigma_\varepsilon^2 = \frac{\beta_{t-1}^2\left(1+\sigma_\varepsilon^{-2}\right)^2\sigma_a^2 + \beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)^2\sigma_\varepsilon^2}{\left(\sigma_a^{-2} + \beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right)\right)^2}$$
$$\propto \frac{\beta_{t-1}^2\sigma_a^2 + \beta_{t-1}\sigma_\varepsilon^2}{\sigma_a^{-4} + 2\sigma_a^{-2}\beta_{t-1}\left(1+\sigma_\varepsilon^{-2}\right) + \beta_{t-1}^2\left(1+\sigma_\varepsilon^{-2}\right)^2}. \tag{1.221}$$

Taking the derivative with respect to β_{t-1} and dropping the denominator leads to the simplified expression

$$2\beta_{t-1}\sigma_a^{-2} + \sigma_a^{-4}\sigma_\varepsilon^2 + \beta_{t-1}^2\left(1-\sigma_\varepsilon^2\right)\left(1+\sigma_\varepsilon^{-2}\right) \tag{1.222}$$

which is positive for all values of β_{t-1} for $\sigma_\varepsilon^2 \leq 1$.

(ii) Denote $\Delta k_{jt+1} = k_{jt+1} - k_{jt}$. Then

$$\Delta k_{jt+1} = \Delta\theta \ln \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\} + \Delta \ln \int_0^1 \tilde{\mathbb{E}}\left\{A_{jt}|s_{ijt} = z_{jt}, z_{jt}\right\} dj + \Delta K_{t+1}. \tag{1.223}$$

Deriving the variance of Δk_{jt+1} across firms yields

$$Var\left(\Delta k_{jt+1}\right) = \theta^2\left(\omega_a^2\sigma_a^2 + \omega_\varepsilon^2\sigma_\varepsilon^2\right) \tag{1.224}$$

which is monotonically increasing in β_{t-1} for $\sigma_\varepsilon^2 \leq 1$ as in (i). □

Proof of Corollary 1.4. The derivation is analogous to Lemma 1.4 (i). The reduction in information-sensitivity of the trading decision dominates a possible increase in investment if the sectoral elasticity of substitution θ_n is large enough as in Proposition 1.4. Therefore, a positive sentiment shock can discourage information production and increase capital misallocation. □

Proof of Corollary 1.5. Assume first that aggregate investment is fixed ($\delta \to \infty$). Then, a positive shock in sector / country n leads to an increase in investment in sector / country n analogously to (1.25). Because the aggregate level of investment is fixed, it must be that investment in other sectors / countries decreases, decreasing the incentive for information production similarly to the *absolute scale* channel in Lemma 1.4. However, if aggregate investment is variable, then an expansion of one sector / countries shrinks investment in other sectors / countries if their respective goods are sufficiently close substitutes ($\theta \to \infty$). In this case, the positive shock discourages investment in non-shocked sectors / countries, lowering information production and increasing capital misallocation in non-shocked sectors / countries. □

Proof of Proposition 1.11. The social planner's allocation is given by equalizing the marginal products of capital for each firm given the market signals $\{z_{jt}\}$. The maximization problem of the social planner for firm capital allocation is therefore

$$\max_{K_{j1}} \mathbb{E}\left\{\left(\int_0^1 Y_{j1}^{\frac{\theta-1}{\theta}} dj\right)^{\frac{\alpha\theta}{\theta-1}} |z_{jt}\right\} - R_1^{SP} K_{j1}, \tag{1.225}$$

for some interest rate R_1^{SP}. The resulting first-order condition for firm capital is

$$K_{j1}^{SP} = \left(\alpha Y_1^{\alpha_Y} \frac{\mathbb{E}\{A_{j0}|z_{j0}\}}{R_1^{SP}}\right)^{\theta}. \tag{1.226}$$

Integrating on both sides yields

$$R_1^{SP} = \alpha Y_1^{\alpha_Y} \left(\int_0^1 \mathbb{E}\{A_{j0}|z_{j0}\}^{\theta} dj\right)^{\frac{1}{\theta}} \left(K_1^{SP}\right)^{-\frac{1}{\theta}}. \tag{1.227}$$

Substituting R_1^{SP} out of K_{j1}^{SP} yields (1.114). Following the same steps as in the proof for Proposition 1.3, leads to

$$Y_1^{SP} = A_0^{SP} K_1^{\alpha}, \quad \text{where } A_0^{SP} = \left(\int_0^1 \mathbb{E}\{A_{j0}|z_{j0}\}^{\theta} dj\right)^{\frac{\alpha}{\theta-1}}, \tag{1.228}$$

as in (1.115). Substituting Y_1^{SP} out of the expression for R_1^{SP} then leads to the interest rate (1.116). Consumption follows using (1.117) in (1.110). Finally, taking K_1^{SP} as given and plugging aggregate capital investment in in Y_1^{SP} in (1.115) in (SPFull), (1.118) follows after taking the derivative of (1.115) with respect to β_0. □

Proof of Proposition 1.12. I will show that the decentralized allocations coincide with the social planner's allocations. The proof follows the same steps as the derivation of the equilibrium in the main section. Households receive from firm j the dividend

$$\hat{\Pi}_{j1} = \tau^{Bias}(z_{j0})\,\Pi_{j1} = \frac{\mathbb{E}\{A_{j0}|z_{j0}\}}{\tilde{\mathbb{E}}\{A_{j0}|s_{ij0} = z_{j0}, z_{j0}\}}\alpha Y_1^{\alpha_Y} A_{j0} K_{j1}^{\frac{\theta-1}{\theta}} \tag{1.229}$$

and the marginal trader expects the dividend to be

$$\tilde{\mathbb{E}}\left\{\hat{\Pi}_{j1}|s_{ij0} = z_{j0}, z_{j0}\right\} = \alpha Y_1^{\alpha_Y}\mathbb{E}\{A_{j0}|z_{j0}\} K_{j1}^{\frac{\theta-1}{\theta}}. \tag{1.230}$$

The price is using $P_{j0} = K_{j1}$,

$$P_{j0} = \frac{1}{R_1}\tilde{\mathbb{E}}\left\{\hat{\Pi}_{j1}|s_{ij0} = z_{j0}, z_{j0}\right\} = \left(\frac{1}{R_1}\alpha Y_1^{\alpha_Y}\mathbb{E}\{A_{j0}|z_{j0}\}\right)^{\theta}. \tag{1.231}$$

This allows to express expected dividends substituting the capital stock,

$$\tilde{\mathbb{E}}\left\{\hat{\Pi}_{j1}|s_{ij0} = z_{j0}, z_{j0}\right\} = \left(\frac{1}{R_1}\right)^{\theta-1}\left(\alpha Y_1^{\alpha_Y}\mathbb{E}\{A_{j0}|z_{j0}\}\right)^{\theta}, \tag{1.232}$$

which is then used for the interest rate R_1

$$\begin{aligned} R_1 &= \frac{\int_0^1 \tilde{\mathbb{E}}\left\{\hat{\Pi}_{j1}|s_{ij0} = z_{j0}, z_{j0}\right\} dj}{\int_0^1 P_{j0} dj} \\ &= \left(\frac{1}{R_1}\right)^{\theta-1}\left(\alpha Y_1^{\alpha_Y}\right)^{\theta}\int_0^1 \mathbb{E}\{A_{j0}|z_{j0}\}^{\theta} dj K_1^{-1} \\ \iff R_1 &= \alpha Y_1^{\alpha_Y}\left(\int_0^1 \mathbb{E}\{A_{j0}|z_{j0}\}^{\theta} dj\right)^{\frac{1}{\theta}} K_1^{-\frac{1}{\theta}}, \end{aligned} \tag{1.233}$$

Using this result again for the price yields

$$K_{j1}^{DE} = \frac{\mathbb{E}\{A_{j0}|z_{j0}\}^{\theta}}{\int_0^1 \mathbb{E}\{A_{j0}|z_{j0}\}^{\theta} dj} K_1 = K_{j1}^{SP}. \tag{1.234}$$

Plugging this into the expression for the interest rate R_1 and substituting $Y_1^{\alpha_Y}$ leads to

$$R_1^{DE} = \alpha \left(\int_0^1 \mathbb{E} \left\{ A_{j0} | z_{j0} \right\}^\theta dj \right)^{\frac{1}{\theta}} K_1^{\alpha-1} = R_1^{SP}. \tag{1.235}$$

This result also leads directly to $K_1^{DE} = K_1^{SP}$. Finally, the first-order condition for information production of trader ij is

$$\begin{aligned} \widetilde{MB}\left(\beta_{ij0}, \beta_{j0}\right) &= \frac{\partial IA^{DE}}{\partial \beta_{ij0}} \\ &= \tau^{Info}\left(\beta_{ij0}, \beta_{j0}\right) \frac{\partial IA\left(\beta_{ij0}\right)}{\partial \beta_{ij0}} \\ &= \frac{\widetilde{MB}\left(\beta_{ij0}, \beta_{j0}\right)}{\left.\frac{\partial Y_1}{\partial \beta_0}\right|_{\beta_0=\beta_{ij0}}} \frac{\partial IA\left(\beta_{ij0}\right)}{\partial \beta_{ij0}} \\ \iff \left.\frac{\partial Y_1}{\partial \beta_0}\right|_{\beta_0=\beta_{ij0}} &= \frac{\partial IA\left(\beta_{ij0}\right)}{\partial \beta_{ij0}} \end{aligned} \tag{1.236}$$

which is the same first-order condition as in (1.118) and therefore $\beta_0^{DE} = \beta_0^{SP}$. □

Proof of Corollary 1.6. *(i)* First, denote $\omega_{s\varepsilon} = \frac{\sqrt{\beta_0}}{\sigma_a^{-2}+\beta_0\left(1+\sigma_\varepsilon^{-2}\right)}$ as the weight on the correlated noise in the private signal. The transaction tax/subsidy $\tau^{Trans}(\varepsilon_0) = \exp\left\{-\omega_{s\varepsilon}\varepsilon_0\right\}$ leads to traders paying P_{j0} but only $\tau^{Trans} P_{j0}$ is collected by the firm. The transaction tax/subsidy is aimed to stabilize aggregate asset prices with respect to aggregate sentiment shocks. It is a tax when traders are exuberant and a subsidy when they are depressed. The proof follows the same steps as for Proposition 1.12 with the difference that $K_{j1} = \tau^{Trans}(\varepsilon_0) P_{j0}$ and therefore capital of firm j is

$$K_{j1} = \left(\frac{\tau^{Trans}(\varepsilon_0)}{R_1} \alpha Y_1^{\alpha_Y} \tilde{\mathbb{E}} \left\{ A_{j0} | s_{ij0} = z_{j0}, z_{j0} \right\} \right)^\theta \tag{1.237}$$

Following the same steps as before, the interest rate is

$$R_1 = \alpha Y_1^{\alpha_Y} \left(\int_0^1 \tilde{\mathbb{E}} \left\{ A_{j0} | s_{ij0} = z_{j0}, z_{j0} \right\}^\theta dj \right)^{\frac{1}{\theta}} \tau^{Trans}(\varepsilon_0) K_1^{-\frac{1}{\theta}}. \tag{1.238}$$

Since $\left(\int_0^1 \tilde{\mathbb{E}}\left\{A_{j0}|s_{ij0}=z_{j0},z_{j0}\right\}^\theta dj\right)^{\frac{1}{\theta}} \propto \exp\left\{\omega_{s\varepsilon}\varepsilon_0\right\}$, it follows that the transaction tax/subsidy $\tau^{Trans}(\varepsilon_0) = \exp\left\{-\omega_{s\varepsilon}\varepsilon_0\right\}$ keeps the interest rate R_1 from moving with the aggregate sentiment shock ε_0 and stabilizes, therefore, aggregate investment with respect to sentiment shocks.

(ii) Similarly, allow now the transaction tax to vary with the share price,

$$\tau^{Trans}\left(P_{j0}\right) = \frac{\mathbb{E}\left\{A_{j0}|z_{j0}\right\}}{\tilde{\mathbb{E}}\left\{A_{j0}|s_{ij0}=z_{j0},z_{j0}\right\}}. \tag{1.239}$$

Same as before, the traders pays P_{j0} but only $\tau^{Trans}\left(P_{j0}\right)P_{j0}$ is collected by the firm. Using $K_{j1} = \tau^{Trans}\left(P_{j0}\right)P_{j0}$, capital of firm j is then equal to

$$\begin{aligned} K_{j1} &= \left(\frac{\tau^{Trans}\left(P_{j0}\right)}{R_1}\alpha Y_1^{\alpha_Y}\tilde{\mathbb{E}}\left\{A_{j0}|s_{ij0}=z_{j0},z_{j0}\right\}\right)^\theta \\ &= \left(\frac{1}{R_1}\alpha Y_1^{\alpha_Y}\mathbb{E}\left\{A_{j0}|z_{j0}\right\}\right)^\theta \end{aligned} \tag{1.240}$$

Integrating on both sides leads to

$$R_1 = \alpha Y_1^{\alpha_Y}\left(\int_0^1 \mathbb{E}\left\{A_{j0}|z_{j0}\right\}^\theta dj\right)^{\frac{1}{\theta}} K_{j1}^{-\frac{1}{\theta}} \tag{1.241}$$

Using this interest in (1.240) leads to the capital allocation

$$K_{j1} = \frac{\mathbb{E}\left\{A_{j0}|z_{j0}\right\}^\theta}{\int_0^1 \mathbb{E}\left\{A_{j0}|z_{j0}\right\}^\theta dj}K_1 \tag{1.242}$$

which coincides with the efficient or social planner allocation (1.114). It follows the transaction tax $\tau^{Trans}\left(P_{j0}\right)$ corrects for the mispricing between firms and as well as stabilizes aggregate investment with respect to the aggregate sentiment shock. □

Chapter 2

OVERCONFIDENCE AND INFORMATION ACQUISITION IN FINANCIAL MARKETS

2.1 Introduction

Financial markets are among the most efficient aggregators of dispersed information in the modern economy. When traders learn new information, they seek to profit by buying when the information is positive and selling when it is negative. Such informed speculation ensures that prices quickly incorporate new information, which decreases the profitability of trading on private information in the first place. However, not all trading occurs for fundamental reasons, e.g., liquidity trading, which can help mask informed trading and maintain the incentives for costly information acquisition.[1] Although such noise trading allows financial markets to be liquid, the mechanisms behind noise trading are little understood.

[1] In the absence of market-wide noise, asset prices can become fully revealing and destroy the incentive for costly information acquisition. This is called the Grossman-Stiglitz paradox (Grossman and Stiglitz 1980).

In this paper, I study the role of overconfidence in the form of correlation neglect[2] for incentivizing costly information acquisition in a financial market that aggregates dispersed information. For this purpose, I develop a model along the lines of Grossman and Stiglitz (1980) and Albagli, Hellwig, and Tsyvinski (2021). In such models, privately informed arbitrageurs seek to profit from mispricing. Usually, unobservable variations in the asset supply keep prices from being fully revealing. Instead, I assume that traders are overconfident about their signals' informativeness as they perceive the signal to be overly independent of other information sources. I choose this approach for two reasons. First, overconfidence[3] and correlation neglect[4] are well-documented behavioral biases that have been observed for traders, financial managers, and experiment participants.[5] Second, assuming that traders are overconfident yields a more tractable model, addressing a significant concern in this literature. The gained tractability allows exploring additional settings while keeping the information structure normal.

The overconfidence assumption affects how traders process their private information, which contains two sources of noise. The first component is idiosyncratic, which can stem from private sources of information or the imperfect understanding of public information. The second noise component is perfectly correlated across traders and can be viewed as a form of *sentiment*. Therefore, traders can be collectively optimistic because the asset's fundamental value increased or because market sentiment is exuberant. In this model, overconfidence takes the form of correlation neglect, i.e., traders believe that their signal is more idiosyncratic than it truly is. This bias leads to an overweighting of private information, which traders believe contains information not already reflected in the market.

I show that an infinitesimal amount of overconfidence is sufficient to generate trade when private signals are exogenous. Although the price aggregates all private information perfectly from the perspective of an uninformed observer, a small amount of over-

[2]Correlation neglect arises when agents learn from multiple signals but do not fully account for the signals' correlation structure. For instance, when two pieces of information are perceived as independent, although they stem from the same, possibly biased, source.

[3]See Biais et al. (2005), Allen and Evans (2005), Malmendier and Tate (2005), and Ben-David, Graham, and Harvey (2013).

[4]See Brandts, Giritligil, and Weber (2015), Eyster and Weizsäcker (2016), Eyster et al. (2018), Grimm and Mengel (2020), Enke and Zimmermann (2019), and Chandrasekhar, Larreguy, and Xandri (2020).

[5]See Glaser and Weber (2010) for a broad overview on the evidence on overconfidence and Hirshleifer (2015) for a survey on behavioral finance.

confidence lets traders not discard their private information. However, when information is costly, substantial amounts of overconfidence are necessary to motivate information acquisition. The reason is that the traders' private signals serve a dual function. First, traders learn about an asset's fundamental value. Second, traders learn about correlated noise from their private signal, which is useful for filtering information learned from the asset price. As the trader exerts effort to reduce both idiosyncratic and correlated noise in her signal, she learns more about the asset's fundamental value and less about correlated noise. Therefore, traders can choose not to exert effort to maximize their information about the correlated noise, which amounts to *free-riding* on the information acquisition of other traders. For the *free-riding* strategy to be unattractive, traders must believe that their private signal is relatively uninformative about correlated noise, which is the case for strongly overconfident traders. Because free-riding cannot be an equilibrium strategy, a large amount of overconfidence is necessary for the equilibrium to exist.

I use the model to study several applications. First, the model accommodates heterogeneity in financial markets when two groups of traders differ in their information technologies and degrees of overconfidence. For example, including rational and boundedly-rational traders allows to model noise in financial markets more carefully. Second, the model's approach to noise trading is especially amenable for macroeconomic applications in which aggregate resource constraints have to be observed. The reason is that boundedly-rational traders are still utility-maximizing and observe budget constraints, whereas noise traders exogenously add and remove resources from the economy.

Finally, I study a setting in which traders' demands depend on their trading capital and the asset price instead of imposing exogenous position limits. Such funding constraints leave the ability of financial markets to aggregate information unchanged while dampening the effect of information change on the price. Because traders learn only about the next dividend payment and have limited trading capital, a higher resale price, e.g., due to lower interest rates, limits their ability to exploit their information and discourages information acquisition.

When traders face funding constraints, disagreement plays an important role in determining the asset price. If prices are high on average, most traders need to buy to clear the market, leading to a relatively pessimistic price-setting trader. Higher disagreement leads to more extreme beliefs, which lowers the price as the price-setting trader becomes more pessimistic. The reverse is true when prices are low on average, as disagreement increases asset prices by exacerbating the price-setting trader's optimism. Due to common priors, disagreement is hump-shaped in information precision. Therefore, a change in information precision also affects the *average* price level. Whether the price level increases or falls depends on the initial precision and on whether trading capital is relatively scarce or abundant.

2.1.1 Literature

There is a large literature studying information in financial markets that goes back to seminal papers such as Grossman and Stiglitz (1980), Hellwig (1980), and Diamond and Verrecchia (1981). The most closely related paper is Albagli, Hellwig, and Tsyvinski (2021), who provide a framework in which traders learn about a variable that summarizes the fundamental value of an asset, e.g., earnings or the probability of default, instead of learning about the dividend directly. The result is that the model allows studying assets with non-linear payoffs (e.g., equity or debt). This paper adds to the literature by considering overconfidence and correlation neglect as a source of noise trading, which yields an especially tractable model.

The paper belongs to the literature that studies the role of behavioral biases in financial markets. A standard result in the CARA-normal framework is that overconfidence, i.e., the perception that the trader's information is more accurate than it truly is, *improves* the informational efficiency of financial markets (e.g., Ko and Huang 2007; Peress 2014), as traders take more aggressive positions. In contrast, higher overconfidence *distorts* asset prices and can crowd out information acquisition in this model, worsening the efficiency of financial markets. Correlation neglect in particular has been mainly studied in experimental settings (Brandts, Giritligil, and Weber 2015; Eyster et al. 2018; Enke and Zimmermann 2019; Grimm and Mengel 2020; Chandrasekhar, Larreguy, and Xandri 2020), but received little attention as a mechanism in theoretical models. This paper shows that correlated noise or sentiments together with correlation neglect can provide incentives for information acquisition while maintaining an otherwise standard model structure as in Albagli, Hellwig, and Tsyvinski (2021).

The last part of this paper studies the role disagreement and funding constraints as in

Fostel and Geanakoplos (2012) and Simsek (2013, 2021). Usually, this class of models assumes two groups of traders with exogenous information who agree to disagree, such that prices are uninformative from their perspective. This approach allows studying settings where the relative wealth of optimists and pessimists is an important determinant of the asset price. This paper adds to the literature by providing a tractable model in which overconfident traders acquire information and learn from prices, yet funding constraints can be introduced. In contrast to "cash-in-the-market"-pricing as in Allen and Gale (1994), these funding constraints can also lead to asset prices being too high if optimists buy the total asset float. Moreover, I show that the effect of disagreement on asset prices depends on the beliefs of the price-setting trader. For example, if the price-setting trader is an optimist, more disagreement makes her even more optimistic, which drives up the price.

The rest of the paper is structured as follows. In section 2.2 I introduce the model and elaborate its equilibrium properties in 2.3. Then, section 2.4 showcases a number of applications. Section 2.5 studies a setting with funding constraints in more detail. Finally, section 2.6 concludes.

2.2 Model

The financial market is populated by a unit mass of risk-neutral traders indexed by $i \in [0,1]$. Time is static. Traders have deep pockets, and their utility function is

$$U_i = \mathbb{E}_i C_i - IA(\beta_i), \tag{2.1}$$

where C_i is trader i's end-of-period consumption and $IA(\beta_i)$ are convex information acquisition costs depending on signal precision β_i. Each trader can either buy up to two units of a perfectly divisible risky asset or invest in a risk-less bond with a unit return. Short-selling is ruled out.[6] Traders use their private signal and learn from the price to make their trading decision.

[6]Ruling out short-selling is not central to the analysis, and allowing limited short-selling is straightforward.

2.2.1 Assets

The risky asset has a monotonically increasing payoff function $\pi(\theta)$ where $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$ can be viewed as a sufficient statistic which pins down the asset's fundamental value (e.g., earnings or distance-to-default). For all numerical simulations $\pi(\theta) = \theta$.

2.2.2 Information Structure

Traders receive a private signal

$$s_i = \theta + \frac{\kappa\eta_i + \sqrt{1-\kappa^2}\varepsilon}{\sqrt{\beta_i}}, \tag{2.2}$$

where $\kappa \in [0, 1)$ is the share of idiosyncratic noise.[7] The signal contains two sources of noise, where $\eta_i \overset{iid}{\sim} \mathcal{N}(0, 1)$ is idiosyncratic noise and independently distributed between traders. In contrast, $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ is correlated noise that affects all traders equally. Traders can choose β_i subject to a convex cost $IA(\beta_i)$ to increase the information content of their private signal.

For traders to acquire costly information, they must believe in having an advantage over other traders, allowing them to recuperate the information acquisition costs. To this end, I assume that traders are overconfident.

Assumption 2.1 (Overconfidence)**.** *Trader i believes the information structure to be*

$$s_i = \theta + \frac{\tilde{\kappa}\eta_i + \sqrt{1-\tilde{\kappa}^2}\varepsilon}{\sqrt{\beta_i}}$$
$$s_{-i} = \theta + \frac{\kappa\eta_{-i} + \sqrt{1-\kappa^2}\varepsilon}{\sqrt{\beta_{-i}}},$$

where $\tilde{\kappa} \in (0, 1]$ and $\tilde{\kappa} > \kappa$.

Each trader believes that she is less exposed to correlated noise than other traders. Consequently, traders think they have the edge over other traders, as they expect to buy with greater probability when negatively correlated noise shocks depress the price below

[7] For $\kappa = 1$, prices are fully revealing which leads to well-known problems of inexistence as discussed in Grossman and Stiglitz (1980).

its fundamental value. In the following, expectations that suffer from *overconfidence* as in Assumption 2.1 are denoted by $\tilde{\mathbb{E}}(\cdot)$.

2.2.3 Trader's Problem

The problem of trader i is

$$\max_{\beta_i \geq 0} \tilde{\mathbb{E}} \left\{ \max_{x_i \in [0,2]} x_i(\pi(\theta) - P) | s_i, P \right\} - IA(\beta_i). \tag{P2.1}$$

Trader i takes two decisions. First, before trading takes place, trader i chooses information precision β_i subject to a convex information cost $IA(\beta_i)$ to maximize expected trading profits. Second, once prices are realized, traders use their private signal s_i and additionally learn from the price P to take the optimal buying decision subject to the position limit $x_i \in [0, 2]$.[8] Alternatively, traders can invest in a risk-less bond with unit return. Note that the expectation $\tilde{\mathbb{E}}(\cdot)$ misperceives the distribution of s_i according to Assumption 2.1.

2.2.4 Equilibrium

Trading The trader's problem is solved in reverse chronological order, starting with the trading decision. Due to linear preferences, the optimal demand of trader i taking information precision β_i as given is

$$x_i = \begin{cases} 0 & \tilde{\mathbb{E}}(\pi(\theta)|s_i, P) < P \\ \in [0,2] & \tilde{\mathbb{E}}(\pi(\theta)|s_i, P) = P \\ 2 & \tilde{\mathbb{E}}(\pi(\theta)|s_i, P) > P. \end{cases} \tag{2.3}$$

Trader i buys up to the limit whenever she expects the risky asset to yield a higher return than the risk-less bond.

Information Acquisition Plugging (2.3) into Trader i's problem (P2.1) allows to

[8]The assumption of exogenously given position limits is standard in the literature to avoid unbounded demands by risk-neutral traders. This assumption is relaxed in section 2.5

write the information acquisition problem as

$$\max_{\beta_i \geq 0} \quad \widetilde{EU}(\beta_i, \beta) - IA(\beta_i) \tag{P2.2}$$

$$\widetilde{EU}(\beta_i, \beta) = \tilde{\mathbb{E}} \left\{ \mathcal{P}(x_i = 2)(\pi(\theta) - P) \right\}, \tag{2.4}$$

where $\widetilde{EU}(\beta_i, \beta)$ denotes the expected trading profits depending on both the individual choice β_i and the symmetric choice of all other traders β. The probability of buying $\mathcal{P}(x_i = 2)$ stems from evaluating the expectations for the buying decision (2.3) over idiosyncratic noise η_i and the remaining expectations are taken over fundamental θ and correlated noise ε. Since the perceived probability of buying deviates from the true buying probability due to Assumption 2.1, also expected utility $\widetilde{EU}$ is distorted. As a result, traders expect to buy with a high probability when the price P is depressed due to a negative correlated noise shock ε. A higher β_i changes the likelihood of buying in a given state (θ, ε) as captured through the first-order condition for an interior solution $\beta_i > 0$,

$$\widetilde{MB}(\beta_i, \beta) = \tilde{\mathbb{E}} \left\{ \frac{\partial \mathcal{P}(x_i = 2)}{\partial \beta_i} (\pi(\theta) - P) \right\} = IA'(\beta_i). \tag{2.5}$$

When choosing their information precision β_i, traders weigh the benefit of a possibly more profitable buying decision with the effort cost of acquiring more precise information. The marginal benefit also depends on the information precision of all other traders, β. Intuitively, the incentive to acquire information individually crucially depends on how efficient the market is already. In the extreme case, when the market is perfectly efficient ($\lim_{\beta \to \infty} P(\theta) = \pi(\theta)$), acquiring private information becomes worthless.

As will become clear shortly, the noise in the private signal and the information learned from the price P are correlated. In this case, a less noisy private signal may make s_i less informative about θ, as traders also use s_i to learn about the correlated noise ε. In turn, the information about ε can be used to better filter information about θ from the price. Since s_i is most informative about ε when trader i decides to put only infinitesimal effort in reducing noise in her signal, the following condition needs to be checked to rule out a corner solution,

$$EU(\beta_i^*, \beta) - IA(\beta_i^*) > \lim_{\beta_i \to 0} EU(\beta_i, \beta), \tag{2.6}$$

where β_i^* is the interior solution according to (2.5). In the following, ruling out the corner solution $\beta_i \to 0$ will be important for the existence of a symmetric equilibrium.

Market-Clearing As I verify ex-post, at the symmetric equilibrium ($\forall i : \beta_i = \beta$), traders buy two units whenever their private signal is above some threshold $\hat{s}(P)$. In other words, a higher realization of the private signal leads to a higher private valuation, and due to linear preferences, all traders with a private valuation higher than the price buy. Total demand is then equal to the mass of traders with a signal $s_i > \hat{s}(P)$ times the upper position limit:

$$D\left(\theta, P\right) = 2\left(1 - \Phi\left(\frac{\sqrt{\beta}}{\kappa}\left(\hat{s}\left(P\right) - \theta\right) - \frac{\sqrt{1-\kappa^2}}{\kappa}\varepsilon\right)\right). \tag{2.7}$$

For the market to clear, total demand must be equal to the normalized asset supply of one,

$$D\left(\theta, P\right) = 1. \tag{2.8}$$

The threshold $\hat{s}\left(P\right)$ can directly be derived from (2.8).

Price Signal Since there is a one-to-one relationship between the threshold $\hat{s}\left(P\right)$ and the price P, observing the price P is informationally equivalent to observing the public signal

$$z = \hat{s}\left(P\right) = \theta + \sqrt{\frac{1-\kappa^2}{\beta}}\varepsilon. \tag{2.9}$$

In the following, I will refer to z as the price signal, and expectations will condition on z rather than P. Note that the price signal z reveals the joint information set of traders by washing out idiosyncratic noise, $z = \int_0^1 s_i di$. If traders processed their private information correctly as in (2.2), they would disregard their private signal s_i after observing z. However, due to Assumption 2.1, traders think that their private signal s_i remains informative.[9]

Orthogonalized Signal As both the price signal z and private signal s_i contain correlated noise ε for $\tilde{\kappa} < 1$, observing $\{s_i, z\}$ is informationally equivalent to observing

[9]Note that the price signal z was derived by assuming that traders have different private valuations. Since the public signal z aggregates all private information perfectly, rational traders would discard their private signal s_i and instead have homogenous valuations. Therefore, Assumption 2.1 avoids artificially breaking the indifference in the trading decision along the now uninformative private signal s_i to maintain the informative equilibrium.

$\{\bar{s}_i, z\}$, where

$$\bar{s}_i = \theta + \frac{\tilde{\kappa}}{\sqrt{\beta_i}\left(1 - \sqrt{\frac{(1-\tilde{\kappa}^2)\beta}{(1-\kappa^2)\beta_i}}\right)}\eta_i \tag{2.10}$$

is the orthogonalized signal, which has been cleaned from correlated noise ε. This confirms the initial conjecture that at the symmetric equilibrium ($\forall i : \beta_i = \beta$), a more positive realization of the private signal s_i also leads to a higher private valuation since $\bar{s}_i$ is increasing in η_i.

Uniqueness Finally, the price is equal to the private valuation of the marginal trader, who is indifferent between buying or not as in (2.3) and observed the private signal $s_i = z$. Therefore, trying to orthogonalize s_i with z as in (2.10) leaves the signal unchanged and the marginal trader conditions directly on s_i as if it was $\bar{s}_i$.[10] This result is captured in the following proposition.

Proposition 2.1. *Given $\beta > 0$, observing the price P is informationally equivalent to observing the signal $z \sim \mathcal{N}(\theta, \frac{1-\kappa^2}{\beta}\sigma_\varepsilon^2)$. In the unique equilibrium with $\forall i : \beta_i = \beta$, in which demand $x(s_i, P)$ is non-increasing in P, the price is equal to the valuation of the trader with the private signal $s_i = \hat{s}(P)$, leading to the price*

$$P(z) = \tilde{\mathbb{E}}\left\{\pi(\theta) | s_i = z, z\right\}. \tag{2.11}$$

As can be seen now, Assumption 2.1 does not only motivate trade and costly information acquisition but also distorts asset prices. Traders who suffer from stronger overconfidence perceive their private signal as being more independent of the price signal z and, therefore, put a larger weight on it when forming their valuations. This leads to a price distortion, as a rational trader would discard her private signal s_i after observing z, as s_i becomes fully uninformative.[11] As a result, the price will *overreact* to both negative and positive realizations of z.[12]

[10] More formally, the orthogonalization for (2.10) is done through $\bar{s}_i = \frac{s_i - az}{1-a}$ where $a = \sqrt{\frac{(1-\tilde{\kappa}^2)\beta}{(1-\kappa^2)\beta_i}}$. As $s_i = z$ for the marginal trader, it follows that $\bar{s}_i = s$.

[11] That s_i becomes uninformative after observing z is easy to see from the fact that at the symmetric equilibrium ($\forall i : \beta_i = \beta$) $s_i = z + \frac{\kappa}{\sqrt{\beta}}\eta_i$, i.e., s_i is a noisier version of z.

[12] See Albagli, Hellwig, and Tsyvinski (2021) for a more in-depth discussion of this distortion in a model with rational and noise traders.

2.3 Equilibrium Characterization

2.3.1 Minimal Degree of Overconfidence

An often articulated critique of models with behavioral frictions is that the severity and persistence of behavioral deviations from the rational benchmark are unrealistic. Therefore, it is desirable to keep behavioral frictions to a minimum while maintaining a similar set of results. I investigate in the following how far $\tilde{\kappa}$ needs to deviate from κ for the equilibrium to exist. I first investigate a setting where agents receive a costless signal with fixed information precision β and consider then the case with costly information acquisition. Moreover, I conduct comparative statics on β.

Exogenous Information and The Rational Limit

If traders receive a costless signal with exogenous precision β, then any $\tilde{\kappa} > \kappa$ is sufficient to generate trade. This result follows directly from (2.10), as the orthogonalized signal $\bar{s}_i$ remains informative as long as each trader believes her private signal to be at least infinitesimally more idiosyncratic than other traders' signals.

How traders process their private signal s_i naturally has an effect on the market-clearing price. If traders' overconfidence increases, they will perceive their private signal as being more independent of the price signal z and, therefore, put a larger weight on it when forming their valuations. As traders become increasingly rational ($\tilde{\kappa} \to \kappa$), this distortion vanishes, and in this sense, asset prices become more efficient:

$$\lim_{\tilde{\kappa} \to \kappa} P(z) = \tilde{\mathbb{E}} \left\{ \pi(\theta) | z \right\}. \tag{2.12}$$

Asymptotically rational traders put an infinitesimal weight on their private signal, sufficient to generate trading but insufficient to distort asset prices. Consequently, the market price behaves as if the only source of information was the public signal z as in (2.9). Note that (2.12) is identical to the price that a rational market maker would set who observed all private signals $\{s_i\}$.

Endogenous Information and Private Information as a Signal of Correlated Noise

When traders instead need to acquire costly information, substantially stronger overconfidence is needed to guarantee the existence of an equilibrium. To see this, note that the private signal s_i serves a dual function. Whenever β_i is sufficiently low ($\beta_i < \hat{\beta} = \frac{1-\tilde{\kappa}^2}{1-\kappa^2}\beta$), a positive realization of s_i is more likely to stem from high correlated noise ε than a high fundamental θ. In this case, an increase in the private signal s_i can *decrease* the private valuation of trader i when holding the price signal z fixed.[13] This negative effect can be seen from (2.10), which is decreasing in idiosyncratic noise η_i when β_i is sufficiently low. If β_i goes to zero, the signal $\bar{s}_i$ can be reformulated as

$$\lim_{\beta_i \to 0} \bar{s}_i = \theta - \frac{\tilde{\kappa}}{\sqrt{1-\tilde{\kappa}^2}} \sqrt{\frac{1-\kappa^2}{\beta}} \eta_i \tag{2.13}$$

Even as s_i becomes an infinitely noisy signal of θ, it becomes more informative of correlated noise ε. This information can be used to filter the price signal z, which also contains correlated noise ε, thus yielding a more precise estimate of θ.

Therefore, traders face not only a trade-off between precision and cost when choosing β_i, but also between learning about ε and θ. In particular, initially increasing β_i can leave traders with *less* information about θ than they would learn if they chose $\beta_i \to 0$ to maximize their information about ε. The net effect of increasing β_i is only positive once it crosses the threshold $\hat{\beta}$.

The trade-off is highlighted in Figure 2.1. In the left panel, the expected utility (2.4) is initially decreasing in β_i and has a local maximum at $\beta_i \to 0$. The right panel shows that the marginal benefit (2.5) is initially negative and turns only positive for $\beta_i > \hat{\beta}$. Therefore, any symmetric equilibrium ($\forall i : \beta_i = \beta$) must be robust to traders free-riding on the public signal z ($\beta_i \to 0$) and using s_i to learn about the correlated noise ε. In this case, the conditions for the existence of the symmetric equilibrium as in (2.6) are fulfilled.

Note that $\beta \to 0$ cannot be an equilibrium, as the public signal z becomes infinitely noisy and free-riding on z impossible. As prices become uninformative about θ, the

[13] In a first pass, a more positive realization of s_i gets translated into higher posterior beliefs about θ and ε. This new information about ε is used to filter the price signal z, which *decreases* the posterior estimate of θ through z. If s_i is sufficiently informative about ε (and uninformative about θ), this effect can dominate.

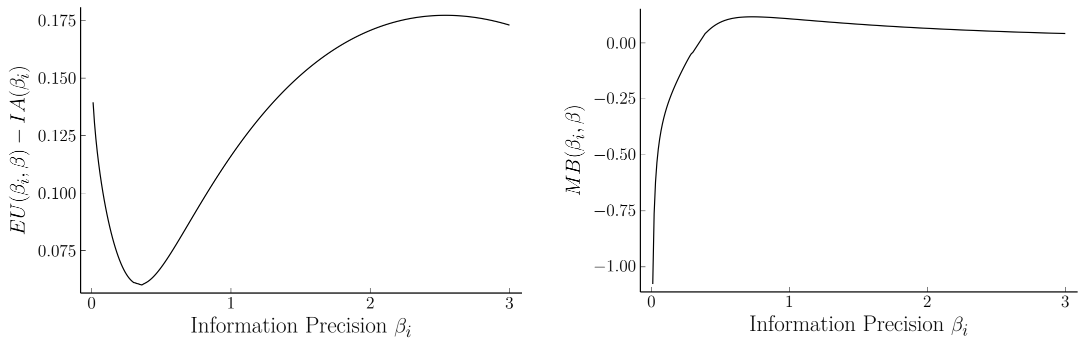

(a) Expected Utility

(b) Marginal Benefit

Figure 2.1: Expected Utility and Marginal Benefit for $\tilde{\kappa} < 1$.

Notes: Expected utility $\widetilde{EU}(\beta_i, \beta)$ minus information acquisition costs $IA(\beta_i)$ and marginal benefit $\widetilde{MB}(\beta_i, \beta)$ of increasing β_i for a given level of market information precision β.

private signal s_i can no longer be used to filter correlated noise ε out of z, and traders switch to increasing the precision of their private signal. Formally, this can be seen from (2.13) becoming infinitely noisy for $\beta \to 0$, such that any $\beta_i > 0$ would yield a more informative signal. Therefore, if the marginal costs of acquiring the initial units of information are sufficiently low, a no-information equilibrium cannot exist. However, a symmetric equilibrium with $\beta > 0$ does not have to exist either. If free-riding on the price signal z is attractive at the symmetric equilibrium, no equilibrium exists at all.

Minimal Overconfidence

As traders become increasingly overconfident ($\tilde{\kappa} \to 1$), they believe that their private signal s_i does not contain any correlated noise. Therefore, traders believe that increasing β_i always leads to more precise information on θ. As (2.13) shows, the closer $\tilde{\kappa}$ is to unity, the less valuable is the private signal for learning about correlated noise. Using continuity arguments, a symmetric equilibrium can exist whenever $\tilde{\kappa}$ is above some minimal value $\hat{\kappa} < 1$, as free-riding on the price signal z becomes less attractive as $\tilde{\kappa}$ increases.

Figure 2.2 shows how $\hat{\kappa}$, the minimal $\tilde{\kappa}$ for which the symmetric equilibrium exists, depends on κ, σ_θ^2, and σ_ε^2. The first result is that for different combinations of parameters, overconfidence needs to be substantial to incentivize information acquisition and deter free-riding on public information ($\hat{\kappa} >> \kappa$). Furthermore, for the parameters considered, $\hat{\kappa}$ decreases in σ_θ^2 and σ_ε^2. Figure 2.3 provides the corresponding equilibrium β for each combination of parameters for $\tilde{\kappa} = \hat{\kappa}$. The main takeaway is that β is increasing in uncertainty through higher σ_θ^2 and σ_ε^2. A lower κ increases the amount of correlated noise that enters the price as seen in (2.9), which makes acquiring precise information more attractive as the price deviates more strongly from its fundamental value.

Finally, the cost of information acquisition affects $\hat{\kappa}$ in two ways. For an individual trader, lower costs make it relatively more attractive to acquire more precise information in comparison to free-ride on the price signal z, lowering $\hat{\kappa}$. However, if all traders acquire more precise information, the whole market becomes more efficient, which lowers the perceived trading profits of each individual trader. Around the symmetric equilibrium ($\forall i : \beta_i = \beta$), an increase in β *decreases* the precision of $\bar{s}_i$ as in (2.10), but *increases* the precision of the information that free-riding traders can extract as in (2.13). This channel tends to increase $\hat{\kappa}$ if the marginal cost of information acquisition increases. In Figure 2.4, the second effect dominates, and $\hat{\kappa}$ is increasing in the marginal cost of information acquisition.

2.3.2 Information Acquisition and Overconfidence

Whereas overconfidence motivates information acquisition in the first place, it is less clear whether a marginal increase in overconfidence ($\tilde{\kappa}$) encourages further information acquisition or instead leaves traders satisfied with a less precise private signal. Note that $\tilde{\kappa}$ affects price informativeness only through the information choice β_i, as $\tilde{\kappa}$ has no direct influence on z and, therefore, on what an objective observer learns from the price. The following Lemma captures the main channel through which $\tilde{\kappa}$ affects the information acquisition decision.

Lemma 2.1. *Denote the orthogonalized signal as* $\bar{s}_i = a + \frac{\eta_i}{\zeta_i}$ *where*

$$\zeta_i = \frac{\sqrt{\beta_i}}{\tilde{\kappa}} \left(1 - \sqrt{\frac{(1-\tilde{\kappa}^2)\,\beta}{(1-\kappa^2)\,\beta_i}} \right).$$

Then,

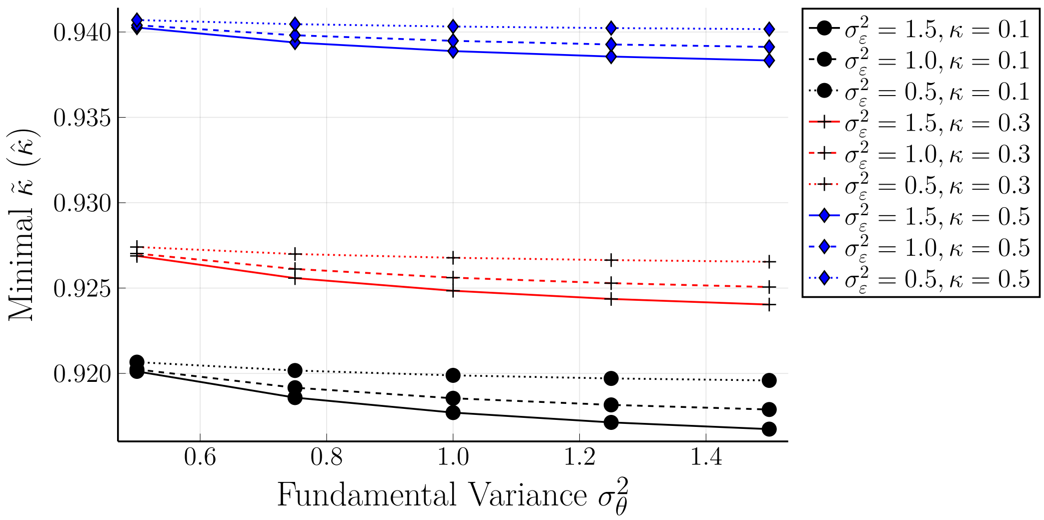

(a) Minimal $\tilde{\kappa}$ ($\hat{\kappa}$) depending on σ_θ^2

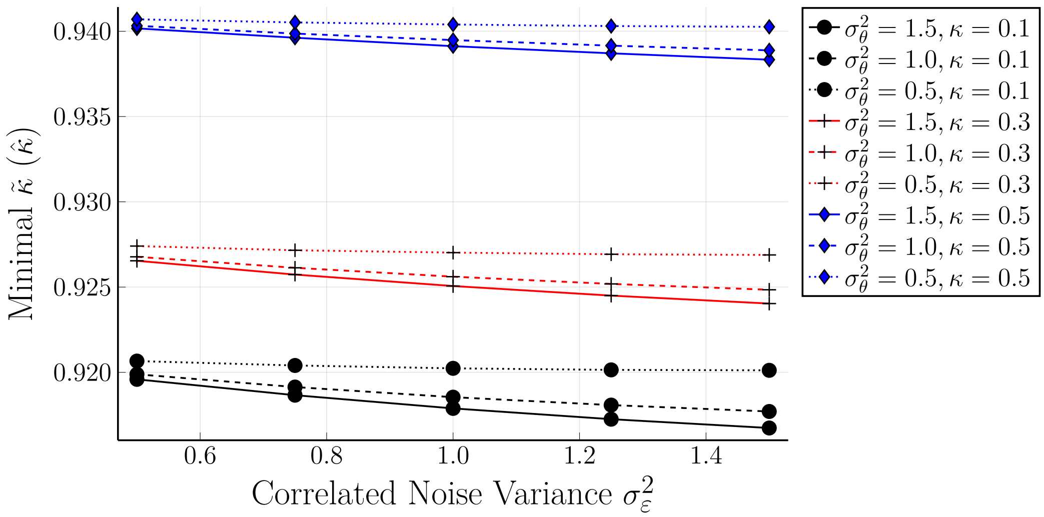

(b) Minimal $\tilde{\kappa}$ ($\hat{\kappa}$) depending on σ_ε^2

Figure 2.2: Comparative Statics for Minimal Overconfidence $\hat{\kappa}$.

Notes: Substantial amounts of overconfidence are needed when information acquisition is endogenous. The minimal $\tilde{\kappa}$ is falling in σ_θ^2 and σ_ε^2, but increasing in κ.

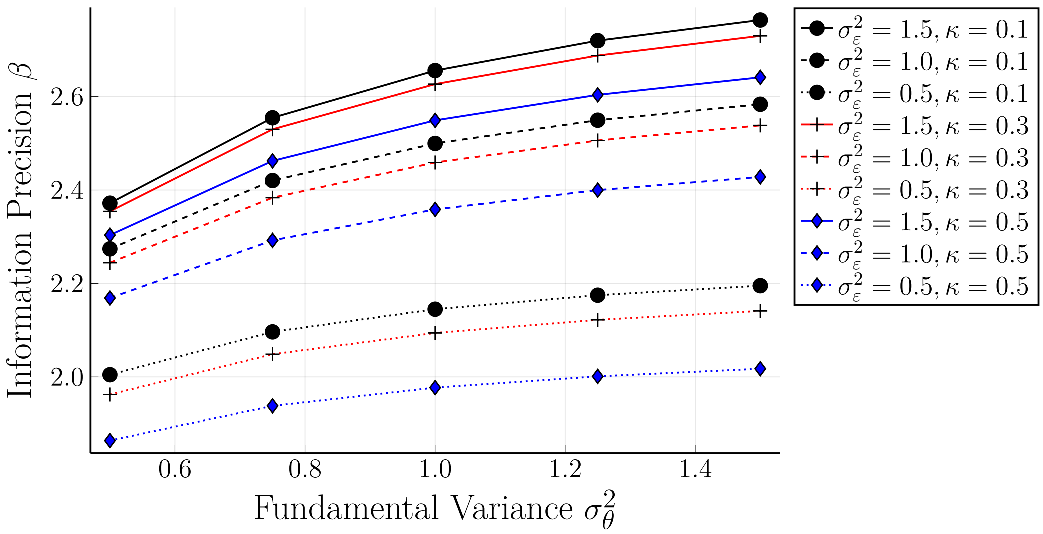

(a) Information precision β depending on σ_θ^2.

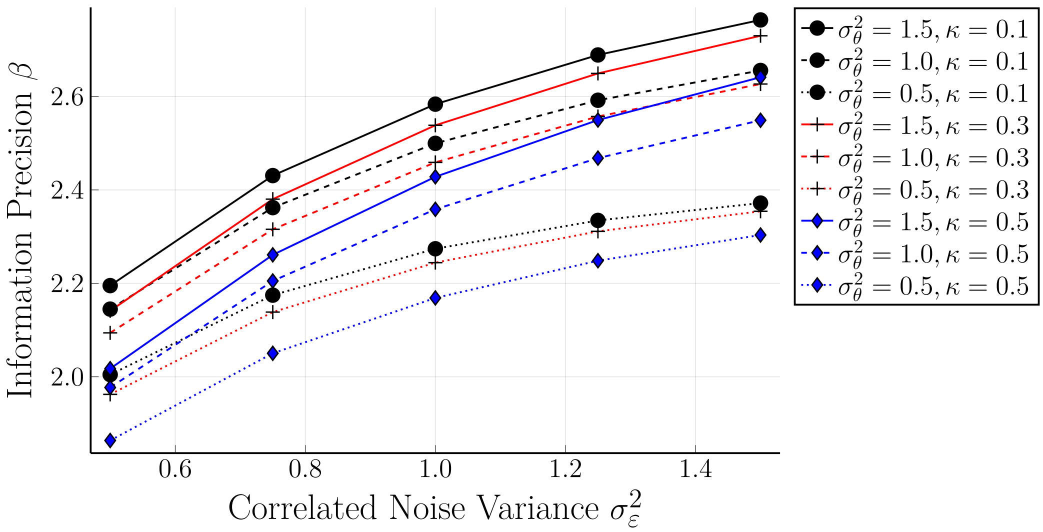

(b) Information precision β depending on σ_ε^2.

Figure 2.3: Comparative Statics for β at $\tilde{\kappa} = \hat{\kappa}$.

Notes: Equilibrium information precision β depending on σ_θ^2, σ_ε^2, and κ at the minimal $\tilde{\kappa}$ ($\hat{\kappa}$).

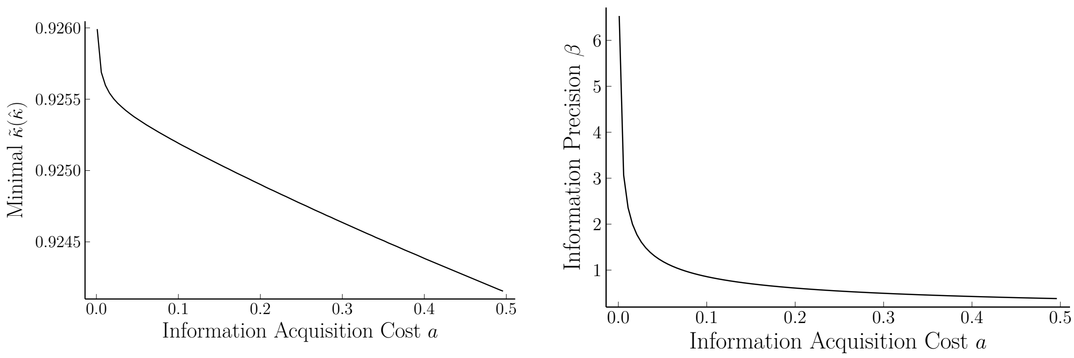

(a) Minimal $\tilde{\kappa}$ ($\hat{\kappa}$)

(b) Information Precision β

Figure 2.4: Comparative Statics on Information Acquisition Cost.

Notes: Define the information acquisition cost as $IA(\beta_i) = a\beta_i^b$. Then, the minimal $\tilde{\kappa}$ ($\hat{\kappa}$) is decreasing and information acquisition β is increasing in a. For this graph $\kappa = 0.3$, $\tilde{\kappa} = \hat{\kappa}$, $\sigma_\theta^2 = \sigma_\varepsilon^2 = 1$, $b = 2$.

(i) a higher β_i increases the precision of $\bar{s}_i$ once $\beta_i \geq \hat{\beta} = \frac{1-\tilde{\kappa}^2}{1-\kappa^2}\beta$:

$$\left.\frac{\partial \zeta_i}{\partial \beta_i}\right|_{\beta_i \geq \hat{\beta}} > 0.$$

(ii) at the symmetric equilibrium ($\forall i : \beta_i = \beta$), higher overconfidence $\tilde{\kappa}$ increases the precision of $\bar{s}_i$:

$$\left.\frac{\partial \zeta_i}{\partial \tilde{\kappa}}\right|_{\beta_i = \beta} > 0.$$

(iii) higher overconfidence $\tilde{\kappa}$ leads to smaller marginal effects of increasing β_i on the precision of $\bar{s}_i$:

$$\frac{\partial^2 \zeta_i}{\partial \beta_i \partial \tilde{\kappa}} < 0.$$

The above Lemma shows that both an increase in $\beta_i \geq \hat{\beta}$ and $\tilde{\kappa}$ at the symmetric equilibrium lead to trader i believing that her private signal s_i became more informative about θ. An increase in β_i makes trader i's private signal s_i less noisy, whereas an increase in $\tilde{\kappa}$ makes trader i believe that s_i is less correlated to the price signal z, making information extraction about θ easier. However, a larger β_i increases the precision of trader i's private information by less the more overconfident (larger $\tilde{\kappa}$) trader i is. The reason is that an increase in β_i reduces the noise in s_i and at the same time lowers s_i's correlation with the price signal z, which both increase the precision of trader i's private information. However, if the correlation is already small, then the correlation-reducing effect is limited, decreasing the overall precision-increasing effect of choosing a higher β_i. The last point leads directly to the following proposition.

Proposition 2.2. *If expected trading profits* (2.4) *are concave in the precision of trader i's private signal $\bar{s}_i$ around the symmetric equilibrium ($\beta_i = \beta$), then the individual information choice β_i is decreasing in $\tilde{\kappa}$,*

$$\left.\frac{\partial \beta_i}{\partial \tilde{\kappa}}\right|_{\beta_i=\beta} < 0.$$

The intuition for Proposition 2.2 is that overconfidence $\tilde{\kappa}$ and information acquisition β_i are substitutes, which discourages information acquisition when overconfidence $\tilde{\kappa}$ is high. Formally, an increase in overconfidence $\tilde{\kappa}$ increases the precision of $\bar{s}_i$, which lowers the marginal benefit of further increasing its precision.[14] If the marginal effect of increasing β_i on $\bar{s}_i$'s information precision was constant, this alone would lead to a decrease in trader i's choice of β_i to realign the lower marginal benefit with the unchanged marginal cost. Additionally, Lemma 2.1 (iii) shows that the marginal effect of increasing β_i is *decreasing* in overconfidence $\tilde{\kappa}$, which amplifies the crowding-out effect of higher overconfidence on information acquisition.

Whereas this mechanism operates on the individual level, Figure 2.5 shows that the symmetric information choice β can also decrease when all traders become more overconfident at the same time. Therefore, markets in which traders suffer more from severe overconfidence can be less informative. This finding contrasts with the effect of overconfidence in models with risk-averse traders (e.g., Ko and Huang 2007; Peress 2014), where overconfident traders take more aggressive trading positions, making prices more informative.

[14] The marginal benefit of increasing the precision of $\bar{s}_i$ must eventually tend to zero, as the potential trading profits are bounded from above due to exogenous position limits.

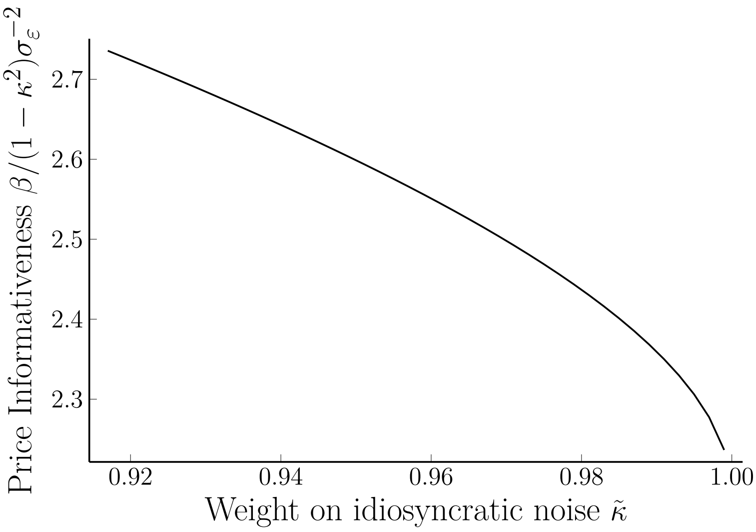

Figure 2.5: Information Precision Choice β and $\tilde{\kappa}$.

Notes: As traders become more overconfident, they find it less attractive to acquire precise information. The result is that price informativeness decreases in the market. For this figure, $\kappa = 0.3$ and $\hat{\kappa} = 0.92$.

2.3.3 Can Overconfidence Persist?

Another natural objection to models with behavioral frictions is that long-lived agents should eventually learn from their mistakes and become rational. In this case, traders would eventually learn the true composition of their private information if they repeatedly traded. Therefore, imperfect and possibly biased priors about noise composition in the private signal alone cannot sustain informed trading in this model.

To make this point, consider an infinite repetition of the static model in discrete time, $t = 1, 2, 3 \ldots \infty$. Long-lived traders can buy a risky asset every period. The properties of the asset and market are otherwise the same. Let trader i be uncertain about the composition of noise κ in her private signal (2.2). After observing $\{\theta_t, \varepsilon_t, s_{it}\}_{t=1,\ldots,T}$ for a long time, estimating

$$\sqrt{\beta_{it}}(s_{it} - \theta_t) = \alpha + \gamma\varepsilon_t + \nu_t \tag{2.14}$$

using OLS yields consistent estimates of α and γ. As traders gather more experience, the estimates converge to $\hat{\alpha} \to 0$ and $\hat{\gamma} \to \sqrt{1-\kappa^2}$. The error term $\nu_t = \kappa\eta_{it}$ is independent from ε_t by assumption.

The issue of overconfident traders learning from their mistakes can be circumvented methodologically by assuming that traders are short-lived and each period replaced by new overconfident traders. However, the literature shows that agents may not learn the truth under certain assumptions even after many periods, e.g., when agents have uncertainty about the underlying distribution of variables (Acemoglu, Chernozhukov, and Yildiz 2016).

Even if overconfident traders did not learn, a related argument is that they may not survive in the market as their mistakes are costly, which eventually deplete their wealth. Several papers suggest that less than fully rational traders can survive for a number of reasons (for an overview, see Dow and Gorton 2006). For example, arbitrageurs' unwillingness to trade aggressively against noise, e.g., due to short horizons (De Long et al. 1990), limits-to-arbitrage such as risk-aversion of arbitrageurs or imperfect information (Shleifer and Vishny 1997) or adjustments in the trading strategy by rational traders (Kyle and Wang 1997; Benos 1998). Moreover, overconfident or noise traders may use riskier strategies, which yield a higher return in the short-term (e.g., De Long et al. 1990). Finally, Hirshleifer and Luo (2001) argue that overconfident traders may be better able to exploit mispricing caused by noise or liquidity traders, allowing them survive in the long-run.

2.4 Applications

Studying the relationship between trading and information acquisition in financial markets and behavioral biases yields relevant insights in its own right, as such deviations from rationality receive increasing attention in financial settings over the last decades (for a survey, see Hirshleifer 2015). Apart from this direct interpretation, the model can be used to study more general settings and questions, in which overconfidence plays the role of a modeling device to motivate trading and information acquisition. As it turns out, the model adds tractability by simplifying the market-clearing condition, which is showcased in the example of trader heterogeneity and aggregate resource constraint. A larger extension is covered in the following section, in which traders have position limits that depend on the price and their trading capital.

2.4.1 Trader Heterogeneity

Not all traders in financial markets have access to the same information technologies. The most striking difference is between retail traders, who trade in their free time, and specialized hedge funds that may use elaborate machine learning algorithms and large quantities of data to inform their trading decisions. The model allows studying the effects of such heterogeneity in information technologies and their impact on market efficiency.

The setup is identical to before, except that two types of traders $j \in \{A, B\}$ are active in the market. The two groups are subject to group-specific sentiment shocks $\varepsilon^j \sim \mathcal{N}(0, \sigma^2_{\varepsilon^j})$ and information precisions β^j.[15] Moreover, the overconfidence assumption is simplified such that traders believe their signal to be fully idiosyncratic.[16]

Assumption 2.2 (***Overconfidence***)**.** *Trader i of group j believes the information structure to be*

$$s^j_i = \theta + \frac{\eta^j_i}{\sqrt{\beta^j_i}}$$

$$s^j_{-i} = \theta + \frac{\eta^j_{-i} + \varepsilon^j}{\sqrt{\beta^j_{-i}}}.$$

Following the same steps as before, demands by both groups clear the market,

$$D^A(\theta, \varepsilon^A, P) + D^B(\theta, \varepsilon^B, P) = 1, \tag{2.15}$$

where demand by group j can be derived by assuming that all traders in group j with a private signal $s^j_i > \bar{s}^j(P)$ buy one unit of the asset:

$$D^j(\theta, \varepsilon^j, P) = 1 - \Phi\left(\sqrt{\beta^j}\left(\hat{s}^j(P) - \theta\right) - \varepsilon^j\right). \tag{2.16}$$

Rearranging and applying the inverse of the standard normal cdf leads to the price signal as a function of the thresholds $\hat{s}^j(P)$

[15] Traders may also vary in how overconfident they are, as in Assumption 2.1. A more general case is covered in Section 2.A.2.

[16] This assumption guarantees that traders process their private signal only as being informative about θ but not about the correlated noise ε^j.

$$z^{Het} = \left(1 + \sqrt{\frac{\beta^B}{\beta^A}}\right)^{-1} \left[\hat{s}^A(P) + \sqrt{\frac{\beta^B}{\beta^A}}\hat{s}^B(P)\right] = \theta + \frac{\varepsilon^A + \varepsilon^B}{\sqrt{\beta^A} + \sqrt{\beta^B}} \tag{2.17}$$

The price signal z^{Het} has a similar form to the case with one group of traders as in (2.9). An increase in the information precision of either group reduces the correlated noise of both types in the price signal. Since the thresholds are linked through the market-clearing condition, they only signal information jointly. The uniqueness result of Proposition 2.1 also extends to this case, and the price is equal to the marginal trader's valuation of either group:

$$j \in \{A, B\} : \tilde{\mathbb{E}}^j \left\{\pi(\theta)|s_i^j = \hat{s}^j(P), z\right\} = P. \tag{2.18}$$

The relevant spillovers from one group of traders to another are the same as in the model with a single trader type. If one group acquired more precise information (e.g., β^A increased), prices become more informative, and traders in group B adjust their information acquisition in response. If information precisions across groups are substitutes, then an increase of information acquisition in one group leads to decreased information acquisition in the other group. Similarly, in a model with one trader type, lower information acquisition costs may encourage information acquisition at the individual level. However, a more precise price signal due to more market-wide information acquisition dampens this effect.

A parsimonious setup entails one rational group A that is unaffected by correlated noise ($\sigma^2_{\varepsilon^A} \to 0$) and another boundedly-rational group B with overconfident beliefs as in Assumption 2.1 or 2.2. An appealing feature of this setup is that all traders behave identically ex-ante as everyone believes to be a member of group A. Moreover, this setup shows that not all traders need to be overconfident to motivate trade and information acquisition. Any particular split in the population between rational and overconfident traders is merely chosen to maintain the normality of the price signal z^{Het}.

The model allows studying various settings with heterogeneity between traders, which is relevant when some shocks affect one group more strongly than the other. For example, the recent abundance of data and more sophisticated algorithms may benefit more institutional investors. In contrast, retail investors may not be able to process data but participate due to overconfidence. Furthermore, including a group of rational traders may be attractive as it allows to compute a measure of trading profits and think about heterogeneity and inequality as in Mihet (2020), as rational traders would make profits at the expense of overconfident traders.

2.4.2 Aggregate Resource Constraints

Conventional models of noisy financial markets study two types of agents. First, rational traders with private information, who are limited in their ability to eliminate mispricing due to limits to arbitrage (Shleifer and Vishny 1997). Second, noise traders buy or sell assets randomly and keep prices from being fully revealing, providing an incentive to acquire information for rational traders. Whereas this approach can be used productively for partial equilibrium analysis, it can lead to difficulties when considering general equilibrium settings. In particular, noise traders add and remove resources from the economy, which is an unappealing feature when the economy is otherwise closed.

This problem can be avoided by letting traders invest in many markets and allowing for an endogenously determined interest rate as in Chapter 1. With boundedly-rational traders, a more positive aggregated realization of correlated noise increases the traders' demand. The corresponding price increase is dampened by a higher interest rate, which leads to a heavier discounting of future payoffs. In this way, prices can never exceed the traders' total resources, which allows using the model as a building block in macroeconomic models. Still, aggregate noise shocks can be a source of uncertainty, as traders do not necessarily have perfect information about aggregate shocks.[17]

2.5 Funding Constraints

The baseline model assumed exogenous position limits in terms of units of the asset, an assumption that can be found in many models with risk-neutral traders (Dow, Goldstein, and Guembel 2017; Albagli, Hellwig, and Tsyvinski 2021). In reality, however, traders have finite private capital, and margin requirements limit the funds that traders can raise for investment (Brunnermeier and Pedersen 2009).

[17] A similar result is derived in Albagli, Hellwig, and Tsyvinski (2017). In their setting, households split their savings between an informed hedge fund sector and a loss-making and randomly trading mutual fund sector, whereas in Chapter 1 households manage their investments autonomously.

The following presents a model with position limits that depend on private trading capital and the asset price while maintaining the normality of the price signal. After introducing the model, I focus on the role of disagreement in forming asset prices.

2.5.1 Model

In this model, heterogeneously informed traders are limited in their investment capacity by their private trading capital.[18] Limited trading capital by itself may depress asset prices akin to "cash-in-the-market"-pricing (Allen and Gale 1994), where the asset may trade below its fundamental value due a lack of liquidity. Paired with heterogeneous beliefs as in Fostel and Geanakoplos (2012) and Simsek (2021), price-dependent limits to traders' investment capacity will play an important role in determining asset prices *at all times*. For example, if traders have plenty of trading capital, optimists will buy up the whole asset supply, thereby inflating the price. In contrast, if traders have relatively little trading capital, optimists alone will not clear the market, and increasingly pessimistic traders need to buy to absorb the total asset supply, which depresses the market-clearing price. The model is introduced more formally in the following.

2.5.2 Traders

There are overlapping generations of traders indexed by $i \in [0, 1]$. Time is discrete and infinite. Traders live for two periods, are risk-neutral and patient. Trader i's utility function, who is born in period t is

$$U_{it} = \tilde{\mathbb{E}}_{it} \left\{ C_{it+1} \right\} - IA(\beta_{it}), \tag{2.19}$$

where C_{it+1} is trader i's consumption at the end of period $t+1$ and $IA(\beta_{it})$ are information acquisition costs for a given information precision β_{it}. When young, traders each receive wealth W, which they can use to buy assets or invest in a risk-less bond with return $R > 1$. Traders cannot short-sell.[19]

[18] Abstracting away from borrowing simplifies the analysis while maintaining the main intuition. Traders' private wealth can be thought of as their maximal capital, including borrowing.

[19] Restricting short-selling simplifies the analysis, but can be introduced as long as the volume of short-selling is constrained by the amount of trading capital.

2.5.3 Assets

Each period a single perfectly dividable asset with monotonically increasing payoff function $\pi(\theta_t)$ is sold by the old to the young, where $\theta_t \overset{iid}{\sim} \mathcal{N}\left(\bar{\theta}, \sigma_\theta^2\right)$ determines the asset's fundamental value. The payoff is weakly positive for all realizations of θ_t, such that $\forall \theta_t \in \mathbb{R} : \pi(\theta_t) \geq 0$.[20] Additionally, I assume that the asset does not have a guaranteed payoff, such that $\lim_{\theta_t \to -\infty} \pi(\theta_t) = 0$.[21] Traders learn about the fundamental θ_t in period t, but the corresponding payoff $\pi(\theta_t)$ is realized at the beginning of period $t+1$ before the old sell the asset to the young.

2.5.4 Information Structure

Traders can exert effort to acquire a noisy signal of θ_t,

$$s_{it} = \theta_t + \frac{\eta_{it} + \varepsilon_t}{\sqrt{\beta_{it}}}, \tag{2.20}$$

where $\eta_{it} \overset{iid}{\sim} \mathcal{N}(0,1)$ is idiosyncratic noise, which is independently distributed among traders. In contrast, correlated noise $\varepsilon_t \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right)$ affects all traders equally. The precision of trader i's signal is captured by β_{it}. The correlated noise component ε_t can be interpreted as a form of sentiment which drives fluctuations in asset prices away from their fundamental value. To provide incentives for trading and information acquisition, traders are assumed to be overconfident.

Assumption 2.3 (*Overconfidence*). *Trader i believes the information structure to be*

$$\begin{aligned} s_{it} &= \theta_t + \frac{\eta_{it}}{\sqrt{\beta_{it}}} \\ s_{-it} &= \theta_t + \frac{\eta_{-it} + \varepsilon_t}{\sqrt{\beta_{-it}}}. \end{aligned}$$

This simplified version of the overconfidence assumption guarantees that traders perceive the noise in their private signal s_{it} as fully independent of what they learn from

[20]Ruling out negative payoffs does not substantially affect the results. Assets with negative payoffs can be studied as long as their payoff is positive in some states. In that case, some infinitely optimistic traders will always attribute a positive value to the asset, which will keep the asset price positive in all states.

[21]Ruling out safe payoffs allows to present the mechanism in the cleanest way.

the asset price ($\tilde{\kappa} = 1$ in the baseline model).[22] This overconfidence motivates traders to exert costly effort to increase their information precision, as they expect to be able to buy when asset prices are depressed due to a negative correlated noise shock.

2.5.5 Trader's Problem

Traders solve the following problem,

$$\max_{\beta_{it}} \quad \tilde{\mathbb{E}}\left\{\max_{x_{it}} \tilde{\mathbb{E}}\left\{x_{it}\left(\pi\left(\theta_t\right)+P_{t+1}-RP_t\right)|s_{it},P_t\right\}\right\}+RW-IA\left(\beta_{it}\right) \quad \text{(P2.3)}$$

$$s.t. \quad x_{it} \in \left[0, \frac{W}{P_t}\right] \quad (2.21)$$

$$\beta_{it} \geq 0. \quad (2.22)$$

The trader decides first on information precision β_{it}, which determines the precision of her private signal s_{it} and, therefore, the perceived ability to trade profitably. At the trading stage, trader i decides on how many units of the asset to buy conditional on realizations of the private signal s_{it} and asset price P_t. Different from the baseline model, the position limits in (2.21) now depend on both the price P_t and the available trading capital W. As a result, traders can buy at most W/P_t units of the asset.

2.5.6 Equilibrium

Trading The trader's problem is solved in reverse chronological order. First, taking the information precision β_{it} as given, the buying decision is

$$x\left(s_{it},P_t\right)=\begin{cases}0 & \text{if } \tilde{\mathbb{E}}\left\{\pi\left(\theta_t\right)+P_{t+1}|s_{it},P_t\right\}<RP_t \\ \in\left[0,\frac{W}{P_t}\right] & \text{if } \tilde{\mathbb{E}}\left\{\pi\left(\theta_t\right)+P_{t+1}|s_{it},P_t\right\}=RP_t \\ \frac{W}{P_t} & \text{if } \tilde{\mathbb{E}}\left\{\pi\left(\theta_t\right)+P_{t+1}|s_{it},P_t\right\}>RP_t\end{cases}. \quad (2.23)$$

Trader i buys zero units if her valuation is below the price, is indifferent between buying or not when her valuation equals the price, and spends her whole wealth W if her valuation exceeds the price.

[22] This formulation is also used in Chapter 1.

Information Acquisition Trader i chooses her information precision β_{it} to improve her ability to identify profitable trading opportunities. The first-order condition for the information production decision is obtained after plugging (2.23) into (P2.3). Evaluating expectations with respect to the realizations of idiosyncratic noise η_{it} leads to the probability of buying $\mathcal{P}\{x_{it} = \frac{W}{P_t}\}$. Taking the information acquisition decision of other traders as given ($\beta_{-it} = \beta_t$) and taking the partial derivative with respect to β_{it} leads to the marginal benefit of information acquisition,

$$\widetilde{MB}\left(\beta_{it}, \beta_t\right) = \tilde{\mathbb{E}}\left\{\frac{\partial \mathcal{P}\left\{x_{it} = \frac{W}{P_t}\right\}}{\partial \beta_{it}} \frac{W}{P_t}\left(\pi\left(\theta_t\right) + P_{t+1} - RP_t\right)\right\}. \tag{2.24}$$

The marginal benefit of acquiring information consists of three parts. The first is the change in the probability of buying in state $(\theta_t, \varepsilon_t)$ given information choices (β_{it}, β_t). The position size W/P_t determines the stake of trader i and scales any trading profits. Finally, trading profits given by the difference between the payoff plus the resale price $\pi\left(\theta_t\right) + P_{t+1}$ and the opportunity cost of buying RP_t.

Market-Clearing In the symmetric equilibrium ($\forall i : \beta_{it} = \beta_t$), traders spend their total trading capital W whenever their private signal is above some threshold, $\hat{s}\left(P_t, W\right)$. Normalizing the asset supply to one and summing up demands of all traders with a private signal about the threshold leads to the market-clearing condition,

$$\frac{W}{P_t}\left(1 - \Phi\left(\sqrt{\beta_t}\left(\hat{s}\left(P_t, W\right) - \theta_t\right) - \varepsilon_t\right)\right) = 1, \tag{2.25}$$

which allows solving for the threshold directly,

$$\hat{s}\left(P_t, W\right) = \theta_t + \frac{\varepsilon_t + \Phi^{-1}\left(1 - \frac{P_t}{W}\right)}{\sqrt{\beta_t}}. \tag{2.26}$$

The identity of the marginal trader now also depends on the ratio between price and wealth. If the price P_t is large relative to total wealth W, most traders need to buy to clear the market. Therefore, the marginal trader needs to be close to the bottom of the trader distribution, i.e., she must be pessimistic about the asset.

Price Signal Traders learn from the price P_t, which is equivalent to observing the noisy signal,

$$z_t = \hat{s}\left(P_t, W\right) - \frac{\Phi^{-1}\left(1 - \frac{P_t}{W}\right)}{\sqrt{\beta_t}} = \theta_t + \frac{\varepsilon_t}{\sqrt{\beta_t}}. \tag{2.27}$$

I call z_t the *price signal* and expectations condition on z_t instead of P_t. Note that although prices depend on trading capital W, the price signal is invariant to changes in W and remains always normally distributed.[23] The normality of the price signal is maintained due to the focus of a single trader type, which keeps the following analysis tractable.

Uniqueness The following proposition shows that the equilibrium is unique for a given symmetric information precision β_t. Moreover, the price P_t is equal to the valuation of the *marginal trader* who is just indifferent between buying or not buying and who observed the private signal $s_{it} = \hat{s}\left(P_t, W\right)$. Any trader who is more optimistic than the marginal trader ($s_{it} > \hat{s}\left(P_t, W\right)$) buys W/P_t shares, whereas more pessimistic traders invest their trading capital in the risk-less bond.

Proposition 2.3. *Given $\beta_t > 0$, observing P_t is equivalent to observing the signal $z_t \sim \mathcal{N}(\theta_t, \sigma_\varepsilon^2/\beta_t)$. In the unique equilibrium, in which demand $x\left(s_{it}, P_t, W\right)$ is non-increasing in P_t, the price is equal to the valuation of the trader with the private signal $s_{it} = \hat{s}\left(P_t, W\right)$, leading to the price*

$$P\left(z_t, W\right) = \frac{1}{R}\tilde{\mathbb{E}}\left\{\pi\left(\theta_t\right) + P_{t+1} | s_{it} = \hat{s}\left(P_t, W\right), z_t\right\}. \tag{2.28}$$

The price reflects the beliefs of the marginal trader $\hat{s}\left(P_t, W\right)$ and the price signal z_t. If due to positive news the price signal z_t increases, then all traders become more optimistic, and the price must increase in response. However, an increase in the price requires a change in the identity of the marginal trader. Due to the price increase, the joint trading capital of the previous buyers is insufficient, and more pessimistic traders need to buy to clear the market. Consequently, the effect of an increase in z_t on the price P_t is *dampened*, because the marginal trader becomes *more pessimistic* (i.e., $\hat{s}(P_t, W)$ as in (2.26) is decreasing in P_t).

Resale Price For simplicity, I assume that θ_t and ε_t are iid, which leads to the expected resale price

$$\mathbb{E}\left\{P_{t+1}\right\} = \frac{1}{R-1}\mathbb{E}\left\{\tilde{\mathbb{E}}\left\{\pi\left(\theta_t\right) | s_{it} = \hat{s}\left(P_t, W\right), z_t\right\}\right\}. \tag{2.29}$$

[23] I demonstrate in the Appendix 2.A.3 that normality is lost if the model is populated by rational and noise traders.

For a given β_t, the expected resale price (2.29) is uniquely determined. Whereas the LHS is monotonically increasing in $\mathbb{E}\left\{P_{t+1}\right\}$, the right hand side is monotonically decreasing. An increase in the resale price increases the price today, which leads to a downward shift in the identity of the marginal trader. As a result, the now more pessimistic marginal trader values the payoff $\pi(\theta_t)$ less.

2.5.7 Equilibrium Characterization

This section investigates the relationship between disagreement and the average price level and shows that a high expected resale price discourages information acquisition.

Information Precision and Disagreement

Changes to the information precision β_t have two effects on the market-clearing price. First, more precise information lets traders put more weight on the price signal z_t and their private information. Therefore, a higher β_t makes the price react more strongly to changes in the price signal z_t. Second, the precision of private information determines the degree of disagreement between traders. If traders disagree more, any change in the identity of the marginal trader to maintain market-clearing will have a stronger effect on the price. In particular, disagreement among traders must be hump-shaped in information precision β_t as seen in Figure 2.6a. As traders share a common prior, they perfectly agree if they do not acquire private information ($\beta_t = 0$). Similarly, if traders acquired perfect information, they would all learn the truth. Therefore, disagreement must take its maximum for an intermediate value of information precision β_t.

These two effects of changes in information precision β_t do not wash out when averaging over different realizations of the price signal z_t as seen in Figures 2.6b and 2.6c. In these figures, the expected price $\mathbb{E}\left\{P_{t+1}\right\}$ is shown as a function of a constant information precision in all coming periods ($\forall s \geq t : \beta_s = \beta$). Although the price does not feature informational feedback (e.g., through a managerial decision based on z_t), the expected *price level* moves with the constant information precision β.[24]

[24]The expected price would also depend on information precision β, if position limits were given in units of the asset, e.g., as in Albagli, Hellwig, and Tsyvinski (2021). The reason is a failure of the law of iterated expectations due to imperfect information aggregation, and its severity depends on the information precision β. This channel is also present here, but not of main interest.

Whether the expected price is initially decreasing or increasing in β_t depends on how abundant or scarce trading capital is among traders. If traders have abundant resources to buy the asset, optimists ($\eta_{it} > 0$) will be able to buy up the market, which inflates the price. As traders acquire initial units of information, optimists become relatively more optimistic as traders disagree more intensely. Therefore, the price may increase initially in β_t as in 2.6b. In contrast, if trading capital is scarce, most traders need to buy to clear the market. Mechanically, the marginal trader must be relatively pessimistic ($\eta_{it} < 0$), and an initial increase in β_t depresses the price as seen in 2.6c. Conversely, a decrease in β_t may be associated with an increase in the expected price, as an uninformed trader would value the asset more highly than an informed yet pessimistic trader.

Information Acquisition and Resale Price

A high expected resale price can discourage information acquisition as it limits the trader's ability to exploit her information. Today's information helps forecast the payoff in the proximate periods but loses precision for the distant future if θ_t is not fully persistent. For example, an interest rate fall leads to a price increase primarily driven by less discounting on distant payoffs. As a result, traders must buy fewer units at a higher price, limiting their ability to speculate on proximate payoffs for which their information is most valuable.

Moreover, a higher resale price limits potential trading losses. From (2.23), trader i buys whenever she expects to turn a profit, i.e., $\tilde{\mathbb{E}}\{\pi|s_{it}, P_t\} > RP - \mathbb{E}\{P_{t+1}\}$. However, an increase in the expected resale price $\mathbb{E}\{P_{t+1}\}$ translates to a less than one-for-one increase in P_t as the marginal trader's identity needs to adjust to maintain market-clearing. As a result, the next payoff stream $\pi(\theta_t)$ is valued by a more pessimistic trader and the RHS is decreasing in $\mathbb{E}\{P_{t+1}\}$ and trader i expects to turn a profit also for lower realizations of her private signal s_{it}. In the extreme case, as $\frac{1}{R}\mathbb{E}\{P_{t+1}\} \to W$, trader i finds it optimal to buy irrespective of her private information, leading to an information-insensitive buying decision and, therefore, no incentive for information acquisition.

Taken together, an increase in the expected resale price can discourage information acquisition, as seen in Figure 2.7. In the figure, the expected resale price is changed exogenously, and traders decide on their information precision β_{it} before trading takes place in period t.

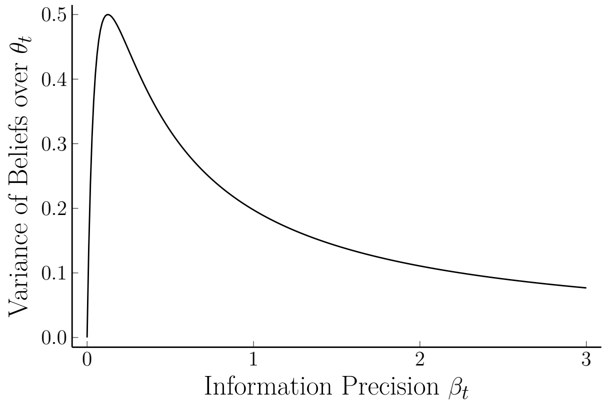

(a) Disagreement is hump-shaped in β_t.

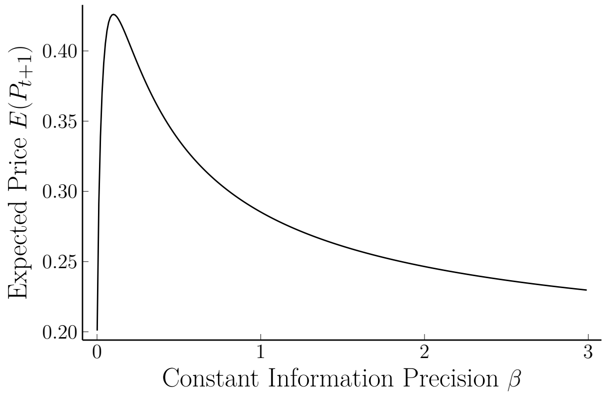

(b) Marginal traders are mostly optimists and trading capital is abundant.

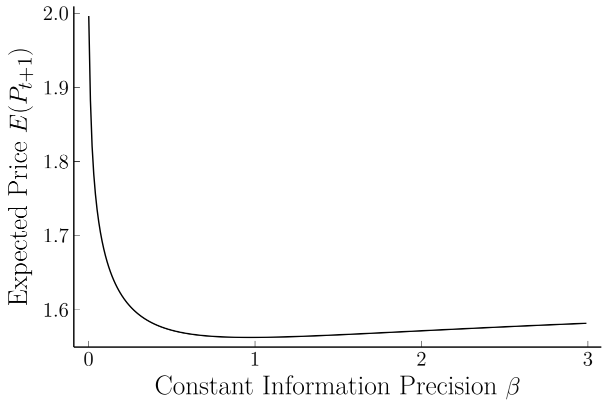

(c) Marginal traders are mostly pessimists and trading capital is scarce.

Figure 2.6: Disagreement and Expected Price Depending on Information Precision.

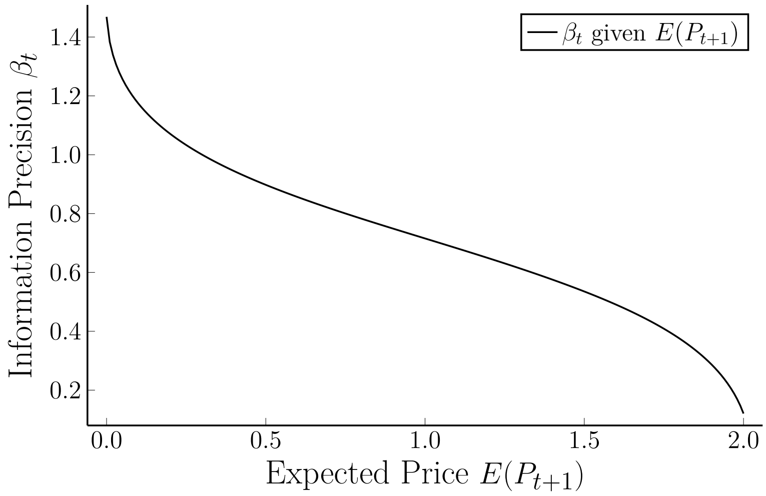

Figure 2.7: Information Precision Choice β_t is Decreasing in $\mathbb{E}\{P_{t+1}\}$.

Discussion

The presented model is stylized and seeks to capture the main mechanism linking disagreement through varying degrees of information acquisition and funding constraints, limiting the positions that traders can take. In the following, I discuss other topics that are not formally incorporated in the model yet relevant.

Interest Rates In this model, the interest rate R is exogenously given, although it is natural to assume that it is connected to the rate of return on investment. For example, when trading capital is scarce, expected returns should be high, and interest rates should increase. As a result, asset prices fall, and trading capital becomes relatively less scarce. Therefore, endogenously determined interest rates are an offsetting force to the relative scarcity or abundance of trading capital.

Funding Constraints and General Equilibrium Throughout, I have assumed that traders' trading capital W is constant. Suppose the described financial market is to be understood as a market for a single specific asset. In that case, liquidity cannot remain scarce indefinitely. At some point, traders will redistribute their capital towards illiquid markets to earn a premium. Instead, the asset is more abstract and representative of the whole stock market, the economy should accumulate capital over time as returns are high. In both cases, markets should adjust in the long-run and provide capital where its return is the highest.

Borrowing Constraints Borrowing between traders is a natural mechanism to redistribute trading capital from pessimistic to optimistic traders as in Simsek (2013). Such borrowing can avoid depressed prices due to a lack of trading capital in the hands of optimists, which could move asset prices closer to their objective valuation. The main complication in considering borrowing between traders is that the price signal z is normally distributed when trading capital is equally distributed among buyers, which is otherwise not guaranteed. One possibility is to consider an ex-ante security design problem, in which buyers issue and sell a junior tranche or credit default swaps to more pessimistic traders. As a result, buyers borrow from other traders and the same amount of trading capital.

Complexity Finally, the model sheds light on the incentives of asset originators in creating complex assets (Asriyan, Foarta, and Vanasco 2020) when investors are heterogeneously informed. Complex assets are more difficult to learn about, i.e., information acquisition costs are higher for such assets. If trading capital is sufficiently scarce (see Figure 2.6c), the model predicts that the asset originator may want to make information acquisition more costly, i.e., make the asset *more complex*. Such an increase in complexity can increase asset prices by reducing disagreement. On the other hand, if funding is abundant (see Figure 2.6b), an asset originator may prefer an intermediate level of information acquisition, which maximizes the price-inflating influence of optimistic traders. To achieve this goal, the asset originator may increase or decrease complexity, depending on the initial level of information acquisition.

2.6 Conclusion

I presented a model of financial markets with dispersed information similar to Albagli, Hellwig, and Tsyvinski (2021), in which overconfidence motivates trade and information acquisition. Traders overestimate the precision of their private signal, as they underestimate the correlation of their signal with the information that can be learned from the market-clearing price. I parametrize the degree of overconfidence, which governs the discrepancy between the true and perceived distribution of the private signal.

Whereas an infinitesimal amount of overconfidence is sufficient to generate trade when information is free, overconfidence must be substantial for equilibrium existence when information is costly. The reason is that the traders' private signals serve a dual function. The private signal is informative about the asset's fundamental and correlated noise, which affects the price. Information about the correlated noise can be used to better distinguish between price changes driven by fundamentals or noise.

As acquiring information reduces the noise in the signal, traders need to balance the cost and benefit of gathering more precise information and decide on learning about correlated noise or the fundamental. By not acquiring information, traders can choose to *free-ride* on the information acquisition of other traders. Since *free-riding* cannot be an equilibrium strategy, equilibrium existence requires that traders find information acquisition and learning about the fundamental sufficiently attractive. This is the case when traders believe their signal to be relatively uninformative about correlated noise, i.e., which is true for strongly overconfident traders.

I use the model to study several applications, for example, trader heterogeneity. Traders that suffer from varying degrees of overconfidence or have different levels of exposure to correlated noise can interact through the financial market while maintaining the normality of the price signal. This setup can be used to study the effects of technological changes (e.g., availability of new data sources and processing techniques) that disproportionally affect one group of traders (e.g., institutional investors in comparison to retail traders). Moreover, the model can be used as a building block in general equilibrium models, as shown in Chapter 1.

Finally, I study a setting in which traders' position limits depend on their available trading capital and the asset price. With such funding constraints, the effect of shocks on the asset price is dampened. For instance, positive news about the fundamental increase the asset price. As a result, more traders need to buy to clear the market. As the most optimistic traders buy first, these new buyers must be relatively more pessimistic than the previous buyers. As a result, the initial price increase is reduced. The model makes predictions on the relationship between disagreement and the asset price, as disagreement determines how extreme the beliefs of optimists or pessimists are. For instance, disagreement increases the asset price if price-setting traders are mostly optimists.

Appendix

2.A Derivations

2.A.1 Market-Clearing with One Trader Type

All traders with $s_i > \hat{s}(P)$ buy two units of the risky asset, such that

$$s_i > \hat{s}\left(P\right) \iff \eta_i > \frac{\sqrt{\beta}}{\kappa}\left(\hat{s}\left(P\right) - \theta\right) - \frac{\sqrt{1-\kappa^2}}{\kappa}\varepsilon \tag{2.30}$$

Given that $\eta_i \sim \mathcal{N}(0,1)$, the probability of buying can be written as

$$P(s_i > \hat{s}(P)) = 1 - \Phi\left(\frac{\sqrt{\beta}}{\kappa}\left(\hat{s}\left(P\right) - \theta\right) - \frac{\sqrt{1-\kappa^2}}{\kappa}\varepsilon\right) \tag{2.31}$$

Equating total demand to a normalized asset supply of one leads to the market-clearing condition, which allows to solve for $\hat{s}(P)$ directly

$$\begin{aligned} 2\left(1 - \Phi\left(\frac{\sqrt{\beta}}{\kappa}\left(\hat{s}\left(P\right) - \theta\right) - \frac{\sqrt{1-\kappa^2}}{\kappa}\varepsilon\right)\right) &= 1 \\ \iff \Phi\left(\frac{\sqrt{\beta}}{\kappa}\left(\hat{s}\left(P\right) - \theta\right) - \frac{\sqrt{1-\kappa^2}}{\kappa}\varepsilon\right) &= \frac{1}{2} \\ \iff \frac{\sqrt{\beta}}{\kappa}\left(\hat{s}\left(P\right) - \theta\right) - \frac{\sqrt{1-\kappa^2}}{\kappa}\varepsilon &= 0 \\ \iff \hat{s}\left(P\right) = \theta + \sqrt{\frac{1-\kappa^2}{\beta}}\varepsilon, \end{aligned} \tag{2.32}$$

which is also equal to the market signal z.

2.A.2 Market-Clearing with Two Trader Types

Assume now instead that there are two groups of traders indexed by $j \in \{A, B\}$. The signal structure is as before but with group-specific weight on idiosyncratic noise κ^j, information precision β^j, and correlated noise shock ε^j. Furthermore, assume that the parameters for both groups are such that either both groups get more optimistic or pessimistic as their private signal increases. The former is the case when the private signal is sufficiently more informative about the fundamental θ than correlated noise ε. The latter is true for the opposite case.

The market-clearing condition is

$$\sum_{j \in \{A,B\}} D^j(\theta, \varepsilon^j, P) = 1 \tag{2.33}$$

In the case when the traders' private signals are more informative about θ, all traders above the threshold $\hat{s}^j(P)$ buy,

$$D^j(\theta, \varepsilon^j, P) = 1 - \Phi\left(\frac{\sqrt{\beta^j}}{\kappa^j}\left(\hat{s}^j(P) - \theta\right) - \frac{\sqrt{1-(\kappa^j)^2}}{\kappa^j}\varepsilon^j\right). \tag{2.34}$$

In the other case, when the private signal is more informative about ε than θ, all traders with a signal below $\hat{s}^j(P)$ buy,

$$D^j(\theta, \varepsilon^j, P) = \Phi\left(\frac{\sqrt{\beta^j}}{\kappa^j}\left(\hat{s}^j(P) - \theta\right) - \frac{\sqrt{1-(\kappa^j)^2}}{\kappa^j}\varepsilon^j\right). \tag{2.35}$$

In both cases, rearranging (2.33) and applying the inverse of the standard normal cdf leads after some algebra to the price signal

$$\begin{aligned} z &= \left(1 + \sqrt{\frac{\beta^B}{\beta^A}}\frac{\kappa^A}{\kappa^B}\right)^{-1}\left[\hat{s}^A(P) + \sqrt{\frac{\beta^B}{\beta^A}}\frac{\kappa^A}{\kappa^B}\hat{s}^B(P)\right] \\ &= \theta + \frac{\sqrt{1-(\kappa^A)^2}\varepsilon^A + \frac{\kappa^A}{\kappa^B}\sqrt{1-(\kappa^B)^2}\varepsilon^B}{\sqrt{\beta^A} + \sqrt{\beta^B}\frac{\kappa^A}{\kappa^B}}, \end{aligned} \tag{2.36}$$

Finally, since marginal traders in both groups must be indifferent between buying or not, the thresholds can be derived from

$$\mathbb{E}^j(\pi(\theta)|s_i^j = \hat{s}^j(P), P) = P. \tag{2.37}$$

2.A.3 Limited Trading Capital and Noise Trading

In contrast to section 2.5, assume that traders receive a signal of the form

$$s_{it} = \theta_t + \frac{\eta_{it}}{\sqrt{\beta_{it}}}, \tag{2.38}$$

where $\eta_{it} \sim \mathcal{N}(0,1)$ is idiosyncratic noise. Additionally, noise traders demand a random number $\frac{W}{P_t}\Phi(u_t)$ of assets where $u_t \sim \mathcal{N}(0, \sigma_u^2)$. Following the same steps as before, the market-clearing condition is

$$\frac{W}{P_t}\left(1 - \Phi\left(\sqrt{\beta_t}\left(\hat{s}\left(P_t, W\right) - \theta_t\right)\right)\right) + \frac{W}{P_t}\Phi(u_t) = 1, \tag{2.39}$$

where the threshold can be derived as

$$\hat{s}\left(P_t, W\right) = \theta + \frac{\Phi^{-1}\left(1 - \frac{P_t}{W} + \Phi(u_t)\right)}{\sqrt{\beta}}. \tag{2.40}$$

From here it is evident that the threshold $\hat{s}\left(P_t, W\right)$ is not a linear function of normally distributed random variables and therefore itself not a normally distributed signal of θ.

2.A.4 Proofs

Proof of Proposition 2.1. The proof is identical to the proof of Proposition 1.1 in Chapter 1 or the proof to Proposition 1 in Albagli, Hellwig, and Tsyvinski (2021). The only difference is the information structure. It is sufficient to show that at the symmetric equilibrium ($\forall i : \beta_i = \beta$), a more positive realization of s_i also leads to a higher private valuation. Since the public signal z as in (2.12) and the private signal s_i as in (2.2) both contain correlated noise ε, s_i can be orthogonalized to derive a signal that is independent of ε,

$$\bar{s}_i = \frac{s_i - \sqrt{\frac{(1-\tilde{\kappa^2})\beta}{(1-\kappa^2)\beta_i}}z}{1 - \sqrt{\frac{(1-\tilde{\kappa^2})\beta}{(1-\kappa^2)\beta_i}}} = \theta + \frac{\tilde{\kappa}}{\sqrt{\beta_i}\left(1 - \sqrt{\frac{(1-\tilde{\kappa}^2)\beta}{(1-\kappa^2)\beta_i}}\right)}\eta_i. \tag{2.41}$$

Therefore, observing $\{s_i, z\}$ is equivalent to observing $\{\bar{s}_i, z\}$. It follows that traders with a more positive realization of $\bar{s}_i$ indeed have a higher valuation, as their posterior on θ is increasing in $\bar{s}_i$. □

Proof of Lemma 2.1. Write the orthogonalized signal as

$$\bar{s}_i = a + \frac{\eta_i}{\zeta_i} \tag{2.42}$$

where

$$\zeta_i = \frac{\sqrt{\beta_i}}{\tilde{\kappa}}\left(1 - \sqrt{\frac{1-\tilde{\kappa}^2}{1-\kappa^2}}\sqrt{\frac{\beta}{\beta_i}}\right). \tag{2.43}$$

The results are then straightforward to show. (i) follows from

$$\frac{\partial \zeta_i}{\partial \beta_i} = \frac{1}{2\beta_i \tilde{\kappa}} > 0. \tag{2.44}$$

(ii) can be derived as

$$\begin{aligned}
\left.\frac{\partial \zeta_i}{\partial \tilde{\kappa}}\right|_{\beta_i=\beta} &= \sqrt{\beta}\left(\frac{\partial}{\partial \tilde{\kappa}}\frac{1}{\tilde{\kappa}} - \frac{\partial}{\partial \tilde{\kappa}}\sqrt{\frac{\tilde{\kappa}^{-2}-1}{1-\kappa^2}}\right) \\
&\propto -\frac{1}{\tilde{\kappa}^2} - \frac{\left(\tilde{\kappa}^{-2}-1\right)^{-\frac{1}{2}}\left(-2\tilde{\kappa}^{-3}\right)}{2\sqrt{1-\kappa^2}} \\
&= \frac{1}{\tilde{\kappa}^3\sqrt{\tilde{\kappa}^{-2}-1}\sqrt{1-\kappa^2}} - \frac{1}{\tilde{\kappa}^2} \\
&= \frac{1}{\tilde{\kappa}^2}\left(\frac{1}{\sqrt{1-\tilde{\kappa}^2}\sqrt{1-\kappa^2}} - 1\right) \\
&> 0.
\end{aligned} \tag{2.45}$$

Finally, (iii) stems simply from

$$\frac{\partial^2 \zeta_i}{\partial \beta_i \partial \tilde{\kappa}} = -\frac{1}{2\beta_i \tilde{\kappa}^2} < 0 \tag{2.46}$$

□

Proof of Proposition 2.2. Starting from a symmetric equilibrium ($\beta_i = \beta$), an increase in β_i cannot make trader i worse off. Given that the orthogonalized signal (2.10) becomes more precise as $\beta_i \geq \beta$ increases, traders can always add noise to their signal to maintain the same signal precision. However, it must be that $\beta_i \to \infty$ provides the highest level of utility, as trader i can then realize all trading profits, which is unattainable with a noisy signal. Therefore, the marginal benefit of increasing $\beta_i \geq \beta$ must be weakly positive everywhere and strictly positive somewhere.

Assume that expected trading profits as in (2.4) are concave in ζ_i around the symmetric equilibrium $\beta_i = \beta$. Moreover, a higher β_i increases expected trading profits around the symmetric equilibrium as $\zeta_i > 0$ and $\frac{\partial \zeta_i}{\partial \beta_i} > 0$ as in Lemma 2.1 (i). Then, an increase in $\tilde{\kappa}$ also increases ζ_i as in Lemma 2.1 (ii), which decreases the marginal benefit of increasing ζ_i further holding β_i constant due to the concavity of expected trading profits. Reinforcing this effect, the marginal effect of increasing β_i on ζ_i is also lower after an increase in $\tilde{\kappa}$ as follows from Lemma 2.1 (iii).

Taking these results together, if $\tilde{\kappa}$ increased for trader i, then her marginal benefit of increasing β_i must decrease around $\beta_i = \beta$. As a result, trader i chooses a lower information precision β_i. □

Proof of Proposition 2.3. The proof is identical to the proof of Proposition 1.1 in Chapter 1 or the proof to Proposition 1 in Albagli, Hellwig, and Tsyvinski (2021). The only difference is that traders' positions now depend on the price P_t and trading capital W. Therefore, it has to be verified that the price P_t is increasing in the price signal z_t. The proof begins in the following.

There must be a threshold $\hat{s}(P_t, W)$ such that all traders with $s_{it} \geq \hat{s}(P_t, W)$ find it profitable to buy $\frac{W}{P_t}$ units of the risky asset and otherwise abstain from buying. Different from before, the threshold also depends on trading capital W, as traders may be able to buy different quantities of the asset depending on how much capital they have and how expensive the asset is.

The price is equal to the valuation of the marginal trader as in (2.23) who is just indifferent between buying or not

$$P_t = \frac{1}{R}\tilde{\mathbb{E}}\left\{\pi(\theta_t) + P_{t+1} | s_{it} = \hat{s}(P_t, W), P_t\right\}. \tag{2.47}$$

The price is now the solution to the implicit function above, as the price determines the threshold $\hat{s}(P_t, W)$ on the right-hand-side. This monotone demand schedule leads to total demand

$$D(\theta_t, \varepsilon_t, P_t, W) = \frac{W}{P_t}\left(1 - \Phi\left(\sqrt{\beta_t}\left(\hat{s}\left(P_t, W\right) - \theta_t\right) - \varepsilon_t\right)\right), \tag{2.48}$$

where $\Phi(\cdot)$ is the standard-normal cdf. Equalizing total demand with a normalized supply of one leads to the market-clearing condition

$$\frac{W}{P_t}\left(1 - \Phi\left(\sqrt{\beta_t}\left(\hat{s}\left(P_t, W\right) - \theta_t\right) - \varepsilon_t\right)\right) = 1, \tag{2.49}$$

with the unique solution

$$\hat{s}\left(P_t, W\right) = \theta_t + \frac{\varepsilon_t + \Phi^{-1}\left(1 - \frac{P_t}{W}\right)}{\sqrt{\beta_t}}. \tag{2.50}$$

As is evident now, the threshold cannot be expressed purely in terms of the fundamental shock θ_t, correlated noise ε_t, and information precision β_t. Therefore, it is also not possible to express the price P_t explicitly.

Nonetheless, the price is uniquely pinned down by the threshold $\hat{s}\left(P_t, W\right)$ and P_t is equal to the trading capital of all traders with a signal above the threshold. The price signal z_t can be extracted from the threshold,

$$z_t = \hat{s}\left(P_t, W\right) - \frac{\Phi^{-1}\left(1 - \frac{P_t}{W}\right)}{\sqrt{\beta_t}} = \theta_t + \frac{\varepsilon_t}{\sqrt{\beta_t}}. \tag{2.51}$$

The price is also invertible with respect to z_t. Consider an increase in z_t, which also increases the valuation of the marginal trader and the price P_t. However, the previous buyers are not able to clear the market anymore. Therefore, $\hat{s}\left(P_t, W\right)$ needs to shift down, accommodating the higher price P_t. It follows that there is a bijective mapping between the threshold $\hat{s}\left(P_t, W\right)$ and the price signal z_t as shown above. Therefore, the price P_t is invertible in z_t.

It follows that observing P_t is equivalent to observing $z_t \sim \mathcal{N}\left(\theta_t, \beta_t^{-1}\sigma_\varepsilon^2\right)$. Traders treat the signal z_t and their private signal $s_{it} \sim \mathcal{N}\left(\theta_t, \beta_{it}^{-1}\right)$ as mutually independent. Conditioning on the price signal z_t allows to rewrite the price as

$$P_t = \frac{1}{R}\tilde{\mathbb{E}}\left\{\pi(\theta_t) + P_{t+1} | s_{it} = \hat{s}\left(P_t, W\right), z_t\right\}. \tag{2.52}$$

where posterior expectations of trader ij are given by

$$\theta_t | s_{it}, z_t \sim \mathcal{N}\left(\frac{\sigma_\theta^{-2}\bar{\theta} + \beta_{it} s_{it} + \beta_t \sigma_\varepsilon^{-2} z_t}{\sigma_\theta^{-2} + \beta_{it} + \beta_t \sigma_\varepsilon^{-2}}, \frac{1}{\sigma_\theta^{-2} + \beta_{it} + \beta_t \sigma_\varepsilon^{-2}} \right). \tag{2.53}$$

The rest of the proof follows the proof of Proposition 1.1 in Chapter 1 or the proof to Proposition 1 in Albagli, Hellwig, and Tsyvinski (2021). □

Chapter 3

DOES DISPERSED SENTIMENT DRIVE RETURNS, TURNOVER, AND VOLATILITY FOR BITCOIN?

3.1 Introduction

A large literature has studied the effect of investor disagreement on returns for different asset classes and periods and with ambiguous results. Generally, the literature discusses two possible opposing mechanisms: (i) in the presence of short-sale constraints, investor disagreement drives up prices, as optimists hold the assets, and returns will be low, and (ii) investor disagreement represents higher uncertainty and thus warrants a higher return for holding the asset.[1] The first mechanism, known as the differences-of-opinion channel, also predicts high turnover and price volatility when investor disagreement is high.

The differences-of-opinion literature is built on the key theoretical insight that if pessimists cannot participate in the market due to high short-sale costs, the asset price

[1] For an extensive discussion, see Diether, Malloy, and Scherbina (2002).

will be higher than the fundamental value, leading to subsequent low returns. Furthermore, as opinions fluctuate and trade becomes more likely, disagreement leads to high volatility and high turnover. These predictions have been derived from a long theoretical literature, starting from Miller (1977) and Harrison and Kreps (1978), later developed into behavioral agree-to-disagree models such as Hong and Stein (2003), Scheinkman and Xiong (2003), Hong, Scheinkman, and Xiong (2006), and more recently Simsek (2013).[2]

A crucial obstacle to testing the predictions of the differences-of-opinion literature has been that investor sentiment is not directly observable. Different proxies have been used, such as analyst opinions or newspaper articles (Sadka and Scherbina 2007). Nowadays, the availability of extensive online discussions about assets allows us to analyze statements and opinions issued by individuals who are potential investors. A seminal paper following this approach is Antweiler and Frank (2004), who analyzed online posts on Yahoo Finance and Raging Bull to predict market volatility and asset returns. Other papers analyzing asset characteristics using different dictionary-based algorithms are Tetlock (2007), Loughran and McDonald (2011), and Jegadeesh and Wu (2013).

In this paper, we exploit the magnitude of online discussion about a highly speculative asset on which opinions are widely divided (Bitcoin) to test a theory of investor disagreement and short-sale constraints. We scrape millions of online comments across a decade of discussion from a Bitcoin-focused online forum and extract sentiment using the lexicon- and rule-based sentiment algorithm called VADER (Hutto and Gilbert 2014), which is specifically trained for online data sets. Our contribution is to explore the joint time-series behavior of this sentiment measure, as well as its dispersion, on the one hand, and Bitcoin's return, turnover, and price volatility, on the other. Our approach allows us to test the predictions of the differences-of-opinion literature in a rich setting of textual data at daily, weekly, and monthly frequency. We argue that Bitcoin is the ideal asset to test these predictions, as it is complicated to judge its value (Bitcoin will never pay dividends). Therefore, opinions on Bitcoin's value differ widely. Moreover, institutionally and due to substantial price volatility, it is difficult to short Bitcoin.

[2]For a full overview of the differences-of-opinion literature, see Hong and Stein (2007) and Xiong (2013). Simsek (2021) provides an overview of the macroeconomic implications of investor disagreement.

We find that there is a significantly negative predictive relationship between disagreement and the return on holding Bitcoin. Disagreement forecasts negative returns into the future at the daily, weekly, and monthly frequency. This empirical finding is consistent with the theoretical predictions of the differences-of-opinion literature. The effect is especially strong and predicts low returns several months into the future if sentiment and returns are measured at a monthly frequency, which we interpret as overpricing resolving slowly over time. The association between contemporaneous returns, average sentiment, and disagreement is economically significant as well: the adjusted R^2 is 0.33 at monthly frequency. A one standard deviation increase in disagreement leads to a negative return of about -9.2% over the following eight weeks. This is around 12% of the standard deviation of the eight-week returns for Bitcoin.

Although disagreement predicts low returns, which can be interpreted as a sign of overpricing, disagreement is not positively related to contemporaneous or past returns. This finding seems at odds with the usual understanding of the differences-of-opinion channel, predicting that an increase in disagreement first leads to positive returns and overpricing. However, the literature usually assumes that an asset's fundamental value is independent of investor beliefs or disagreement, but this might not be the case for a purely belief-driven asset such as Bitcoin. In this case, the emergence of disagreement could erode the coordination of beliefs among Bitcoin investors, which is key to the asset's value proposition. A slow adjustment of beliefs on the side of optimists can then lead to a situation in which disagreement predicts low returns in the medium term without initially increasing the price.

We study the consequences of the easing of short-sale constraints for Bitcoin starting in December 2017.[3] We find that, as the literature would predict, the effect of disagreement on returns diminishes significantly towards the end of our sample. However, shorting Bitcoin remains expensive and risky, as margin requirements are high compared to other assets, and Bitcoin's price is extremely volatile.

Extending our analysis to volatility and turnover, we find that disagreement has a strong and significant effect on price volatility and turnover growth. Higher disagreement leads to persistently more trading at the same time as a short-lived increase in volatility. These findings generally are consistent with the predictions of the differences-

[3]The CME Group started offering futures contracts for Bitcoin only in December 2017 (CME Group 2017) and options on futures in January 2020 (CME Group 2019).

of-opinion literature. Our findings are also economically significant in this case. In the regression, at a monthly frequency, the adjusted R^2 is 0.06 for turnover growth and 0.34 for volatility.

We contribute by extending the literature about disagreement to a speculative asset with a market capitalization that has increased to over a trillion US dollars since 2010.[4] This makes cryptocurrency assets worth serious scientific attention, despite their quirkiness and novelty. The determinants of their pricing and asset characteristics are interesting in their own right, even from a public policy perspective: a collapse of cryptocurrency prices (e.g., optimists could become disillusioned and leave the market) would destroy immense wealth.

The remainder of the paper is organized as follows. Section 3.2 explains the mechanism behind the results in the differences-of-opinion literature in a stylized model and contrasts it with other possible explanations, namely the idea that disagreement is just a symptom of underlying uncertainty and the that disagreement today is simply driven by low past returns. Section 3.3 details how we collected the data and conducted our sentiment analysis. Then, section 3.4 tests the derived relationships empirically and finds substantial support for the predictions of the theoretical literature. In section 3.5 we interpret our results. Section 3.6 concludes.

3.2 Model

3.2.1 Disagreement in a Differences-of-Opinion Model

We present a simple discrete-time model of heterogeneous beliefs and limits to arbitrage[5] to motivate our empirical analysis. There are overlapping generations of risk-neutral traders indexed by i who each live for two periods and maximize end-of-life consumption.[6] The utility function of trader i is

$$U_{it} = \mathbb{E}_{it}\{C_{it+1}\}. \tag{3.1}$$

[4] As of April 2021. The whole market capitalization of all cryptocurrencies has pushed past two trillion USD.

[5] For a comprehensive review of the theoretical differences-of-opinion literature, see Simsek (2021).

[6] Risk-neutrality is chosen to present the differences-of-opinion channel in the cleanest way. Risk-aversion is considered in section 3.2.2.

When young, traders either buy a long-lived asset from the old or invest in a risk-less bond with return $R > 1$. Traders are split into two groups - optimists and pessimists - who have diverging beliefs about the asset's value, in this case, Bitcoin. Traders have deep pockets, such that optimists have sufficient wealth to buy up the asset supply, and pessimists must stay out of the market due to short-sale constraints. Therefore, the equilibrium price will be determined by the optimists' beliefs only.

The features of the model's long-lived asset capture the essence of Bitcoin in reduced form. Bitcoin investors believe that a coordinated and permanent shift in beliefs will make Bitcoin valuable as a store of value and currency with some positive probability.[7] We denote this absorbing event as A_t if it takes place in period t. Because Bitcoin's protocol limits its supply, the value of one Bitcoin in this event can be derived according to a quantity theory of money type equation, which we denote as $\bar{P}$.[8] The probability of the collective shift is fixed over time to $\mathcal{P}(A_t) = \phi \in (0, 1)$. Once the event has taken place, all traders agree on Bitcoin's price being fixed to $\bar{P}$ for all future periods. As our focus is on explaining short- to medium-term fluctuations in the price of Bitcoin and not the long-term trend, we assume that ϕ is fixed over time.

However, traders believe that the probability ϕ is time-variant. In particular, they have heterogeneous beliefs about $\mathcal{P}(A_{t+1})$ but agree that the probability is fixed to ϕ from $t + 2$ onward. Within each group, beliefs are homogeneous.[9] The beliefs of group i are distributed as $\phi^i_s \overset{iid}{\sim} G(\phi^i_s)$, where $G(\cdot)$ is the continuous cumulative distribution function of ϕ^i_s over the interval $[0, 1]$.[10] We refer to the group that attributes a higher probability to A_{t+1} as *optimists* and the other group as *pessimists* ($\phi^o_{t+1} > \phi^p_{t+1}$).

Due to perfect competition between optimists, the price in period t is

$$P_t = \frac{1}{R} \left(\phi^o_{t+1} \mathbb{E}(P_{t+1}|A_{t+1}) + (1 - \phi^o_{t+1}) \mathbb{E}(P_{t+1}|\neg A_{t+1}) \right), \tag{3.2}$$

where $\mathbb{E}(P_{t+1}|A_{t+1}) = \bar{P}$. Optimists believe that with probability ϕ^o_{t+1} a belief shift

[7] For example, as of 9 June 2021, Bitcoin became legal tender in El Salvador.

[8] We do not model the determinants of $\bar{P}$ explicitly, but instead focus on the relationship between disagreement and the price today, while taking $\bar{P}$ as given.

[9] The assumption of homogeneous in-group beliefs is not crucial and can be replaced with heterogeneous beliefs inside the group. Traders agree to disagree and do not learn from the price. This assumption can be relaxed by assuming that traders are overconfident.

[10] The iid assumption is made for simplicity to highlight that prices and beliefs are eventually mean-reverting, which would also hold if beliefs were somewhat persistent over time.

will take place and they will sell the asset at a price $\bar{P}$ when old. Otherwise, they will sell the asset at a price P_{t+1} that depends on the beliefs of tomorrow's optimists. This leads us to the main prediction of the model.

Proposition 3.1. *Returns* $\frac{P_{t+1}-P_t}{P_t}$ *are* decreasing *in disagreement* $(\phi^o_{t+1} - \phi^p_{t+1})$ *holding the average belief* $(\phi^o_{t+1} + \phi^p_{t+1})/2$ *constant.*

The intuition for this result is that when disagreement is high, the overoptimism of optimists is more severe, which depresses returns going into the future. Naturally, such overoptimism increases the price initially, which leads to the following Corollary.

Corollary 3.1. *Past returns* $\frac{P_t-P_{t-1}}{P_{t-1}}$ *are* increasing *in disagreement* $(\phi^o_{t+1} - \phi^p_{t+1})$ *holding the average belief* $(\phi^o_{t+1} + \phi^p_{t+1})/2$ *constant.*

The model can also be extended to yield predictions for turnover when introducing convex costs for short-sales instead of assuming that such costs are infinite. For simplicity, assume that the costs are sufficiently high such that the marginal buyer remains an optimist.

Proposition 3.2. *With convex short-sale costs, turnover is* increasing *in disagreement* $(\phi^o_{t+1} - \phi^p_{t+1})$ *holding the average beliefs* $(\phi^o_{t+1} + \phi^p_{t+1})/2$ *constant.*

Convex short-sale costs are chosen for simplicity to relate disagreement and turnover. Generally, increased disagreement positively influences the perceived gains from trade, leading to additional traders entering the market which can increase turnover even when the amount of short-sale is exogenously fixed per trader.

Volatility has a more ambiguous relationship with disagreement than returns or turnover. For example, if traders were long-lived and beliefs fully persistent, prices would be constant for any level of disagreement. As beliefs are short-lived in our stylized model, the relationship between volatility and disagreement is, in principle, non-linear. If today's optimists are pessimistic relative to the average optimists over time $(\phi^o_{t+1} < \mathbb{E}(\phi^o | \phi^o > \phi^p))$, then higher disagreement can move today's prices closer to the historical average while keeping the average belief constant, decreasing expected volatility.

Still, expected volatility can be positively related to today's disagreement when the distribution of beliefs $G(\phi^i_s)$ is subject to variance shocks. In this case, more disagreement today can indicate more volatile beliefs in the future, which increases expected price volatility. Therefore, we expect volatility to be positively related to disagreement.

Similarly, our predictions for returns and turnover can be extended dynamically if we think about an increase in disagreement stemming from such a mean-reverting variance shock to $G(\phi_s^i)$. Then, disagreement is persistent and overpricing due to disagreement does not resolve immediately, leading to protracted negative returns. Also turnover and volatility remain alleviated for several periods.

Before we turn to our empirical approach, note that the model is deliberately kept simple to provide intuition for the main mechanisms. Moreover, the model is stationary conditional on the base probability of adoption ϕ, whereas asset prices are usually non-stationary as returns follow random walks. Therefore, our simple model is best used to explain fluctuations that happen around the trend in Bitcoin's price. As the price and turnover clearly show non-stationary behavior in Figure 3.1, we will use returns, the growth rate of turnover and dispersion in hourly returns as left-hand-side variables in our empirical analysis.[11]

Before moving to our empirical analysis, we present two alternative models through which the relationship between disagreement and returns can be interpreted.

3.2.2 Disagreement as Uncertainty

Large disagreement among traders can be a sign of fundamental uncertainty. When uncertainty increases, risk-averse traders require higher future returns to absorb the risk, which leads to a fall in the price today.

To capture this intuition, consider a model with overlapping generations of representative traders[12] with CARA-utility

$$U_t = 1 - \exp\left(-\gamma W_{t|t+1}\right). \tag{3.3}$$

[11] Although beliefs almost solely drive price movements and trading in Bitcoin, our approach of extracting sentiment from text is not well-suited to explain the long-term price movements in Bitcoin. Sentiment analysis is more appropriate to measure how prevalent relatively positive or negative sentiment is in a given moment in time, which can be an important determinant for short- to medium-run price movements.

[12] The same result would also hold when considering a model with a mass of heterogeneously informed traders as in Grossman and Stiglitz (1980). In that case, an increase in fundamental risk leads to higher disagreement and posterior uncertainty. A decrease in the precision of private signals can have similar effects as shown in Chapter 2.

where $\gamma \geq 0$ is the coefficient of absolute risk aversion and $W_{t|t+1}$ is end-of-period wealth of the representative trader born in period t, who is free to borrow or lend at interest rate $R > 1$. Without loss of generality, initial wealth is normalized to zero. Every period, the old trader sells a single risky asset to the young trader. Otherwise, the asset characteristics are unchanged.

As before, traders believe that the probability of the adoption event A_{t+1} is time-variant. But this time, their beliefs are uncertain. At the beginning of the period, the representative trader draws a belief over ϕ_{t+1} with finite mean and positive variance. Beliefs are iid across generations.

The price P_t is derived from the representative trader being indifferent between holding the asset or lending out P_t at interest rate R,

$$1 - \mathbb{E}_t \left\{ \exp\left(-\gamma P_{t+1}\right) \right\} = 1 - \exp\left(-\gamma R P_t\right), \tag{3.4}$$

leading to

$$P_t = \frac{-\log \mathbb{E}_t \left\{ \exp\left(-\gamma P_{t+1}\right) \right\}}{\gamma R}. \tag{3.5}$$

The price P_t depends on the representative trader's expectations about the probability of adoption ϕ_{t+1}. Due to risk-aversion, a mean-preserving increase in uncertainty about ϕ_{t+1} must lead to a lower price P_t, which is captured in the following proposition.

Proposition 3.3. *Returns* $\frac{P_{t+1}-P_t}{P_t}$ *are* increasing *in the representative trader's variance of beliefs on the probability of the adoption event* ϕ_{t+1}.

As before, a fall in today's price due to higher uncertainty must also mean that past returns were negative, which leads to the following corollary.

Corollary 3.2. *Past returns* $\frac{P_t-P_{t-1}}{P_{t-1}}$ *are* decreasing *in the representative trader's variance of beliefs on the probability of the adoption event* ϕ_{t+1}.

It follows that viewing disagreement as a proxy of uncertainty leads to exactly opposite predictions on the relationship between disagreement and returns compared to the differences-of-opinion model. When traders are risk-averse, they require a higher return when absorbing greater risks, leading to falling prices. In contrast, in the presence of overconfidence and short-sale constraints, higher disagreement means that price-setting optimists are increasingly over-optimistic, inflating the price today and leading to low future returns.

3.2.3 Sentiment and Disagreement as a Side-Show

As we set out to derive a proxy for sentiment and disagreement from posts in an online forum, reverse causality is a plausible concern. Such posts may merely react to price movements but not reveal any information that could be useful to predict future returns. In the following, we suppose that sentiment is a function of lagged returns.

$$\text{Sent}_{it} = \alpha_i + \sum_{s=0}^{S} \gamma_{is} \left(\frac{P_{t-s} - P_{t-1-s}}{P_{t-1-s}} \right) + \varepsilon_{it} \tag{3.6}$$

where $S < \infty$ and $\varepsilon_{it} \overset{iid}{\sim} \mathcal{N}(0, \sigma_\varepsilon^2)$. Finally, we use in our analysis average sentiment and the dispersion in sentiments as disagreement, formally

$$\text{Sent}_t = \sum_i \text{Sent}_{it}. \tag{3.7}$$

$$\text{Dis}_t = \sqrt{Var(Sent_{it})} \tag{3.8}$$

Naturally, sentiment should react positively to current, and past returns ($\forall s : \gamma_{is} > 0$) as investors profit from positive returns.[13] Less clear is the relationship between disagreement and past returns. One possible explanation draws on *confirmation bias*. Following this idea, investors in Bitcoin may be likely to disregard information that does not match their prior.

Suppose that investors are split into two groups. The first group consists of dogmatically optimistic traders (high α_o), who do not revise their beliefs in the face of new information ($\gamma_{os} \approx 0$).[14] The second group is composed of less optimistic traders ($\alpha_p < \alpha_o$) with more flexible beliefs ($\gamma_{ps} > 0$). As a result, both groups hold more

[13] A positive relationship may also be plausible when interpreting sentiments as expectations about future returns. Greenwood and Shleifer (2014) show that investors increase their expectations of future returns after positive past returns.

[14] It is also possible to assume that there are dogmatic pessimists, but attributing dogmatism to optimists is in line with anecdotal evidence of a fraction of Bitcoin investors who buy Bitcoin and hold it for extended periods irrespective of news. Moreover, dogmatic pessimists should eventually leave the market.

similar beliefs when returns are positive and disagree more intensely when returns are negative. Therefore, we would expect to see a negative relationship between past returns and disagreement.

To derive predictions about the predictability of returns using sentiment and disagreement, we consider two returns processes. First, prices may follow a random walk and returns are white noise. In this case, sentiment or disagreement cannot forecast returns as they are not correlated with the innovations to the price. Second, returns may be autocorrelated, which can lead to sentiment and disagreement predicting future returns.[15] Nonetheless, sentiment and disagreement should lose their predictive power when controlling for lagged returns.

3.3 Data

We use publicly available data from the *Kraken.com* exchange for the opening and closing price of Bitcoin and an aggregated measure of turnover across all major exchanges from *coinmarketcap.com*. We compute daily returns by dividing the difference between closing and opening prices by the opening price and turnover as the daily dollar volume divided by Bitcoin's total market capitalization.[16] We compute a volatility measure as the standard deviation of hourly returns in a given day, week, or month.

We relate Bitcoin's market characteristics to sentiment changes among Bitcoin investors. For this purpose, we scrape the Bitcoin-related online-forum *bitcointalk.org* using the python package *Scrapy*. In particular, we scrape all threads and comments from the *Speculation* subforum, which most closely covers discussions on Bitcoin's price movements and expectations about future price developments. We gathered 1,482,589 comments that were posted between 18 October 2010 and 21 April 2021.

We gathered comments from 54,173 unique accounts, of which 7,183 accounts opened discussion topics. Posting activity follows a power law, as the top ten percent of most active accounts (more than 32 posts) produce more than 84% of all content. In contrast, the median number of posts per account is three. At the same time, no single

[15] Positive autocorrelation may arise when new information is only gradually incorporated (McQueen, Pinegar, and Thorley 1996), whereas negative autocorrelation may stem from overreaction to new information (Lo and MacKinlay 1990).

[16] Since crypto exchanges are open 24/7, the opening and closing prices are the earliest and latest price available in a specific period according to UTC.

account dominates the discussion, as the most active account wrote 1.3% of all posts (19,758 in total or five posts per day). Overall concentration is low with a Herfindahl index of 0.0011.[17] According to *bitcointalk.org*,[18] there are in total over three million registered users, and more than a million page views a day. With this reach, *bitcointalk.org* is an important medium in discussions related to cryptocurrencies.[19]

We use all comments with non-zero valence in our analysis, as the Speculation subforum is already focused on Bitcoin's price movements. The forum allows users to quote other comments in their posts. We filter out such repetitions as quotes and keep only the new part of each post.[20] We run our main specification starting 1 January 2014, as the number of posts per day reaches a higher and more stable level from 2014 on. Figure 3.C.2 provides a word cloud with the most commonly used words of a random sample of 10,000 comments.

Figure 3.1 summarizes the time series of the main variables at a weekly frequency: the price level, turnover, price volatility measured as the standard deviation of hourly returns, average sentiment, and dispersion in sentiments. Additionally, the right upper panel displays the number of posts on the Speculation subforum of *bitcointalk.org*. As both the price and turnover display a clear trend, we will use their growth rates. All time series have substantial time variation, which we exploit in our empirical analysis. Bitcoin's price follows a distinct boom and bust cycle. Turnover and volatility increase when Bitcoin's price increases or decreases rapidly (e.g., the boom leading up to December 2017 or the short-term bust in 2019.). Posts per week are also cyclical and peak at around 7,000 posts per week at the beginning of 2018. Finally, although average sentiment and the standard deviation of sentiment is relatively noisy week-to-week, both time series show persistence at lower frequencies.

[17] We show the time series of the number of active users in Figure 3.C.1a, which is positively correlated to the overall number of posts. In Figure 3.C.1b, we show that the Herfindahl Index is stable over time and decreased during the surge in activity in 2018.

[18] See `https://bitcointalk.org/index.php?action=stats` (accessed 25 June 2021, 20:00).

[19] For example, today's second-largest cryptocurrency *Ether* and its Initial Coin Offering were first announced on *bitcointalk.org* in January 2014: `https://bitcointalk.org/index.php?topic=428589.0` (accessed 25 June 2021, 20:00).

[20] We do not attempt to weigh posts according to importance (e.g., through their number of views or quotes), but instead attribute the same weight to every post. Although an abundance of quotes potentially reflects the greater importance of the quoted post, we find that filtering out quotes increases the explanatory power of our sentiment and disagreement measure in all regressions.

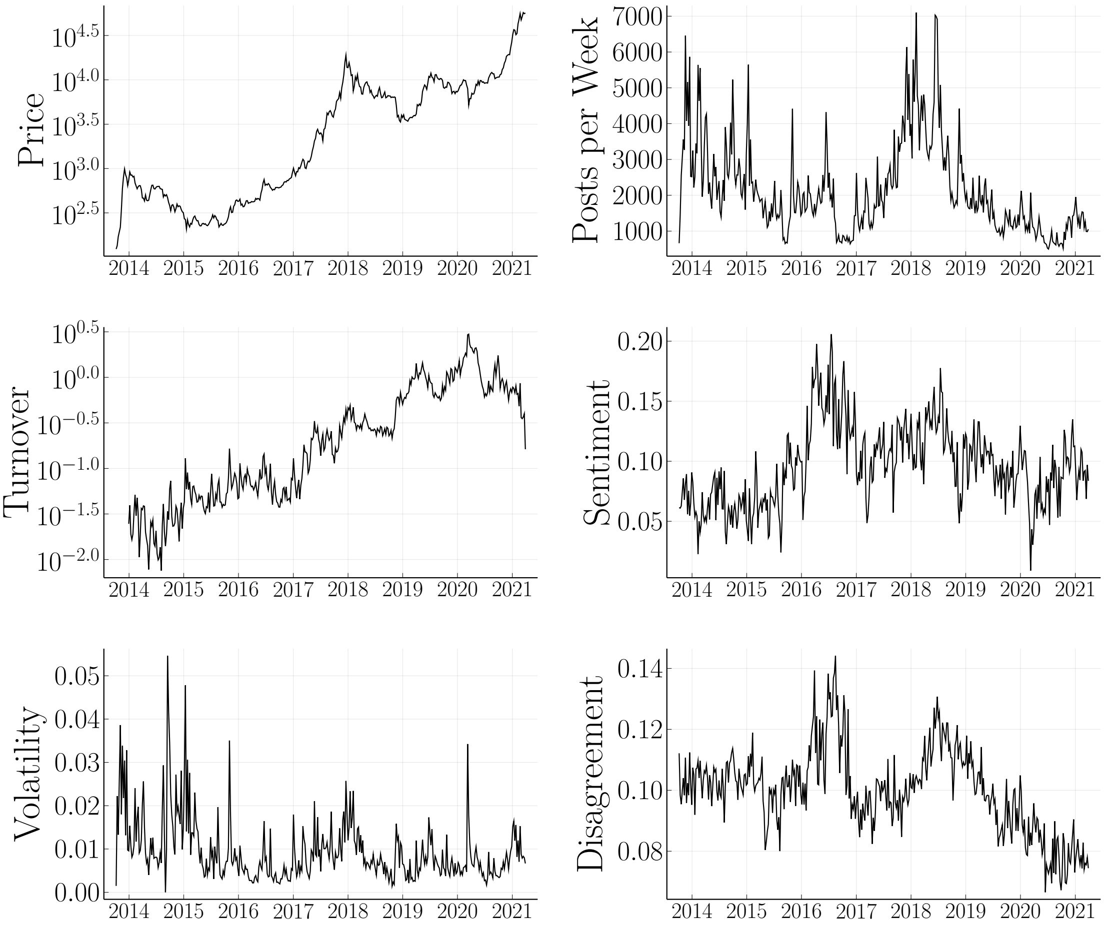

Figure 3.1: Overview over the Main Variables at Weekly Frequency.

Notes: On the left: returns (difference between closing and opening price divided by opening price), turnover (volume divided by market capitalization) and volatility (standard deviation of hourly returns). On the right: number of posts, sentiment (the mean of the comment sentiment distribution) and disagreement (the standard deviation of the comment sentiment distribution).

Table 3.1 provides summary statistics (mean, standard deviation, the first and ninth decile cut-offs, as well as the median) for our main variables of interest: sentiment, disagreement, returns, turnover growth, and volatility. All statistics are shown at daily, weekly, and monthly frequency. Strikingly, mean returns for Bitcoin are quite high, with 7.6% monthly. Volatility is also high, with an average standard deviation of hourly returns of around 0.9 percentage point.

	Frequency	Mean	SD	Q10	Median	Q90
Sentiment	daily	0.097	0.042	0.044	0.096	0.15
	weekly	0.096	0.034	0.054	0.095	0.14
	monthly	0.096	0.031	0.056	0.094	0.136
Disagreement	daily	0.099	0.019	0.076	0.099	0.122
	weekly	0.1	0.014	0.081	0.1	0.118
	monthly	0.1	0.013	0.08	0.101	0.116
Return	daily	0.3%	4 %	-3.9 %	0.2%	4.4 %
	weekly	1.6%	10.7%	-12.4%	1.1%	15.6%
	monthly	7.6%	23.2%	-18.9%	6.2%	37.4%
Turnover Growth	daily	7 %	49 %	-31.5%	-1.3%	45.2%
	weekly	5.5 %	35.8%	-29.3%	-1.5%	46.5%
	monthly	9.3 %	40.6%	-28.7%	0 %	64.1%
Volatility	daily	0.8 %	0.77%	0.22 %	0.58%	1.57%
	weekly	0.91%	0.68%	0.34 %	0.72%	1.68%
	monthly	0.99%	0.64%	0.42 %	0.78%	1.84%

Table 3.1: Summary Statics of Sentiment, Disagreement, Returns, Turnover Growth, and Volatility of Bitcoin.

Notes: Mean, standard deviation, first decile, median and ninth decile of the main variables. Statistics for returns, volatility, and turnover growth are in percentage points.

3.3.1 Sentiment Analysis using VADER

We use a lexicon and rule-based algorithm called VADER (**V**alence **A**ware **D**ictionary and s**E**ntiment **R**easoner) for the sentiment analysis.[21] The underlying lexicon and algorithm are specialized for the analysis of social media posts (see Hutto and Gilbert 2014, for a comparison with other lexica).[22]

A sentiment lexicon is a mapping from "tokens" (words, stems of words, abbreviations, etc.) to a numerical indicator of sentiment. Each token carries a certain valence (negative, neutral, or positive sentiment) irrespective of context. These valence intensities were generated by letting ten independent human raters rate tokens. The final valence is the average of the individual ratings (Wisdom of the Crowd approach). All human raters had been pre-screened, trained, and quality checked. Following this approach, over 9000 tokens were rated on a scale from "[-4] Extremely Negative" to "[4] Extremely Positive" with an option to rate the token as "[0] Neutral." Already existing established lexicons inspired the list of tokens (e.g., LIWC, ANEW, and GI) to which Western-style emoticons (e.g., ":-)"), sentiment-related acronyms and initialism (e.g. "LOL", "ROFL") and commonly used slang (e.g., "nah," "meh") were added. After dropping tokens that ended up with a neutral mean-sentiment rating or a standard deviation of individual ratings higher than 2.5, about 7500 tokens were left and rated on the -4 to $+4$ scale.

Although relying on a lexicon for sentiment analysis, VADER is not a bag-of-words algorithm that neglects the syntax and order of words. Instead, VADER employs five simple rules to improve its sentiment ratings for whole sentences. First, punctuation is included by using the exclamation point (!) as an intensifier. Secondly, capitalization increases the sentiment intensity. Thirdly, modifiers are used to adjust the intensity. With the corresponding valence between -1 (very negative) and 1 (very positive) computed by VADER, "Bitcoin has a bright future" (0.44) is less intense than "Bitcoin has a very bright future" (0.49) and more intense than "Bitcoin has a somewhat bright future" (0.38). Fourthly, the conjunction "but" is used to signal a reversal of semantic

[21] A detailed description of VADER can be found on Github: `https://github.com/cjhutto/vaderSentiment` (accessed 25 June 2021, 20:00).

[22] The authors show that VADER can produce valence ratings with high correlation to human mean-sentiment ratings. In particular, run on a corpus of over 4000 Tweets, sentiment, as calculated by VADER, had the largest correlation (0.88) and R^2 (0.77) to the human mean-sentiment rating.

orientation. For example, "Bitcoin had a great year, but has a lot of problems" (-0.25) conveys negative sentiment, although the initial statement is positive. Lastly, the three words before a sentiment-laden token are included in the sentiment rating to check for words that flip the semantic orientation. For example, "Bitcoin does not have a great future" (-0.34) conveys negative sentiment, although "great" carries positive sentiment.

To sum up, VADER is an appropriate sentiment analysis tool for the domain of our investigation. We use the "compound" measure, a weighted average of sentiment normalized to values between -1 (extremely negative) and 1 (extremely positive). As suggested by the package authors, we compute the sentiment index for each comment on the sentence level and use the mean to compute comment-level sentiment. Finally, we aggregate the sentiment data at different frequencies and use the mean to measure the level of daily, weekly, and monthly sentiment. We use the standard deviation of variance as a proxy for disagreement among investors.

3.4 Empirics

Our empirical approach is to extract a sentiment measure from comments on *bitcointalk.org* and use this measure as a proxy for beliefs about the success of Bitcoin (ϕ^i_{t+1} in the model). In particular, we think of our sentiment measure as being relative to some time-variant base level of expectations (e.g., a time-variant ϕ). In that way, high sentiment can be interpreted as expectations of positive returns at any point in time. Henceforth, we refer to the valence measure as computed by VADER from each comment simply as sentiment.

Whereas we capture an average stance of sentiment through the first moment the sentiment distribution, we define disagreement as the dispersion in sentiment. If comments with positive and negative sentiment are posted during the same period, we interpret such dispersion as a sign of high disagreement. We use these measures to analyze the effect of disagreement, conditional on average sentiment, on the return, turnover growth, and volatility of Bitcoin.

All our regressions are summarized by the following equation,

$$X^j_{t+s} = \alpha^j_s + \sum_{l=0}^{L} (\beta^{j,l}_{\mu,s} \text{Sentiment}_{t-l} + \beta^{j,l}_{\sigma,s} \text{Disagreement}_{t-l}) + \varepsilon_{t+s}. \tag{3.9}$$

We use returns, turnover growth, and the dispersion in hourly returns (volatility) as the left-hand-side variable X^j_{t+s} where j stands for each different variable. We run the regression at different leads and, for returns, lags s. Moreover, we also use long-horizon returns with overlapping observations as the left-hand-side variable, in which case X^j_{t+s} stands for the return between the beginning of period $t+1$ and the end of period $t+s$. We include up to L lags of sentiment and disagreement and estimate (3.9) for each variable at daily, weekly, and monthly frequency. Throughout, we apply HAC-robust standard errors following Newey and West (1987). To address concerns due to the persistence of our regressors, we repeat our forecasting regressions with confidence intervals computed according to Campbell and Yogo (2006).[23] Additionally, for the regression with long-horizon returns, we adjust our confidence intervals according to Hjalmarsson (2011), which additionally increases the bandwidth as our forecasting-horizon lengthens.

3.4.1 Return Regressions

We set out to predict returns of Bitcoin through sentiment and disagreement. As a first step, we present evidence that the price of Bitcoin is indeed predictable while not taking a stance on the specific predictor. For this purpose, we use the variance ratio test of Lo and MacKinlay (1988). The idea of the test is that if prices move randomly, the variance of returns should increase linearly in the horizon. If this assumption is violated, returns are not random and can potentially be predicted.

	2 Lags	3 Lags	4 Lags	5 Lags	6 Lags	10 Lags
Daily	-0.58	-0.43	-0.30	-0.22	-0.03	0.45
Weekly	0.83	1.28	1.55	1.41	1.36	1.55
Monthly	1.81*	1.74*	1.96*	1.93*	1.82*	1.54

Table 3.2: Lo and MacKinlay (1988) Variance Ratio Test for Return Predictability.

Notes: We find that Bitcoin's returns are predictable at lower frequencies. Critical values are as for the two-sided t-test. *: p <0:10, **: p <0:05, ***: p <0:01

As shown in Table 3.2, we find evidence that Bitcoin's returns are indeed predictable

[23]We provide autocorrelation functions of our variables in Section 3.C.1.

at monthly frequency, which we confirm in our regression analysis. Before turning to our regression results, we recap the predictions for the relationship between disagreement and returns for each model in Table 3.3.

	Past Returns	Future Returns
Differences-of-Opinion	>0	<0
Uncertainty	<0	>0
Side-Show	>0	0

Table 3.3: Predictions for the Relationship between Disagreement and Returns.

The presented models of differences-of-opinion in the presence of short-sale constraints and disagreement as uncertainty yield exactly opposite predictions regarding the relationship of returns and disagreement. If optimists price the asset as in the differences-of-opinion model, an increase in disagreement, while holding average sentiment constant, leads to an increase in overpricing. Such overpricing then is predictive of lower returns in the future due to the mean-reversion of overoptimism. In contrast, an increase in disagreement can be viewed as a sign of uncertainty, which leads to risk-averse traders requiring higher returns, thus lowering the price today. Finally, in the model in which sentiment is simply a reflection of past returns, disagreement should have no predictive power when controlling for past returns.

Regressions

As a first step, we estimate the relationship between sentiment and disagreement in period t, and returns in $t-1$, t, and $t+1$ to distinguish between the different models.

	Returns t-1 daily	Returns t daily	Returns t+1 daily	Returns t-1 weekly	Returns t weekly	Returns t+1 weekly	Returns t-1 monthly	Returns t monthly	Returns t+1 monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment	0.73***	0.70***	0.09	3.11***	4.59***	1.04*	8.84***	13.38***	4.84**
	(0.08)	(0.09)	(0.08)	(0.48)	(0.55)	(0.61)	(2.18)	(2.02)	(2.22)
Disagreement	−0.45***	−0.42***	−0.08	−2.86***	−3.52***	−1.14**	−9.56***	−12.21***	−7.69***
	(0.07)	(0.07)	(0.07)	(0.49)	(0.47)	(0.56)	(2.16)	(1.92)	(2.41)
Constant	0.93***	0.85***	0.48	13.26***	13.74***	6.88**	52.85***	59.67***	51.59***
	(0.32)	(0.32)	(0.37)	(3.18)	(2.97)	(3.41)	(14.41)	(13.77)	(16.51)
N	2,635	2,634	2,633	378	377	376	88	87	86
R^2	0.03	0.03	0.001	0.09	0.18	0.01	0.17	0.34	0.09
Adjusted R^2	0.03	0.03	−0.0001	0.09	0.18	0.01	0.15	0.33	0.06

Notes: Returns are the growth rate between the opening and closing price in percentage points. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.4: Regression of Leads and Lags of Returns on Sentiment and Disagreement.

Notes: Sentiment is positively related with returns, whereas disagreement is negatively related to returns. All relationships are stronger at lower frequencies.

Our results are not strictly in line with the proposed models. As seen in Table 3.4, the contemporaneous effect of disagreement on the return is negative, as is the effect one period ahead. For example, a one standard deviation increase in disagreement in month t decreases the return in the subsequent month by 7.7 percentage points. The negative contemporaneous relationship is what we would have expected from the model with risk-averse traders, while the negative predictive effect is in line with the differences-of-opinion model. Therefore, neither model explains the empirical results exactly.

To test whether sentiment and disagreement contain information beyond what is reflected by past returns, we report in Table 3.5 the one-period-ahead predictive regression while controlling for lagged returns. If the "disagreement as a side-show" model was true, we would expect that sentiment and disagreement do not predict returns and that past returns significantly forecast future returns. However, this turns out to be incorrect: past returns have little explanatory power for future returns and the coefficients for sentiment and disagreement remain close to what they where in the regression without lags in Table 3.4. Thus, all in all, none of the three most suggestive models seem to provide a comprehensive explanation for our empirical results.

	Returns t+1								
	daily	daily	daily	weekly	weekly	weekly	monthly	monthly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment	0.10	0.09	0.08	0.94	0.71	0.68	3.98	4.08	4.10
	(0.08)	(0.09)	(0.09)	(0.59)	(0.59)	(0.62)	(2.61)	(2.58)	(2.71)
Disagreement	−0.08	−0.08	−0.07	−1.07*	−0.87	−0.84	−6.91**	−7.01**	−7.16**
	(0.07)	(0.07)	(0.07)	(0.57)	(0.61)	(0.65)	(2.95)	(2.99)	(3.14)
Return t	−0.02	−0.03	−0.03	0.02	0.02	0.02	0.06	0.06	0.07
	(0.03)	(0.03)	(0.03)	(0.07)	(0.07)	(0.07)	(0.11)	(0.10)	(0.11)
Return t-1		0.003	0.002		0.05	0.05		−0.01	−0.02
		(0.03)	(0.03)		(0.07)	(0.07)		(0.10)	(0.11)
Return t-2			0.01			0.03			−0.03
			(0.02)			(0.06)			(0.04)
Constant	0.48	0.47	0.45	6.62*	5.83	5.55	47.79**	48.32**	49.83**
	(0.36)	(0.37)	(0.37)	(3.49)	(3.75)	(3.92)	(18.80)	(18.79)	(20.16)
N	2,624	2,616	2,608	375	374	373	86	86	86
R^2	0.001	0.001	0.001	0.01	0.01	0.02	0.09	0.09	0.10
Adjusted R^2	−0.0000	−0.0004	−0.001	0.004	0.003	0.002	0.06	0.04	0.04

Notes: Returns are the growth rate between the opening and closing price in percentage points. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.5: Regression of Returns on Sentiment and Disagreement Controlling for Lagged Returns.

Notes: The coefficients on all lagged returns are insignificant and including lags reduces the adjusted R^2. The sign on the coefficient of disagreement and sentiment remains stable, but including lags marginally decreases the size and increases the standard errors, leading to a decrease in significance.

An additional testable prediction whether sentiment and disagreement predict returns further into the future. Indeed, if beliefs and, therefore, disagreement are persistent, overpricing as in the differences-of-opinion model might resolve only slowly. As a result, disagreement should predict negative returns for multiple periods ahead.

To this end, we estimate (3.9) for many periods ahead and present the results in Figure 3.2. We find that disagreement has a strongly persistent negative effect on future returns, which is more pronounced at lower frequencies, cancelling out higher-frequency noise. At monthly frequency, disagreement has a significantly negative effect on returns at 95% confidence up to five months into the future. The plots at the daily and weekly frequency show that this effect is not driven by outliers, but that returns are consistently negative. Thus, disagreement predicts lower returns for up to half a year ahead, which, through the lens of the differences-of-opinion model, suggests that prices take a long time to revert back from overoptimistic levels.[24]

As seen in Figure 3.1, sentiment and disagreement are relatively persistent and may feature stochastic trends. Indeed, Augmented Dickey-Fuller tests on disagreement rejects the null of the series featuring a unit root at the daily frequency, however does not reject the null at the weekly and monthly frequency, highlighting persistent low-frequency movements. As is well known, very persistent regressors can lead to t-statistics that are too large. Therefore, we provide estimates of confidence intervals that take into account the persistence of regressors.

To address these concerns, we employ the methodology in Campbell and Yogo (2006) to compute confidence intervals that are robust to the presence of persistent regressors and show the results in Figure 3.3, where the black line shows the central value of the confidence interval. Different to before, we use univariate local projections of returns on disagreement, as Campbell and Yogo (2006) is only applicable to univariate predictive regressions. We find that our results hold at the 90% confidence level.[25]

[24]This long-lasting effect is also found for other markets. Disagreement in the stock market may forecast lower returns for up to a year (Diether, Malloy, and Scherbina 2002).

[25]We also run our regressions in first differences and show the results in Figures 3.B.1 and 3.B.2. Although significance suffers due to the introduction of additional noise through differencing, the basic results continue to hold. However, we focus on the regression in levels due to its straightforward interpretation.

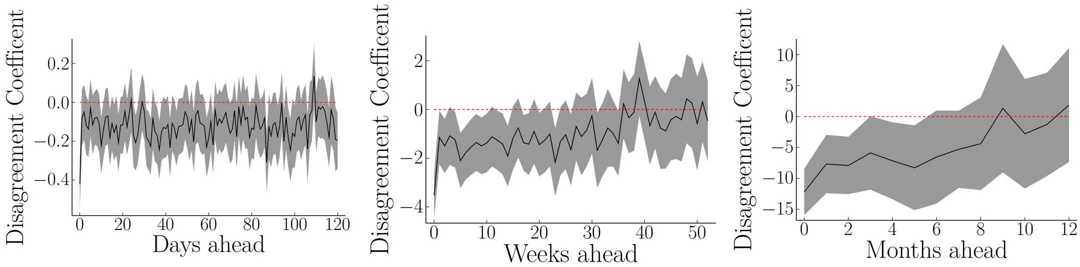

(a) Daily Returns. (b) Weekly Returns. (c) Monthly Returns.

Figure 3.2: Local Projections of Returns on Disagreement controlling for Sentiment with HAC-Robust Standard Errors.

Notes: The shown estimates are the coefficients on disagreement when estimating (3.9) for leads of returns. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

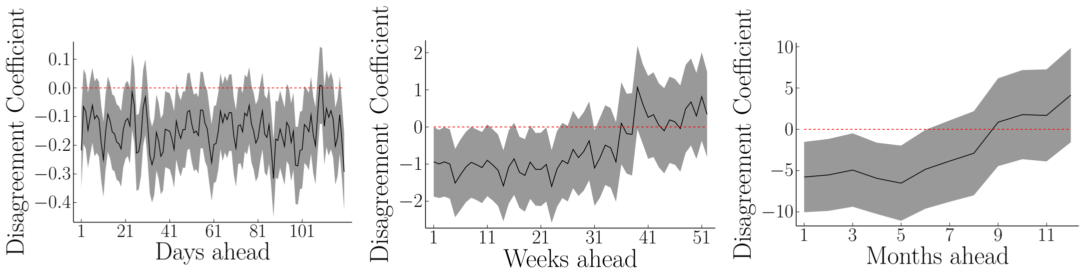

(a) Daily Returns. (b) Weekly Returns. (c) Monthly Returns.

Figure 3.3: Univariate Robust Local Projection of Returns on Disagreement.

Notes: In a univariate regression, we find that disagreement predicts negative returns for several periods at all frequencies. 90% confidence intervals according to Campbell and Yogo (2006).

Another way to express our findings on the persistence of disagreement shocks is to focus on the cumulative returns over multiple months. Similar to Figure 3.3, in Figure 3.4 we run univariate predictive regressions with disagreement in period t as the predictor for the cumulative returns between the opening price in $t+1$ and the closing price in $t+s$. We compute first the confidence intervals as in Campbell and Yogo (2006) and additionally widen them by a factor $\sqrt{s}$ as suggested by Hjalmarsson (2011). Without this adjustment, the implied confidence intervals would be too narrow for long-horizon regression with overlapping observations.

We find that the effect of disagreement on cumulative returns is close to being significant at the 90% confidence level for most horizons and significant at the 90% level for some horizons (e.g., five to nine weeks ahead). This loss in significance compared to Figure 3.3 is somewhat puzzling, but we attribute it to the conservative computation of the confidence intervals. Note also that these are univariate regressions. Given that disagreement and sentiment are positively correlated, and that sentiment is positively related to returns, we would expect that the effect of disagreement on returns is *biased towards zero* when not controlling for sentiment.

For the estimates that are significant at 90% confidence, we find that a one standard deviation shock on disagreement leads to a eight-week return that is about 9.2 percentage points lower, which corresponds to about 13% of the standard deviation of eight-week returns for Bitcoin. Furthermore, though insignificant, we find that the effect of disagreement on cumulative returns only reverts after more than twelve months.

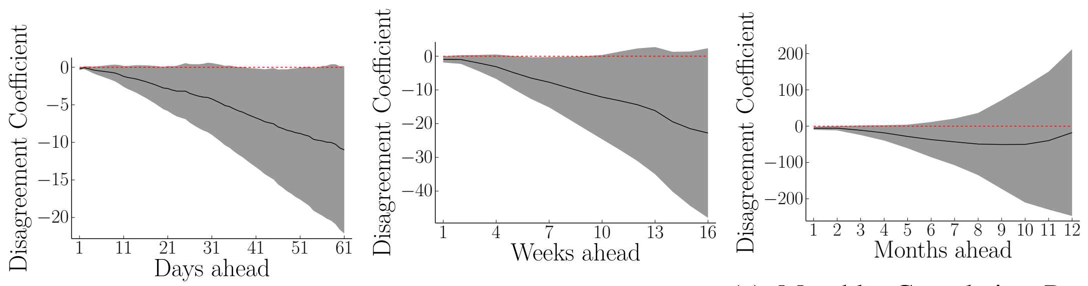

(a) Daily Cumulative Returns.

(b) Weekly Cumulative Returns.

(c) Monthly Cumulative Returns.

Figure 3.4: Long-Horizon Regression of Cumulative Returns on Disagreement.

Notes: We find that a positive one standard deviation shock to disagreement has long-lasting negative effects on returns. Confidence intervals are at 90% according to Hjalmarsson (2011).

Finally, since we hypothesize that the predictive power of disagreement is due to disagreement (and thus overoptimistic beliefs) being persistent, we test whether disagreement remains predictive when controlling for contemporaneous disagreement. In Table 3.B.3, which is shown in the appendix, we run our regression on contemporaneous returns while including contemporaneous sentiment and disagreement, as well as three lags at each frequency. We find that lagged disagreement is insignificant at daily and monthly frequency yet significantly positive at weekly frequency. At the same time, contemporaneous the coefficient of disagreement remains significantly negative at all frequencies.

Again, this finding does not fit well in any of the suggested theories. Whereas the differences-of-opinion model predicts that lagged disagreement inflates yesterday's price with a negative effect on today's return, viewing disagreement as a sign of increasing uncertainty should depress yesterday's price with a positive effect on today's return. We do not find strong evidence for either story.

3.4.2 Turnover and Volatility Regressions

We present our results for the contemporaneous effect of sentiment and disagreement on turnover growth and price volatility in Table 3.6. We find that sentiment is significantly associated with contemporaneous turnover growth and volatility of Bitcoin at all frequencies. Moreover, our results grow in magnitude and explanatory power when looking at lower frequencies, i.e., longer-lasting increases in sentiment or disagreement have greater effects.

Our results are in line with the theoretical predictions of the differences-of-opinion model: disagreement increases trading activity and drives up price volatility. This last finding suggests that increases in disagreement indicate more underlying volatility of beliefs.

On the other hand, we find that sentiment and disagreement do not have much predictive power in explaining turnover growth and volatility one period ahead, as shown in Table 3.7. Disagreement predicts volatility only at the daily and weekly frequency and does not predict turnover growth at all. Note that the lack of mean-reversion in turnover growth means that the effect of disagreement on turnover is relatively persistent. This finding is also confirmed at longer horizons in the local projections in Figure 3.5 when focusing on the effect of disagreement. We provide univariate local projects with confidence intervals according to Campbell and Yogo (2006) in Figure 3.C.3. Tables 3.B.1 and 3.B.2 in the appendix show the contemporaneous and one-period-ahead regression in first-differences for volatility and turnover growth.

	Turnover Growth t			Volatiliy t		
	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)
Sentiment	-6.59^{***}	-7.20^{***}	-12.23^{***}	-0.22^{***}	-0.36^{***}	-0.42^{***}
	(1.09)	(2.48)	(4.58)	(0.03)	(0.07)	(0.12)
Disagreement	5.11^{***}	4.91^{***}	9.19^{**}	0.10^{***}	0.21^{***}	0.23^{**}
	(0.92)	(1.77)	(4.23)	(0.03)	(0.07)	(0.11)
Constant	-4.68	-9.06	-23.09	0.80^{***}	0.47	0.50
	(3.53)	(8.78)	(27.85)	(0.11)	(0.34)	(0.55)
N	2,644	378	86	2,616	378	87
R^2	0.02	0.04	0.08	0.08	0.25	0.36
Adjusted R^2	0.02	0.03	0.06	0.08	0.24	0.34

Notes: Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.6: Contemporaneous Regressions of Turnover Growth and Volatility on Sentiment and Disagreement.

	Turnover Growth t+1			Volatiliy t+1		
	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)
Sentiment	−0.09	2.02	2.88	-0.20^{***}	-0.25^{***}	-0.25^{*}
	(0.99)	(1.79)	(4.16)	(0.03)	(0.07)	(0.14)
Disagreement	−0.62	−2.41	−3.03	0.08^{***}	0.10^{*}	0.09
	(0.86)	(1.68)	(5.22)	(0.03)	(0.06)	(0.11)
Constant	10.48^{**}	16.78^{*}	23.61	0.81^{***}	0.88^{***}	1.03^{*}
	(4.20)	(10.17)	(35.02)	(0.11)	(0.32)	(0.57)
N	2,643	377	86	2,615	377	86
R^2	0.0002	0.005	0.01	0.06	0.11	0.12
Adjusted R^2	−0.001	−0.001	−0.02	0.06	0.11	0.10

Notes: Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.7: Predictive Regressions of Turnover Growth and Volatility on Sentiment and Disagreement.

Notes: Disagreement predicts lower returns at weekly and monthly frequency. Turnover remains alleviated after an increase in disagreement, whereas the effect of disagreement on volatility disappears after a week.

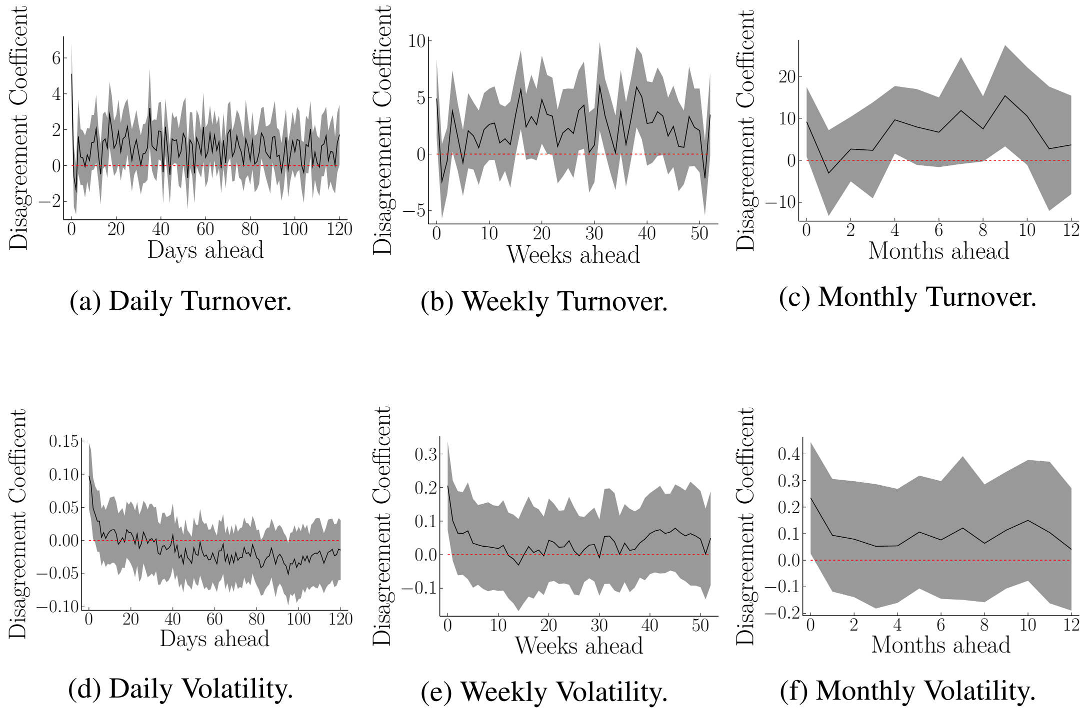

(a) Daily Turnover. (b) Weekly Turnover. (c) Monthly Turnover.

(d) Daily Volatility. (e) Weekly Volatility. (f) Monthly Volatility.

Figure 3.5: Local Projections of Turnover Growth and Volatility on Disagreement with HAC-robust Standard Errors.

Notes: The shown estimates are the coefficients on disagreement when estimating (3.9) for leads of turnover growth and volatility. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

Although the main focus of our analysis is on the effect of disagreement on returns, turnover growth, volatility, we provide for completeness the corresponding local projections focusing on the effect of sentiment in Figure 3.C.4 with HAC-robust standard errors. Figure 3.C.5 shows the results for the univariate regressions with sentiment as the predictor with confidence intervals computed according to Campbell and Yogo (2006).

3.4.3 Introduction of CME Futures

The presented framework analyzed the effect of disagreement in the presence of short-sale constraints. A major event in this context is the introduction of futures trading contract at the Chicago Mercantile Exchange (CME) on 18 December 2017 (CME Group

2017) and options on futures contract started on 12 January 2020 (CME Group 2019). The introduction of futures contracts and options does not only make markets more complete but should also substantially alleviate short-sale constraints.

The differences-of-opinion model predicts that an easing of short-sale constraints through the introduction of futures and options can eliminate the effect of disagreement, as pessimists can voice their opinion by selling short. To study this prediction, we estimate (3.9) contemporaneously and year-by-year. We focus on the contemporaneous specification, as the regression with lagged sentiment and disagreement in Table 3.B.3 suggests that the negative effect of disagreement on future returns stems contemporaneous disagreement. We show the coefficient on disagreement with 95% error bands in Figure 3.6. We also report the monthly specification for completeness, although twelve observations per year are arguably too little to draw solid inference.

We find that the coefficient and its error bands on disagreement change over time. In particular, the negative effect of disagreement on returns is particularly large in 2017 and 2018 at the daily and weekly frequency, whereas no effect can be measured in 2015. Potentially, this result can be related to insufficient variance in disagreement and returns in 2015, such that some episodes can be characterized as more or less speculative.

Starting from 2017, the estimate of the coefficient of disagreement for returns tends toward zero. Moreover, the estimate for 2020 is insignificantly different from zero at all frequencies, and the difference between the coefficients in 2016 and 2020 is statistically significant as shown in Table 3.B.4. This finding can be interpreted as short-sale constraints having sufficiently eased since 2018 such that disagreement does not lead to overpricing anymore.[26]

We also study the role of sentiment more generally over time by showing the R^2 of estimating (3.9) year-by-year in Figure 3.7. Generally, we find that sentiment and disagreement play a larger role at lower frequencies as demonstrated by higher R^2 measures. We also see here that the importance of sentiment changes over time. Although the coefficient on disagreement tends towards zero at the end of the sample, the explanatory power of sentiment and disagreement combined remains high. This is not surprising, as Bitcoin remains a speculative asset also when short-sales are permitted.

[26]We conduct a similar analysis for sentiment in Figure 3.C.6. Similarly, we find that the effect of sentiment changes over time, but does not go to zero towards the end of our sample.

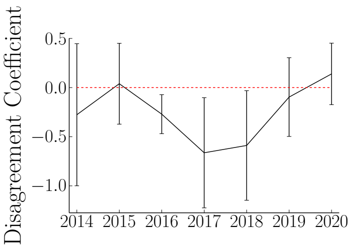

(a) Daily Returns.

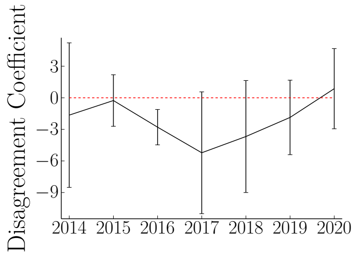

(b) Weekly Returns.

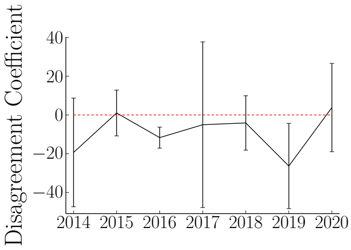

(c) Monthly Returns.

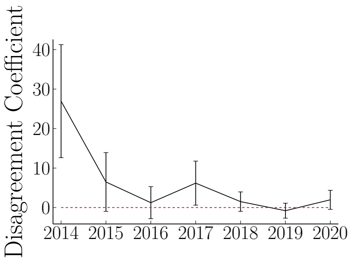

(d) Daily Turnover.

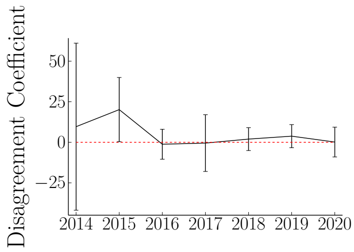

(e) Weekly Turnover.

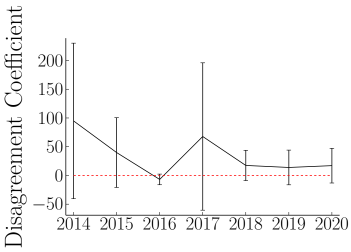

(f) Monthly Turnover.

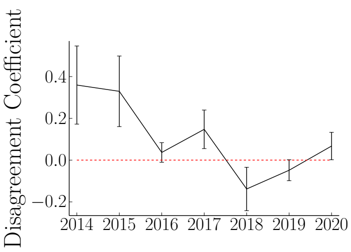

(g) Daily Volatility.

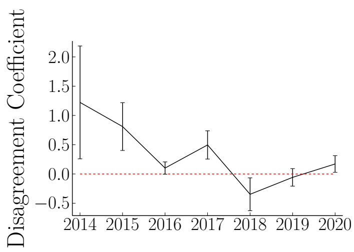

(h) Weekly Volatility.

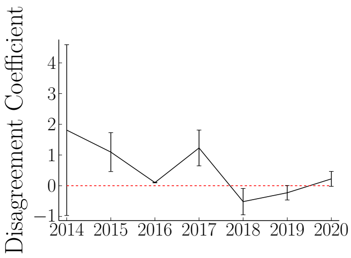

(i) Monthly Volatility.

Figure 3.6: Year-by-Year Coefficient of Disagreement for the Contemporaneous Regression.

Notes: We find that the contemporaneous effect of disagreement is relatively stable over time. Towards the end of the sample, the negative correlation between disagreement and returns vanishes. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

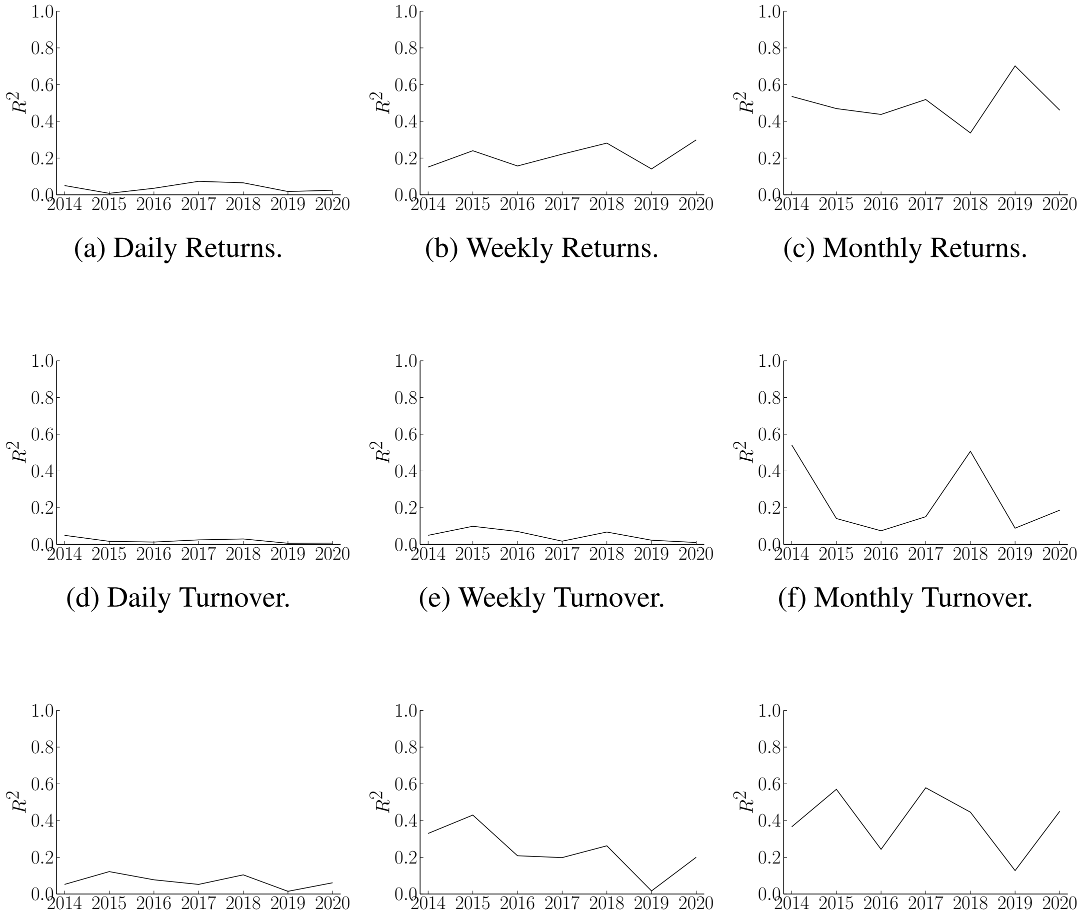

(a) Daily Returns. (b) Weekly Returns. (c) Monthly Returns.

(d) Daily Turnover. (e) Weekly Turnover. (f) Monthly Turnover.

(g) Daily Volatility. (h) Weekly Volatility. (i) Monthly Volatility.

Figure 3.7: Year-by-Year R^2 for the Contemporaneous Regression.

3.5 Discussion

We can summarize our main results as follows: disagreement does predict lower subsequent returns for up to five months into the future and the explained variation is not small (but also not suspiciously large) with an adjusted R^2 of around 6% for the one-month-ahead regression,[27] high disagreement does also come with a negative contemporaneous effect on returns. Robustness checks, such as correcting for regressor persistence and controlling for lagged regressors and returns, respectively, do not alter this result. Furthermore, high disagreement comes with contemporaneously higher price volatility and turnover growth, however, for these variables we do not see a strong predictive effect.

3.5.1 Disagreement and Overpricing

Our main result that returns can be predicted by disagreement is in line with the differences-of-opinion argument that we characterize in 3.2.1: we find that high dispersion in our sentiment measure (i.e., disagreement is high) forecasts long-lasting negative returns at the daily, weekly, and monthly frequency. Through the lens of the differences-of-opinion literature, we would interpret this result as buyers' overoptimism decaying slowly, which could only occur if pessimists' ability or willingness to short-sale is limited. Moreover, we also find positive effects of high disagreement on turnover and volatility, which can be interpreted as market participants trading more often when their opinions are more dispersed.

However, our findings also differ from the standard differences-of-opinion story as portrayed in 3.2.1. Following a standard interpretation of the channel, the fundamental value of the asset is orthogonal to investors' beliefs. Therefore, an increase in overoptimism as reflected by high disagreement leads to overpricing. Thus, from the viewpoint of the differences-of-opinion literature, we would have hypothesized that the contemporaneous effect of disagreement on returns is positive. However, we find that it is significantly negative. This result would be expected if disagreement was just a symptom of underlying uncertainty, as we show in 3.2.2, or if disagreement was caused by negative past and contemporaneous returns, as we discuss in 3.2.3. However, with these two explanations we should not see that disagreement predicts negative returns. If the

[27] The analysis at higher frequencies picks up more noise, which leads to our R^2 being generally largest at monthly frequency.

subsequent negative returns are interpreted as a correction of overpricing caused by mean-preserving disagreement, why do we not see a price increase of equal magnitude leading to this overpricing?[28]

Our empirical results suggest that there has to be some mechanism that goes beyond these straightforward stories that the previous literature has discussed.[29] In the case of an asset entirely supported by beliefs, such as Bitcoin, disagreement could be associated with an erosion of the coordination that makes the asset valuable in the first place. This view reconciles our finding that disagreement is a sign of overpricing with the fact that we do not observe a price increase in the first place. A rise in disagreement lowers the "fair" or "objective" value of the asset while also keeping the price from falling immediately, leading to overpricing.

At the same time, we find that disagreement increases exactly when the price is already falling, for example, as shown in Figure 3.8 during the 2018 bust. Therefore, negative returns could themselves increase disagreement. A possible explanation is that traders filter negative news, which coincide with negative returns, heterogeneously. Consider as an extreme example that traders can be split into two groups: dogmatic believers and skeptics. Traders who believe dogmatically in Bitcoin might not correct their beliefs in the face of negative news. In contrast, other, less convinced traders quickly correct their beliefs downward and sell when the price starts falling. The loss of potential buyers and users of Bitcoin leads to a fall in Bitcoin's medium-term value.

This narrative is in sharp contrast to a view of an asset's fundamental value being unaffected by beliefs or disagreement. In general, asset prices may influence a firm's fundamentals in the presence of financial frictions, such that overvaluation due to the optimism of buyers can fix another inefficiency. In our case, the force behind said overvaluation - disagreement - is possibly detrimental to the asset's fundamentals, which leads to further negative returns.

[28]We only find a positive relationship between past disagreement and contemporaneous returns at weekly returns in Table 3.B.3. Still, the effect is much smaller than the subsequent predicted negative returns as in Figure 3.2 and 3.3.

[29]See Diether, Malloy, and Scherbina (2002) for an overview of the proposed channels.

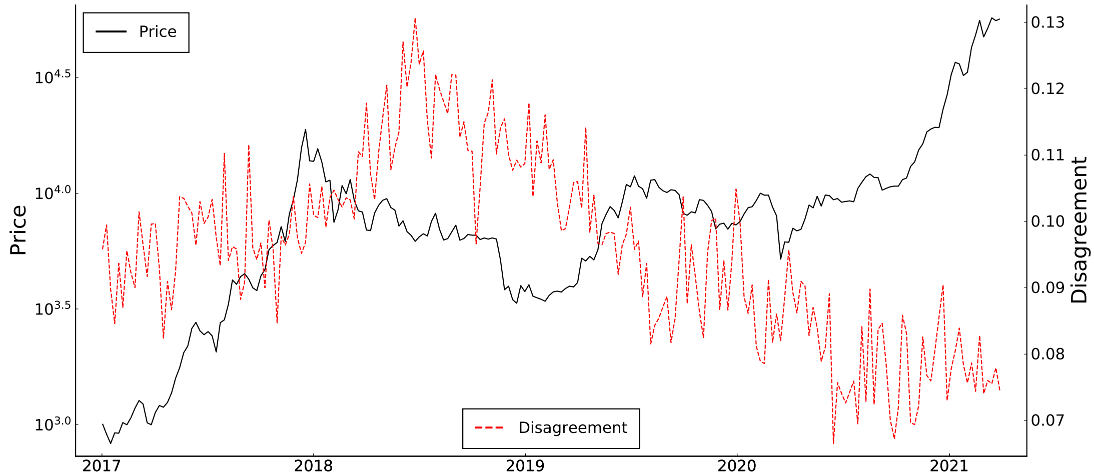

Figure 3.8: Disagreement Peaked after Returns Turned already Negative in 2018.

3.5.2 Easing of Short-Sale Constraints

According to the literature on differences-of-opinion, disagreement leads to overpricing in the presence of short-sale constraints. Being able to short allows pessimists to trade on their belief, which reduces asset prices and offsets the influence of optimists. For Bitcoin, short-sales were difficult for two reasons: (i) the lack of financial instruments, especially through established exchanges accessible to institutional investors, and (ii) the extremely high volatility and explosive price behavior. Since the end of 2017, we have seen the gradual introduction of shorting instruments for Bitcoin, which addresses the first point. It is now possible to borrow Bitcoin on large exchanges,[30], the Chicago Mercantile Exchange introduced future contracts for Bitcoin in December 2017 (CME Group 2017) and options on futures contracts in January 2020 (CME Group 2019) which enabled especially institutional investors to bet on a falling price of Bitcoin.

The introduction of CME's futures contracts coincided with a steep fall in the price of Bitcoin, which supports the narrative that over-optimistic buyers inflated Bitcoin's price, and the introduction of futures contracts eased short-sale constraints considerably. This easing should also diminish the effect of disagreement that we find in our analysis. Indeed, if we compare the coefficient of disagreement in the regression on contemporaneous returns in the years 2016 and 2020, we find that the effect is significantly reduced

[30]The annualized interest rates are between around 12% as of April 2021 on *Kraken.com* while requiring 20% collateral in the form of cash or cryptocurrencies.

as in Table 3.B.4. However, we find that the effect of disagreement on returns is also small in 2014 or 2015, well before short-sale constraints were eased. These small coefficients can potentially be understood as a sign that our disagreement channel is not strong at all times, as our disagreement measure appears to be noisy without a clear trend in 2014 and 2015.

In Figure 3.9 we repeat the local projection from Figure 3.3 to see whether the effect of disagreement on future returns changes after the introduction of futures contracts. In this analysis, we test whether weekly disagreement observed in 2019 predicts returns that stretch until the end of 2020. Indeed, we find that in this time-frame in which futures contract were already well-established, disagreement predicts initially positive returns, which is more in line with the model with risk-averse traders. In reality, aspects of both models are likely to be relevant, and, therefore, the shift of return predictions from negative to positive suggests that relaxed short-sale constraints led to the differences-of-opinion channel being less important.

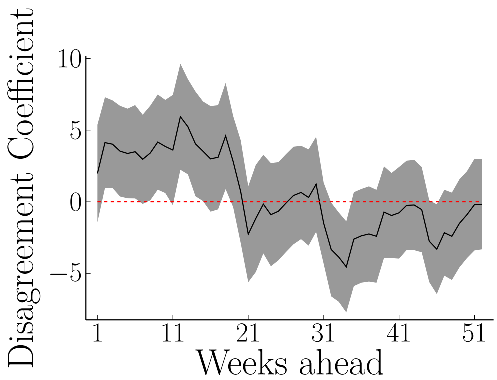

Figure 3.9: Univariate Robust Local Projection of Returns on Disagreement from 2019.

Notes: We find that in contrast to our earlier results, disagreement predicts initially positive returns. 90% confidence intervals according to Campbell and Yogo (2006).

Although our analysis suggests that short-sale constraints loosened, shorting Bitcoin remains costly due to limits-to-arbitrage: the maintenance volatility scan exemplifies this for the CME future contract, which as of April 2021 stands around 60%, much higher than on future contracts for other assets.[31] Relatedly, as of April 2021, the CME requires a maintenance amount for a futures contract over five Bitcoin that corresponds

[31]The maintenance volatility scan is the highest level of change that is "most likely" to occur with the

to 36% to 40% of the current spot price, which is much larger than the respective 5% for S&P 500 futures at CME. In other words, investors need to lock up larger amounts of capital to bet on price movements of Bitcoin than for other assets, which limits the ability of investors to take on larger positions.

3.5.3 Future Research

Our analysis leaves scope for future research. First, we base our measurement of sentiment on the VADER algorithm. This algorithm is trained on online comments but has not been tested on a cryptocurrency domain. Our sentiment measure is positively correlated with returns, suggesting that the algorithm performs reasonably well for our purposes. Moreover, we use this sentiment measure as a proxy for disagreement around a time-variant level of beliefs. Ideally, we would observe beliefs and disagreement directly, which is impossible without detailed surveys. More can be done to tune the sentiment algorithm towards this specific domain.

Secondly, we are observing only a subset of potential Bitcoin investors. The online forum that we analyzed does not contain institutional investors, nor can we be certain that it represents a balanced sample of all bitcoin investors. Investors who visit an online forum and post many comments about Bitcoin might differ from those who trade quietly. Still, relevant for our purposes is only that users of *bitcointalk.org* representatively reflect the sentiments present in the general population of potential bitcoin investors. Future research could seek to analyze the beliefs and trading of institutional investors.

Finally, we did not control for who posted the comments in our regressions: were there some participants who posted many more comments than others and how did the make-up and breadth of the discussion participants change over the years? We only tested that the concentration of comments across individuals at any given time is not too high to introduce an obvious bias to our results (as we report in Figures 3.C.1a and 3.C.1b). Our empirical analysis does not exploit additional information on commenters' identities. However, they represent an interesting topic for future research.

---

underlying volatility affecting each future option's price. If the volatility of an asset is high, the margin requirement will be high as well. For comparison, S&P 500 Futures with a duration of seven months have a volatility scan of 25%.

3.6 Conclusion

We performed sentiment analysis on posts from the online forum *bitcointalk.org* to obtain a measure of investors' sentiment regarding the prospects of Bitcoin. We used the dispersion of these sentiment data points as a proxy for disagreement to explore the empirical predictions of the differences-of-opinion literature. We find that disagreement indeed predicts lower returns while being related to higher turnover and volatility, confirming the theoretical predictions of our model. The most striking result is that disagreement predicts negative returns for up to five months into the future, pointing towards a slow correction of large overvaluations. Moreover, sentiment and disagreement play a large role in explaining returns, with an adjusted R^2 of 0.33 in the regression on contemporaneous returns at a monthly frequency.

We study the change in effects of disagreement after CME introduced futures contracts in December 2017 (CME Group 2017). We find that the effect of disagreement significantly diminishes in the years after the introduction, in line with the view that a combination of disagreement and short-sale constraints is necessary to generate overpricing and subsequent predictable lower returns. An important departure from the previous empirical literature is that we find no strong evidence of a positive effect of disagreement on the contemporaneous or past return of Bitcoin. We hypothesize that disagreement may erode the coordination, which is the foundation of Bitcoin's value proposition, leading to a price decrease.

Future research could seek to understand better the effects of changes in beliefs at the investor level. Also, the network structure of online discussions can be exploited to understand better belief formation and the impact of different kinds of discussions on the price of Bitcoin. One could explore, for example, the narratives that generate disagreement between participants through topical analysis. A better understanding of the determinants of investor disagreement can help refine asset pricing theory.

Appendix

3.A Proofs

Proof of Proposition 3.1. Since the price tomorrow is independent of disagreement or the average belief today given that ϕ is constant over time and beliefs are iid across group of traders and time, it is sufficient to show that the optimists' belief ϕ^o_{t+1} and, therefore, P_t are increasing in disagreement $\phi^o_{t+1} - \phi^p_{t+1}$ when holding the average belief $\frac{\phi^o_{t+1}+\phi^p_{t+1}}{2}$ constant. The optimists' belief can be written as

$$\phi^o_{t+1} = \underbrace{\frac{\phi^o_{t+1} + \phi^p_{t+1}}{2}}_{\textit{Average Belief}} + \frac{1}{2}\underbrace{(\phi^o_{t+1} - \phi^p_{t+1})}_{\textit{Disagreement}}. \tag{3.10}$$

Indeed, if the right-hand-side increases due to an increase in disagreement and the average belief stays constant, the left-hand-side must increase. The increase in optimists' belief ϕ^o_{t+1} leads to higher price P_t as follows from (3.2) and $\bar{P} > \mathbb{E}(P_{t+1}|\neg A_{t+1})$, which lowers future returns. □

Proof of Corollary 3.1. Follows from the proof of Proposition 3.1, except that disagreement increases today's price and, therefore, increases past returns $\frac{P_t - P_{t-1}}{P_{t-1}}$. □

Proof of Proposition 3.2. Since the price is set by optimists with $\phi^o_{t+1} > \phi^p_{t+1}$ and $\bar{P} > \mathbb{E}(P_{t+1}|\neg A_{t+1})$, it must be that the pessimists' valuation is below the current price P_t, i.e.,

$$V^p_t = \frac{1}{R}\left(\phi^p_{t+1}\bar{P} + (1 - \phi^p_{t+1})\mathbb{E}(P_{t+1}|\neg A_{t+1})\right) < P_t, \tag{3.11}$$

Therefore, the maximization problem of pessimists is

$$\max_{d_t \leq 0} \quad d_t(P_t - V_t^p) - c(d_t) \tag{3.12}$$

with the solution $d_t^* : c'(d_t^*) = P_t - V_t^p$. It follows that pessimists short the asset more if V_t^p is lower due to more pessimistic beliefs ϕ_{t+1}^p, which decreases in disagreement as apparent in:

$$\phi_{t+1}^p = \underbrace{\frac{\phi_{t+1}^o + \phi_{t+1}^p}{2}}_{\textit{Average Belief}} - \frac{1}{2} \underbrace{\left(\phi_{t+1}^o - \phi_{t+1}^p\right)}_{\textit{Disagreement}}. \tag{3.13}$$

Therefore, an increase in disagreement while holding the average belief constant means that pessimists are even more pessimistic, which lowers V_t^p and increases the short positions d_t^* and turnover. □

Proof of Proposition 3.3. Tomorrow's price is independent of the belief of the young representative trader given that ϕ is constant over time and beliefs are iid. It remains to show that the price today is decreasing in the uncertainty of the young representative trader when the average belief is held constant.

The young representative trader views ϕ_{t+1} as a random variable X where $|\mathbb{E}(X)| < \infty$ and $Var(X) \in (0, \infty)$. Denote alternative beliefs that are more uncertain than X but have the same mean as $Z = X + Y$ where $\mathbb{E}(Y) = 0$ and $Var(Y) \in (0, \infty)$ and Y is independent of X. It is sufficient to show that a representative trader with more uncertain beliefs has a lower utility holding the asset than a more certain trader with the same mean belief. Given that the utility function (3.3) is concave in ϕ,

$$\begin{aligned} \mathbb{E}\left(U_t(Z)\right) &= \mathbb{E}\left(U_t(X+Y)\right) \\ &\overset{\text{L.I.E.}}{=} \mathbb{E}\left(\mathbb{E}\left(U_t(X+Y)|X\right)\right) \\ &\overset{\text{Jensen's}}{<} \mathbb{E}\left(U_t\left(\mathbb{E}\left(X+Y|X\right)\right)\right) \\ &= \mathbb{E}\left(U_t\left(X\right)\right), \end{aligned} \tag{3.14}$$

where the utility function U_t is written as a function of ϕ. Since a trader with more uncertain beliefs is worse off compared to a trader with more certain beliefs but the same mean, it follows that the price P_t must fall to restore the indifference as in (3.4). As a result, future returns increase in uncertainty. □

Proof of Corollary 3.2. Follows from the proof of Proposition 3.3, except that disagreement decreases today's price and, therefore, decreases past returns $\frac{P_t - P_{t-1}}{P_{t-1}}$. □

3.B Tables

	Return Diff t			Turnover Growth t			Volatiliy Diff t		
	daily	weekly	monthly	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment Diff	0.59***	5.97***	15.68***	−6.06***	−10.16***	−15.26***	−0.07***	−0.21***	−0.28***
	(0.12)	(0.88)	(2.56)	(0.98)	(2.27)	(4.11)	(0.01)	(0.04)	(0.05)
Disagreement Diff	−0.20*	−2.32***	−4.31*	4.65***	7.16***	11.83***	0.02*	0.15***	0.14***
	(0.10)	(0.68)	(2.59)	(0.98)	(2.12)	(4.27)	(0.01)	(0.04)	(0.04)
Constant	0.005	0.02	0.28	7.04***	5.52***	10.03***	−0.0004	0.0000	−0.01
	(0.03)	(0.25)	(2.53)	(0.76)	(1.38)	(3.78)	(0.01)	(0.01)	(0.05)
N	2,623	376	87	2,642	378	86	2,595	378	87
R^2	0.01	0.18	0.29	0.02	0.11	0.21	0.01	0.16	0.25
Adjusted R^2	0.01	0.17	0.27	0.02	0.11	0.19	0.01	0.15	0.23

Notes: Returns are the growth rate between the opening and closing price in percentage points. Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Changes in sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. We compute the difference in a variable as today's value minus yesterday's value. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.B.1: Contemporaneous Regressions in First Differences.

Notes: An alternative to providing robust standard errors is differencing variables until both independent and dependent variables are stationary. We find that the results remain generally unchanged, as changes in disagreement are negatively related to returns and positively related to turnover growth and changes to volatility.

	Return Diff t+1			Turnover Growth t+1			Volatiliy Diff t+1		
	daily	weekly	monthly	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment Diff	-0.62^{***}	-4.11^{***}	-11.54^{***}	1.13	2.51^{*}	9.94^{**}	-0.01	0.06^{**}	0.17^{***}
	(0.12)	(0.88)	(2.66)	(0.97)	(1.44)	(4.08)	(0.01)	(0.03)	(0.06)
Disagreement Diff	0.20^{*}	2.47^{***}	2.38	0.97	-1.83	-4.23	0.02	-0.08^{**}	-0.13^{**}
	(0.12)	(0.74)	(2.51)	(0.96)	(1.57)	(3.47)	(0.01)	(0.03)	(0.06)
Constant	-0.003	0.10	0.56	7.01^{***}	5.33^{***}	8.87^{**}	-0.001	-0.003	-0.01
	(0.03)	(0.38)	(2.69)	(0.75)	(1.58)	(4.02)	(0.01)	(0.02)	(0.05)
N	2,622	375	86	2,641	377	86	2,594	377	86
R^2	0.01	0.10	0.16	0.001	0.01	0.07	0.001	0.02	0.12
Adjusted R^2	0.01	0.09	0.14	0.0002	0.002	0.05	-0.0000	0.02	0.10

Notes: Returns are the growth rate between the opening and closing price in percentage points. Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Changes in sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. We compute the difference in a variable as today's value minus yesterday's value. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.B.2: One-Period-Ahead Regressions in First Differences.

Notes: We find that the effects of sentiment and disagreement revert in comparison to Table 3.B.1. Note that the effect of disagreement of the change in returns one-period-ahead is insignificant and smaller in magnitude than the contemporaneous effect, pointing to a protracted negative effect of disagreement on returns.

	Returns t			Turnover Growth t			Volatiliy t		
	daily	weekly	monthly	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment t	1.08^{***}	8.32^{***}	20.29^{***}	-8.74^{***}	-15.87^{***}	-22.62^{***}	-0.16^{***}	-0.38^{***}	-0.49^{***}
	(0.10)	(0.79)	(3.18)	(1.33)	(3.83)	(6.92)	(0.02)	(0.06)	(0.10)
Sentiment t-1	-0.30^{***}	-4.19^{***}	-8.00^{**}	4.49^{***}	16.24^{***}	21.84^{***}	-0.09^{***}	0.02	0.06
	(0.10)	(0.95)	(3.41)	(1.39)	(3.49)	(8.02)	(0.02)	(0.03)	(0.09)
Sentiment t-2	−0.13	1.00	0.15	1.27	−0.79	−10.48	−0.02	0.004	0.08
	(0.10)	(0.97)	(3.68)	(1.36)	(3.36)	(7.18)	(0.02)	(0.05)	(0.07)
Sentiment t-3	-0.24^{**}	-2.32^{***}	−2.10	−1.46	−3.88	5.44	−0.02	0.01	−0.04
	(0.11)	(0.74)	(3.27)	(1.23)	(3.18)	(8.66)	(0.02)	(0.05)	(0.08)
Disagreement t	-0.39^{***}	-4.89^{***}	-11.60^{***}	6.75^{***}	12.34^{***}	23.06^{**}	0.08^{***}	0.29^{***}	0.45^{***}
	(0.08)	(0.87)	(3.93)	(1.19)	(3.84)	(9.91)	(0.02)	(0.07)	(0.12)
Disagreement t-1	0.09	2.30^{**}	−2.26	-2.65^{**}	-11.44^{***}	-20.20^{**}	0.05^{***}	−0.02	−0.11
	(0.09)	(1.02)	(4.67)	(1.20)	(4.31)	(8.29)	(0.02)	(0.05)	(0.11)
Disagreement t-2	0.05	−0.49	0.36	-3.34^{***}	-7.93^{**}	−7.09	0.01	−0.07	−0.07
	(0.09)	(1.00)	(3.94)	(1.22)	(3.14)	(10.61)	(0.01)	(0.06)	(0.10)
Disagreement t-3	−0.09	0.41	3.22	1.90	10.07^{***}	8.43	0.004	−0.02	−0.12
	(0.09)	(0.77)	(3.56)	(1.36)	(3.52)	(8.70)	(0.01)	(0.04)	(0.08)
Constant	1.13^{***}	12.65^{***}	54.00^{***}	3.23	−4.02	−4.46	0.71^{***}	0.65^{*}	1.02^{*}
	(0.42)	(3.42)	(15.92)	(4.07)	(8.52)	(32.69)	(0.16)	(0.39)	(0.61)
N	2,628	377	87	2,638	378	86	2,610	378	87
R^2	0.04	0.27	0.41	0.03	0.14	0.25	0.09	0.26	0.40
Adjusted R^2	0.04	0.26	0.35	0.03	0.12	0.18	0.09	0.24	0.34

Notes: Returns are the growth rate between the opening and closing price in percentage points. Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.B.3: Contemporaneous Regressions Controlling for Lags of Sentiment and Disagreement.

Notes: The results are broadly consistent with mean-reversion. Whereas sentiment in period t is positively related to returns in period t, sentiment in period $t - 1$ has a *negative* effect on returns in t, albeit the coefficient is at most half as large as the contemporaneous effect. Disagreement exhibits mean-reversion for returns only at weekly frequency.

	Returns t		
	daily	weekly	monthly
	(1)	(2)	(3)
Sentiment	0.93^{***}	6.78^{***}	20.75^{***}
	(0.11)	(0.90)	(3.29)
Disagreement	-0.42^{***}	-4.90^{***}	-17.46^{***}
	(0.10)	(0.90)	(3.05)
Year 2017	1.60	2.87	−91.41
	(1.45)	(21.14)	(146.18)
Year 2018	0.59	−13.11	−98.69
	(1.80)	(20.73)	(73.90)
Year 2019	−0.91	−19.83	-119.45^{**}
	(1.22)	(12.87)	(59.61)
Year 2020	-1.47^{*}	-31.84^{**}	-132.51^{**}
	(0.83)	(12.98)	(58.94)
Year 2021	2.20	179.96^{***}	108.93
	(3.03)	(44.45)	(79.60)
Disagreement x Year 2017	−0.12	0.35	15.04
	(0.28)	(3.17)	(19.50)
Disagreement x Year 2018	−0.10	1.71	11.17
	(0.30)	(2.59)	(8.53)
Disagreement x Year 2019	0.29	3.20^{*}	16.71^{**}
	(0.23)	(1.81)	(8.17)
Disagreement x Year 2020	0.52^{***}	5.77^{***}	21.37^{**}
	(0.18)	(2.19)	(9.38)
Disagreement x Year 2021	−0.31	-32.35^{***}	−18.16
	(0.77)	(8.07)	(13.48)
Constant	−0.33	13.78^{**}	67.32^{***}
	(0.54)	(6.12)	(20.81)
N	1,917	274	63
R^2	0.05	0.27	0.54
Adjusted R^2	0.04	0.23	0.43

Notes: Returns are the growth rate between the opening and closing price in percentage points. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.B.4: Contemporaneous Regression with Year-FEs.

Notes: The effect of disagreement on contemporaneous returns is significantly different in 2020 compared to 2016. The sign on the coefficient is positive taking together the base effect and interaction term. However, disagreement has again a more negative effect in 2021.

3.C Figures

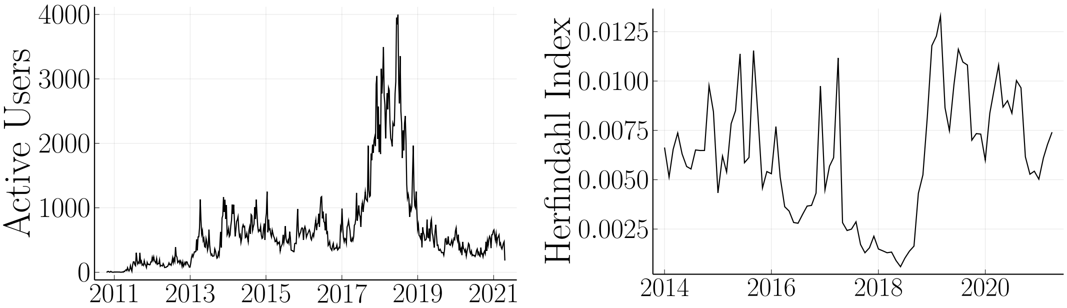

(a) Weekly Active Users.

(b) Monthly Authorship Concentration.

Figure 3.C.1: Activity and Concentration Graphs for *bitcointalk.org*.

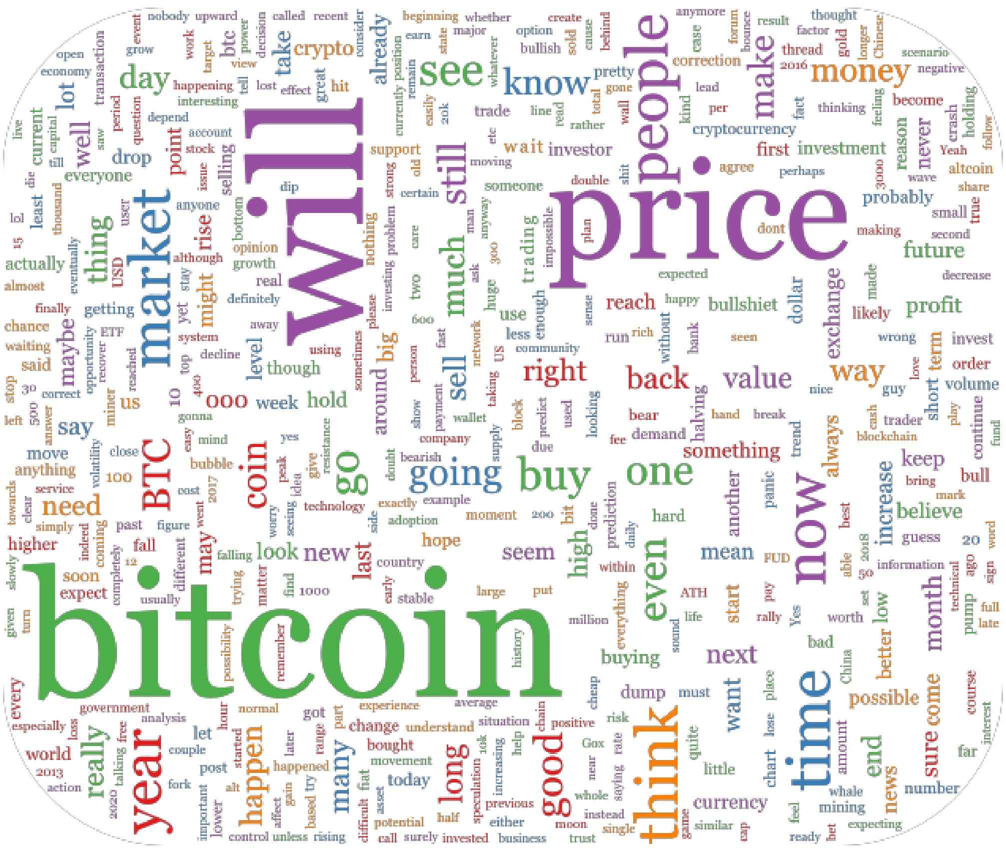

Figure 3.C.2: Most Commonly Used Words out of a Random Sample of 10,000 Comments.

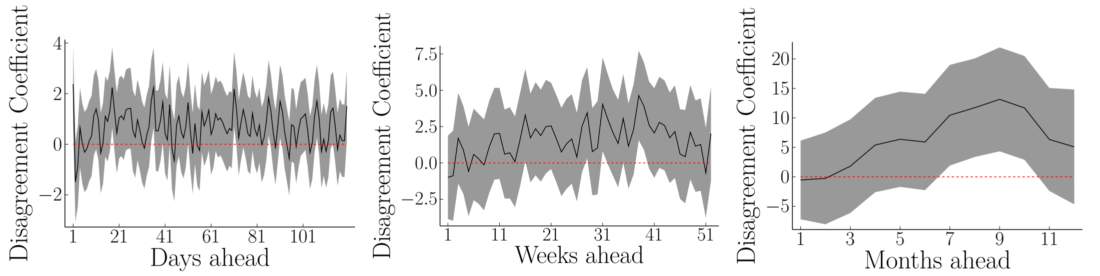

(a) Daily Turnover Growth. (b) Weekly Turnover Growth. (c) Monthly Turnover Growth.

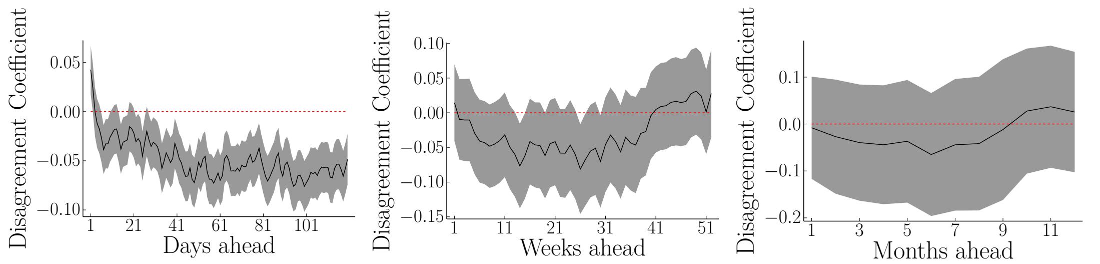

(d) Daily Volatility. (e) Weekly Volatility. (f) Monthly Volatility.

Figure 3.C.3: Univariate Robust Local Projection of Turnover Growth and Volatility on Disagreement.

Notes: In a univariate regression, we find that disagreement does not predict further turnover growth in the following periods. Moreover, the positive effect of disagreement on volatility is short-lived. The error bands are 90% confidence intervals according to Campbell and Yogo (2006).

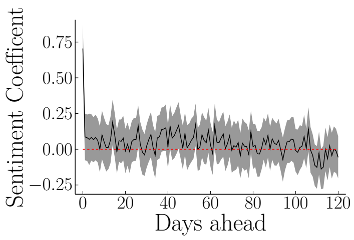

(a) Daily Returns.

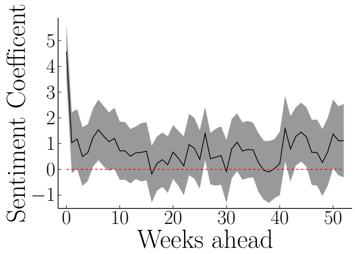

(b) Weekly Returns.

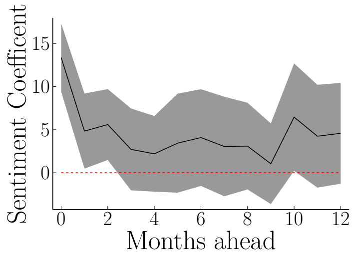

(c) Monthly Returns.

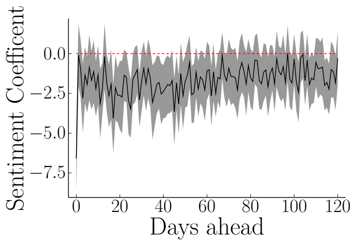

(d) Daily Turnover.

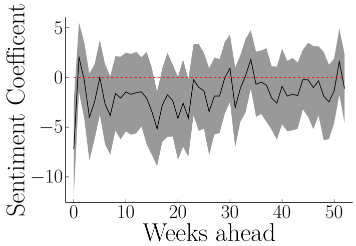

(e) Weekly Turnover.

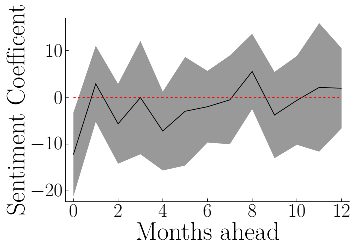

(f) Monthly Turnover.

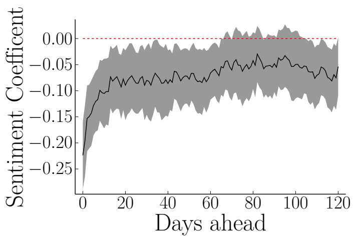

(g) Daily Volatility.

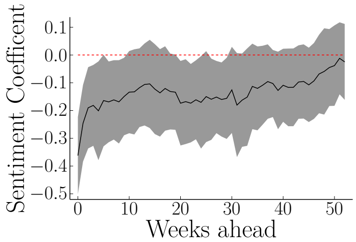

(h) Weekly Volatility.

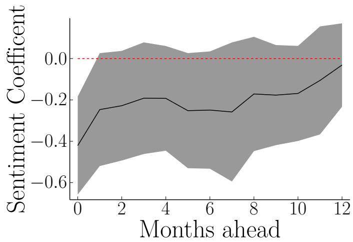

(i) Monthly Volatility.

Figure 3.C.4: Local Projections of Returns on Sentiment Controlling for Disagreement with HAC-Robust Standard Errors.

Notes: The shown estimates are the coefficients on sentiment when estimating (3.9) for leads of returns, turnover growth, and volatility. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

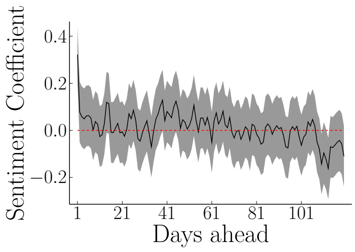

(a) Daily Returns.

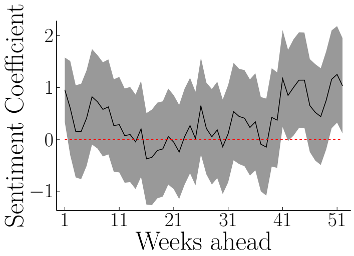

(b) Weekly Returns.

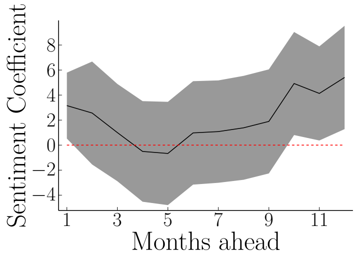

(c) Monthly Returns.

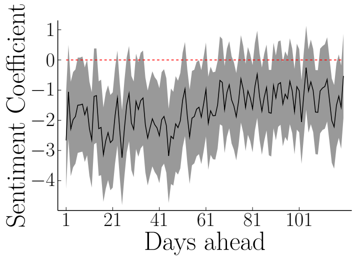

(d) Daily Turnover.

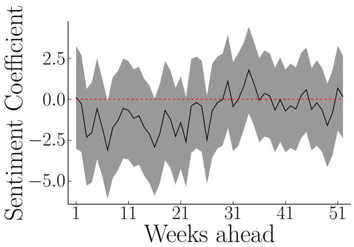

(e) Weekly Turnover.

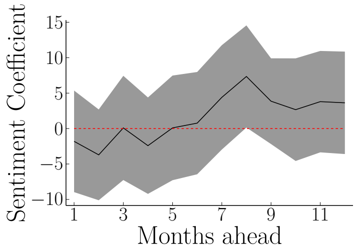

(f) Monthly Turnover.

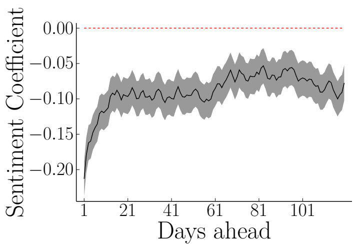

(g) Daily Volatility.

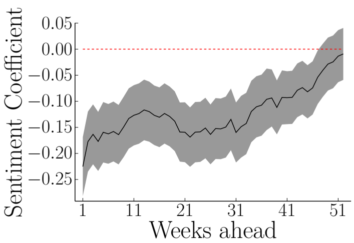

(h) Weekly Volatility.

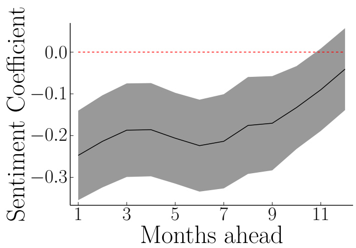

(i) Monthly Volatility.

Figure 3.C.5: Univariate Robust Local Projection of Returns, Turnover Growth, and Volatility on Sentiment.

Notes: In a univariate regression, we find that disagreement predicts negative returns for several periods at all frequencies. The error bands are 90% confidence intervals according to Campbell and Yogo (2006).

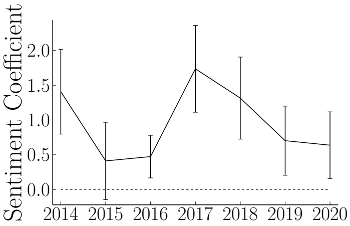

(a) Daily Returns.

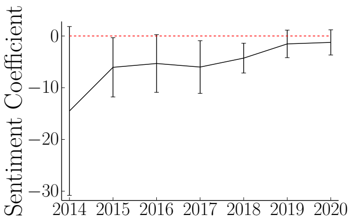

(b) Daily Turnover.

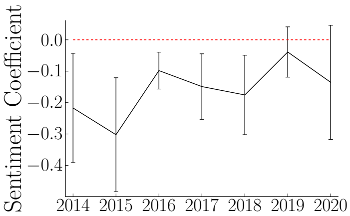

(c) Daily Volatility.

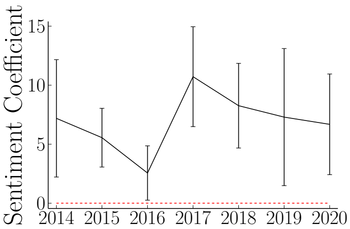

(d) Weekly Returns.

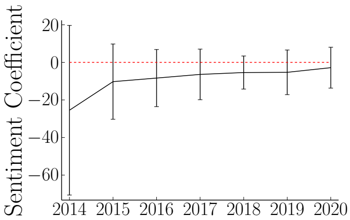

(e) Weekly Turnover.

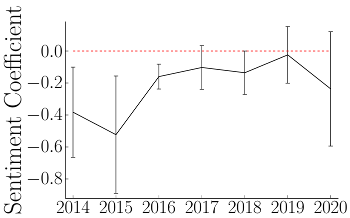

(f) Weekly Volatility.

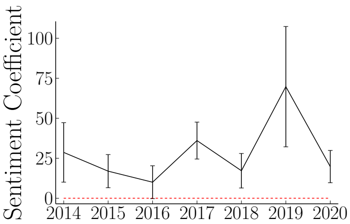

(g) Monthly Returns.

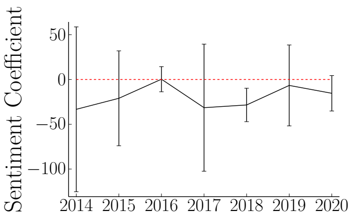

(h) Monthly Turnover.

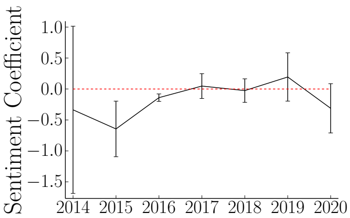

(i) Monthly Volatility.

Figure 3.C.6: Year-by-Year Coefficient of Sentiment for the Contemporaneous Regression.

Notes: We find that the contemporaneous effect of sentiment is relatively stable over time. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

3.C.1 Autocorrelation Plots

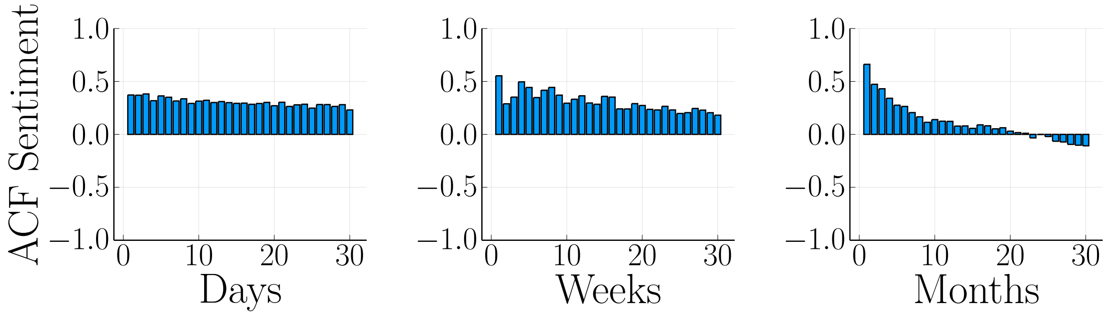

Figure 3.C.7: Autocorrelation Functions Sentiment.

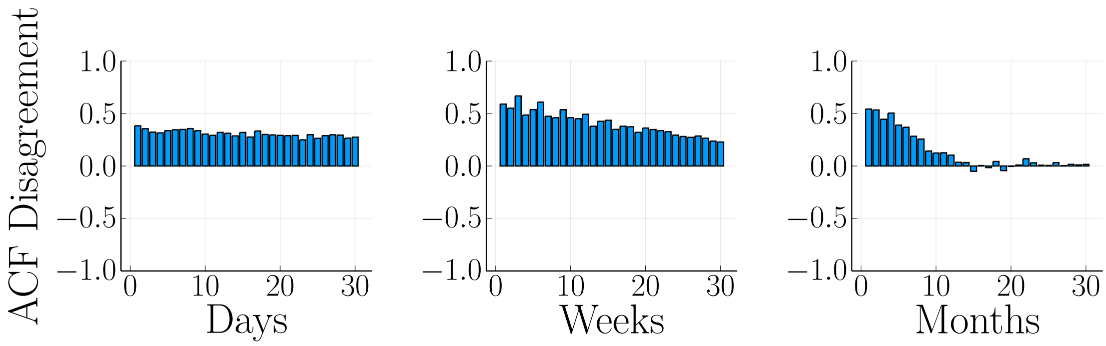

Figure 3.C.8: Autocorrelation Functions Disagreement.

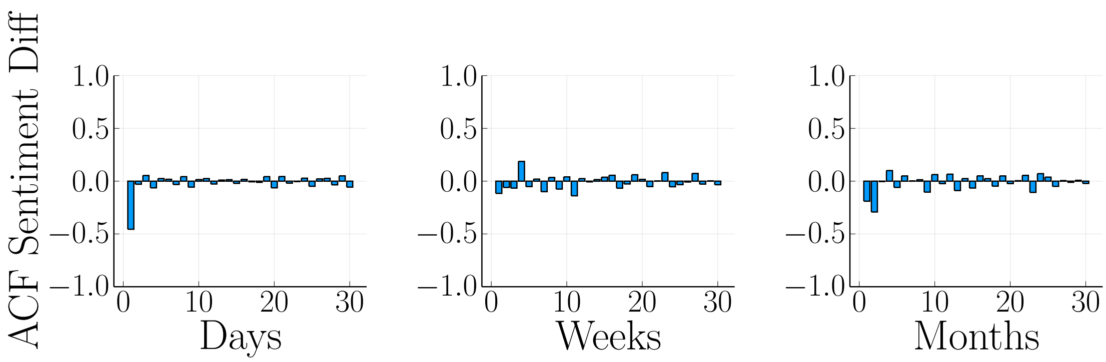

Figure 3.C.9: Autocorrelation Functions Changes in Sentiment.

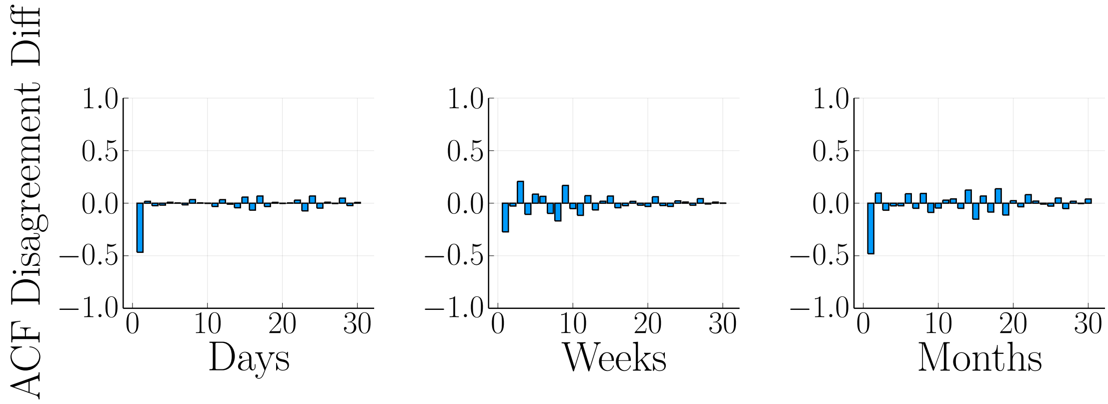

Figure 3.C.10: Autocorrelation Functions Changes in Disagreement.

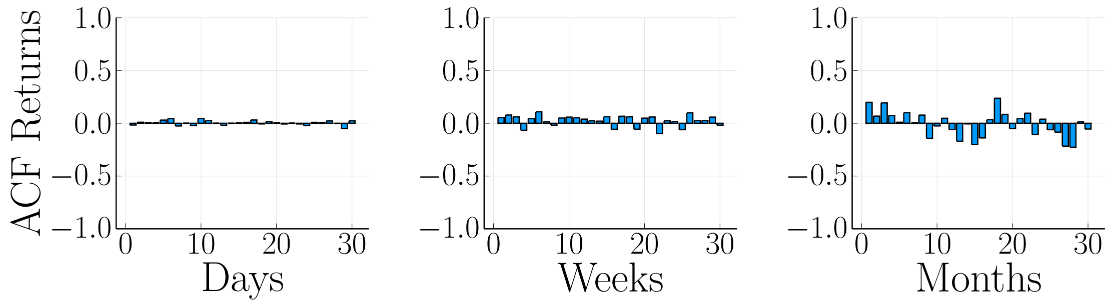

Figure 3.C.11: Autocorrelation Functions Returns.

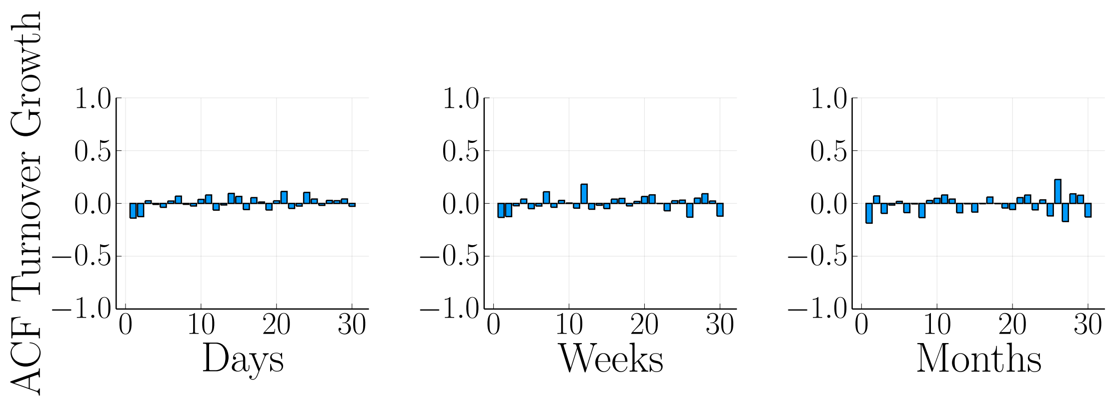

Figure 3.C.12: Autocorrelation Functions Turnover Growth.

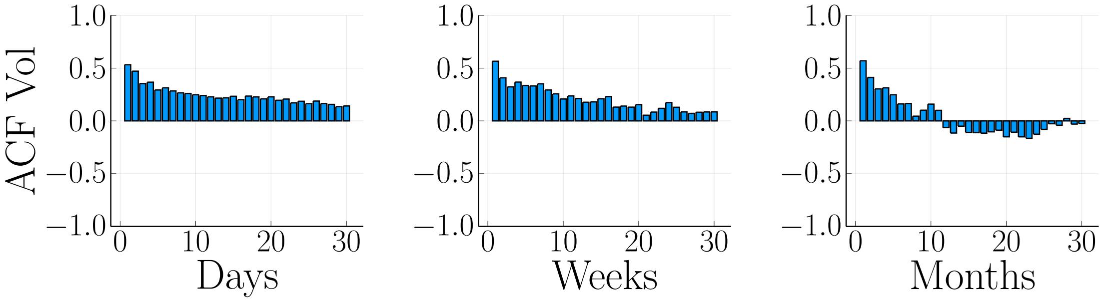

Figure 3.C.13: Autocorrelation Functions Volatility.

Bibliography

Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz (2016). "Fragility of asymptotic agreement under Bayesian learning". *Theoretical Economics* 11.1, pp. 187–225.

Acharya, Viral V, Tim Eisert, Christian Eufinger, and Christian Hirsch (2019). "Whatever It Takes: The Real Effects of Unconventional Monetary Policy". *The Review of Financial Studies* 32.9, pp. 3366–3411.

Albagli, Elias, Christian Hellwig, and Aleh Tsyvinski (2011a). *A Theory of Asset Prices based on Heterogeneous Information*. NBER Working Paper No. 17548. National Bureau of Economic Research.

— (2011b). *Information Aggregation, Investment, and Managerial Incentives*. NBER Working Paper No. 17330. National Bureau of Economic Research.

— (2015). *A Theory of Asset Prices based on Heterogeneous Information*. NBER Working Paper No. 17548, revised 2015. National Bureau of Economic Research.

— (2017). *Imperfect Financial Markets and Shareholder Incentives in Partial and General Equilibrium*. NBER Working Paper No. 23419. National Bureau of Economic Research.

— (2021). "Dispersed Information and Asset Prices". TSE Working Paper n. 21-1172.

Allen, Franklin and Douglas Gale (1994). "Limited Market Participation and Volatility of Asset Prices". *American Economic Review* 84.4, pp. 933–955.

Allen, W. David and Dorla A. Evans (2005). "Bidding and Overconfidence in Experimental Financial Markets". *Journal of Behavioral Finance* 6.3, pp. 108–120.

Ambrocio, Gene (2020). "Rational exuberance booms". *Review of Economic Dynamics* 35, pp. 263–282.

Angeletos, George-Marios, Guido Lorenzoni, and Alessandro Pavan (2010). *Beauty Contests and Irrational Exuberance: A Neoclassical Approach*. NBER Working Paper No. 15883. National Bureau of Economic Research.

Antweiler, Werner and Murray Z. Frank (2004). "Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards". *The Journal of Finance* 59.3, pp. 1259–1294.

Asea, Patrick K and Brock Blomberg (1998). "Lending cycles". *Journal of Econometrics* 83.1-2, pp. 89–128.

Asriyan, Vladimir (2021). "Balance Sheet Channel with Information-Trading Frictions in Secondary Markets". *The Review of Economic Studies* 88.1, pp. 44–90.

Asriyan, Vladimir, Dana Foarta, and Victoria Vanasco (2020). *The Good, the Bad and the Complex: Product Design with Imperfect Information*. CEPR Discussion Paper No. DP14307.

Asriyan, Vladimir, William Fuchs, and Brett Green (2019). "Liquidity sentiments". *American Economic Review* 109.11, pp. 3813–48.

Asriyan, Vladimir, Luc Laeven, and Alberto Martin (forthcoming). "Collateral Booms and Information Depletion". *The Review of Economic Studies*.

Bai, Jennie, Thomas Philippon, and Alexi Savov (2016). "Have financial markets become more informative?" *Journal of Financial Economics* 122.3, pp. 625–654.

Basu, Susanto, John G Fernald, and Miles S Kimball (2006). "Are Technology Improvements Contractionary?" *American Economic Review* 96.5, pp. 1418–1448.

Becker, Bo, Marieke Bos, and Kasper Roszbach (2020). "Bad Times, Good Credit". *Journal of Money, Credit and Banking* 52.S1, pp. 107–142.

Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey (2013). "Managerial Miscalibration". *The Quarterly Journal of Economics* 128.4, pp. 1547–1584.

Bennett, Benjamin, René Stulz, and Zexi Wang (2020). "Does the stock market make firms more productive?" *Journal of Financial Economics* 136.2, pp. 281–306.

Benos, Alexandros V. (1998). "Aggressiveness and survival of overconfident traders". *Journal of Financial Markets* 1.3, pp. 353–383.

Biais, Bruno, Denis Hilton, Karine Mazurier, and Sébastien Pouget (2005). "Judgemental Overconfidence, Self-Monitoring, and Trading Performance in an Experimental Financial Market". *The Review of Economic Studies* 72.2, pp. 287–312.

Blanchard, Olivier and Jordi Galí (2010). "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment". *American Economic Journal: Macroeconomics* 2.2, pp. 1–30.

Bond, Philip, Alex Edmans, and Itay Goldstein (2012). "The Real Effects of Financial Markets". *Annual Review of Financial Economics* 4, pp. 339–360.

Borio, Claudio, Enisse Karroubi, Christian Upper, and Fabrizio Zampoli (2015). *Labour reallocation and productivity dynamics: financial causes, real consequences*. BIS Working Papers No 534. Bank for Interantional Settlements.

Borio, Claudio and Anna Zabai (2018). "Unconventional monetary policies: a re-appraisal". *Research Handbook on Central Banking*. Edward Elgar Publishing.

Brandts, Jordi, Ayça Ebru Giritligil, and Roberto A. Weber (2015). "An experimental study of persuasion bias and social influence in networks". *European Economic Review* 80, pp. 214–229.

Brunnermeier, Markus K and Lasse Heje Pedersen (2009). "Market Liquidity and Funding Liquidity". *The Review of Financial Studies* 22.6, pp. 2201–2238.

Campbell, John Y. and Motohiro Yogo (2006). "Efficient tests of stock return predictability". *Journal of Financial Economics* 81.1, pp. 27–60.

Cecchetti, Stephen Giovanni, Hans Genberg, John Lipsky, and Sushil Wadhwani (2000). *Asset prices and central bank policy*. Centre for Economic Policy Research.

CME Group (2017). *CME Group Self-Certifies Bitcoin Futures to Launch Dec. 18 [Press Release]*. URL: https://www.cmegroup.com/media-room/press-releases/2017/12/01/cme_group_self-certifiesbitcoinfuturestolaunchdec18.html (visited on 06/25/2021).

— (2019). *CME Group Announces Jan. 13, 2020 Launch for Bitcoin Options [Press Release]*. URL: https://www.cmegroup.com/media-room/press-releases/2019/11/12/cme_group_announcesjan132020launchforbitcoinoptions.html (visited on 06/25/2021).

Chandrasekhar, Arun G, Horacio Larreguy, and Juan Pablo Xandri (2020). "Testing models of social learning on networks: Evidence from two experiments". *Econometrica* 88.1, pp. 1–32.

Chousakos, Kyriakos, Gary Gorton, and Guillermo Ordoñez (2020). *The Macroprudential Role of Stock Markets*. NBER Working Paper Series No. 27113. National Bureau of Economic Research.

Daniel, Kent and David Hirshleifer (2015). “Overconfident Investors, Predictable Returns, and Excessive Trading”. *Journal of Economic Perspectives*. Vol. 29. 4. American Economic Association, pp. 61–88.

Da Silva, Luiz Awazu Pereira and Phurichai Rungcharoenkitkul (2017). *QE experiences and some lessons for monetary policy: defending the important role central banks have played*. Speech for the for the Eurofi High Level Seminar 2017, Malta, 5-7 April 2017. URL: https://www.bis.org/speeches/sp170407.pdf (visited on 10/25/2020).

David, Joel M, Hugo A Hopenhayn, and Venky Venkateswaran (2016). “Information, Misallocation, and Aggregate Productivity”. *The Quarterly Journal of Economics* 131.2, pp. 943–1005.

Dávila, Eduardo and Cecilia Parlatore (2021). *Identifying Price Informativeness*. NBER Working Paper No. 25210. National Bureau of Economic Research.

De Long, J Bradford, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann (1990). “Noise Trader Risk in Financial Markets”. *Journal of Political Economy* 98.4, pp. 703–738.

Diamond, Douglas W. and Robert E. Verrecchia (1981). “Information aggregation in a noisy rational expectations economy”. *Journal of Financial Economics* 9.3, pp. 221–235.

Diether, Karl B, Christopher J Malloy, and Anna Scherbina (2002). “Differences of Opinion and the Cross Section of Stock Returns”. *The Journal of Finance* 57.5, pp. 2113–2141.

DNB (2017). *2016 Annual report*. Tech. rep. De Nederlandsche Bank.

Doerr, Sebastian (2018). “Collateral, Reallocation, and Aggregate Productivity: Evidence from the U.S. Housing Boom”. *Available at SSRN 3184162*.

Dow, James, Itay Goldstein, and Alexander Guembel (2017). “Incentives for Information Production in Markets where Prices Affect Real Investment”. *Journal of the European Economic Association* 15.4, pp. 877–909.

Dow, James and Gary Gorton (1997). “Stock Market Eefficiency and Economic Efficiency: is there a Connection?” *The Journal of Finance* 52.3, pp. 1087–1129.

— (2006). *Noise Traders*. NBER Working Paper No. 12256. National Bureau of Economic Research.

Durnev, Art, Randall Morck, and Bernard Yeung (2004). “Value-Enhancing Capital Budgeting and Firm-specific Stock Return Variation”. *The Journal of Finance* 59.1, pp. 65–105.

Enke, Benjamin and Florian Zimmermann (2019). “Correlation Neglect in Belief Formation”. *The Review of Economic Studies* 86.1, pp. 313–332.

Eyster, Erik and Georg Weizsäcker (2016). “Correlation Neglect in Portfolio Choice: Lab Evidence”. *Available at SSRN 2914526*.

Eyster, Erik, Georg Weizsäcker, Klaus M Schmidt, and Matthew Rabin (2018). *An Experiment On Social Mislearning*. Tech. rep. No. 73. CRC TRR 190 Rationality and Competition.

Fajgelbaum, Pablo D, Edouard Schaal, and Mathieu Taschereau-Dumouchel (2017). “Uncertainty traps”. *The Quarterly Journal of Economics* 132.4, pp. 1641–1692.

Farboodi, Maryam and Péter Kondor (2020). *Rational Sentiments and Economic Cycles*. NBER Working Paper Series No. 27472. National Bureau of Economic Research.

Farboodi, Maryam, Adrien Matray, Laura Veldkamp, and Venky Venkateswaran (2020). *Where Has All the Data Gone?* NBER Working Paper Series No. 26927. National Bureau of Economic Research.

Fernandez, Rodrigo, Pablo Bortz, and Nicolas Zeolla (2018). *The Politics of Quantitative Easing*. SOMO Amsterdam. The Centre for Research on Multinational Corporations.

Fostel, Ana and John Geanakoplos (2012). “Tranching, CDS, and asset prices: How financial innovation can cause bubbles and crashes”. *American Economic Journal: Macroeconomics* 4.1, pp. 190–225.

Gagnon, Joseph E and Brian Sack (2018). *QE: A User’s Guide*. Policy Briefs PB18-19. Peterson Institute for International Economics.

García-Santana, Manuel, Enrique Moral-Benito, Josep Pijoan-Mas, and Roberto Ramos (2020). “Growing Like Spain: 1995-2007”. *International Economic Review* 61.1, pp. 383–416.

Glaser, Markus and Martin Weber (2010). “Overconfidence”. *Behavioral Finance: Investors, Corporations, and Markets*. Vol. 6. John Wiley & Sons, pp. 241–258.

Gopinath, Gita, Şebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez (2017). “Capital Allocation and Productivity in South Europe”. *The Quarterly Journal of Economics* 132.4, pp. 1915–1967.

Gorton, Gary and Guillermo Ordoñez (2014). “Collateral Crises”. *American Economic Review* 104.2, pp. 343–378.

— (2020). “Good Booms, Bad Booms”. *Journal of the European Economic Association* 18.2, pp. 618–665.

Greenwood, Robin and Andrei Shleifer (2014). “Expectations of Returns and Expected Returns”. *The Review of Financial Studies* 27.3, pp. 714–746.

Grimm, Veronika and Friederike Mengel (2020). “Experiments on Belief Formation in Networks”. *Journal of the European Economic Association* 18.1, pp. 49–82.

Grossman, Sanford (1976). “On the Efficiency of Competitive Stock Markets Where Trades Have Diverse Information”. *The Journal of Finance* 31.2, pp. 573–585.

Grossman, Sanford and Joseph Stiglitz (1980). “On the impossibility of informationally efficient markets”. *American Economic Review* 70.3, pp. 393–408.

Harrison, J Michael and David M Kreps (1978). “Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations”. *The Quarterly Journal of Economics* 92.2, pp. 323–336.

Hayek, Friedrich August (1945). “The use of knowledge in society”. *American Economic Review* 35.4, pp. 519–530.

Hellwig, Martin F (1980). “On the Aggregation of Information in Competitive Markets”. *Journal of Economic Theory* 22.3, pp. 477–498.

Hirshleifer, David (2015). “Behavioral Finance”. *Annual Review of Financial Economics* 7, pp. 133–159.

Hirshleifer, David and Guo Ying Luo (2001). “On the survival of overconfident traders in a competitive securities market”. *Journal of Financial Markets* 4.1, pp. 73–84.

Hjalmarsson, Erik (2011). “New Methods for Inference in Long-Horizon Regressions”. *The Journal of Financial and Quantitative Analysis* 46.3, pp. 815–839.

Holmström, Bengt and Jean Tirole (1993). “Market Liquidity and Performance Monitoring”. *Journal of Political Economy* 101.4, pp. 678–709.

Hong, Harrison, Jose Scheinkman, and Wei Xiong (2006). “Asset Float and Speculative Bubbles”. *The Journal of Finance* 61.3, pp. 1073–1117.

Hong, Harrison and Jeremy C Stein (2003). “Differences of Opinion, Short-Sales Constraints, and Market Crashes”. *The Review of Financial Studies* 16.2, pp. 487–525.

— (2007). “Disagreement and the Stock Market”. *Journal of Economic Perspectives* 21.2, pp. 109–128.

Hutto, Clayton and Eric Gilbert (2014). “Vader: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text”. *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 8. 1.

Jegadeesh, Narasimhan and Di Wu (2013). “Word power: A new approach for content analysis”. *Journal of Financial Economics* 110.3, pp. 712–729.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor (2011). “Financial crises, credit booms, and external imbalances: 140 years of lessons”. *IMF Economic Review* 59.2, pp. 340–378.

Keys, Benjamin J, Tanmoy Mukherjee, Amit Seru, and Vikrant Vig (2010). “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans”. *The Quarterly Journal of Economics* 125.1, pp. 307–362.

Kindleberger, Charles P and Robert Z Aliber (2015). *Manias, panics and crashes: a history of financial crises*. Palgrave Macmillan.

Ko, K. Jeremy and Zhijian James Huang (2007). “Arrogance can be a virtue: Overconfidence, information acquisition, and market efficiency”. *Journal of Financial Economics* 84.2, pp. 529–560.

Krishnamurthy, Arvind and Tyler Muir (2017). *How Credit Cycles across a Financial Crisis*. NBER Working Paper No. 23850. National Bureau of Economic Research.

Kurtzman, Robert and David Zeke (2020). “Misallocation costs of digging deeper into the central bank toolkit”. *Review of Economic Dynamics* 38, pp. 94–126.

Kyle, Albert S (1985). “Continuous Auctions and Insider Trading”. *Econometrica* 53, pp. 1315–1335.

Kyle, Albert S. and F. Albert Wang (1997). “Speculation Duopoly with Agreement to Disagree: Can Overconfidence Survive the Market Test?” *The Journal of Finance* 52.5, pp. 2073–2090.

Lo, Andrew W. and A. Craig MacKinlay (1988). “Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test”. *The Review of Financial Studies* 1.1, pp. 41–66.

— (1990). “When Are Contrarian Profits Due to Stock Market Overreaction?” *The Review of Financial Studies* 3.2, pp. 175–205.

Loughran, Tim and Bill McDonald (2011). “When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks”. *The Journal of Finance* 66.1, pp. 35–65.

Malmendier, Ulrike and Geoffrey Tate (2005). "CEO Overconfidence and Corporate Investment". *The Journal of Finance* 60.6, pp. 2661–2700.

Martin, Alberto and Jaume Ventura (2018). "The Macroeconomics of Rational Bubbles: A User's Guide". *Annual Review of Economics* 10, pp. 505–539.

McQueen, Grant, Michael Pinegar, and Steven Thorley (1996). "Delayed Reaction to Good News and the Cross-Autocorrelation of Portfolio Returns". *The Journal of Finance* 51.3, pp. 889–919.

Mihet, Roxana (2020). *Financial Technology and the Inequality Gap*. Job Market Paper.

Miller, Edward M (1977). "Risk, Uncertainty, and Divergence of Opinion". *The Journal of Finance* 32.4, pp. 1151–1168.

Mishkin, Frederic S et al. (2011). "How Should Central Banks Respond to Asset-Price Bubbles? The 'Lean' versus 'Clean' Debate After the GFC". *RBA Bulletin* June, pp. 59–70.

Morck, Randall, Bernard Yeung, and Wayne Yu (2000). "The information content of stock markets: Why do emerging markets have synchronous stock price movements?" *Journal of Financial Economics* 58.1-2, pp. 215–260.

— (2013). "R 2 and the Economy ". *Annual Review of Financial Economics* 5.1, pp. 143–166.

Newey, Whitney K and Kenneth D West (1987). "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix". *Econometrica* 55.3, pp. 703–708.

Ningning, Pan and Zhu Hongquan (2014). "A Review of Stock Return Synchronicity". *2014 International Conference on Economic Management and Social Science (ICEMSS 2014)*. Atlantis Press, pp. 241–245.

Nowzohour, Laura and Livio Stracca (2020). "More than a Feeling: Confidence, Uncertainty, and Macroeconomic Fluctuations". *Journal of Economic Surveys* 34.4, pp. 691–726.

Okimoto, Tatsuyoshi (2019). *The Bank of Japan's exchange-traded fund purchases*.

Ordoñez, Guillermo (2013). "The Asymmetric Effects of Financial Frictions". *Journal of Political Economy* 121.5, pp. 844–895.

Peress, Joel (2014). "Learning from Stock Prices and Economic Growth". *Review of Financial Studies* 27.10, pp. 2998–3059.

Roll, Richard (1988). "R2". *The Journal of Finance* 43.3, pp. 541–566.

Sadka, Ronnie and Anna Scherbina (2007). "Analyst Disagreement, Mispricing, and Liquidity". *The Journal of Finance* 62.5, pp. 2367–2403.

Scheinkman, Jose A and Wei Xiong (2003). "Overconfidence and Speculative Bubbles". *Journal of Political Economy* 111.6, pp. 1183–1220.

Schmitt-Grohé, Stephanie and Martin Uribe (2012). "What's News in Business Cycles". *Econometrica* 80.6, pp. 2733–2764.

Shiller, Robert J (2015). *Irrational Exuberance*. Princeton University Press.

— (2017). "Narrative Economics". *American Economic Review* 107.4, pp. 967–1004.

Shleifer, Andrei and Robert W Vishny (1997). "The Limits of Arbitrage". *The Journal of Finance* 52.1, pp. 35–55.

Simsek, Alp (2013). "Belief Disagreements and Collateral Constraints". *Econometrica* 81.1, pp. 1–53.

— (2021). "The Macroeconomics of Financial Speculation". *Annual Review of Economics* 13.

Straub, Ludwig and Robert Ulbricht (2018). *Endogenous Uncertainty and Credit Crunches*. mimeo, Harvard and Boston College.

Tetlock, Paul C (2007). "Giving Content to Investor Sentiment: The Role of Media in the Stock Market". *The Journal of Finance* 62.3, pp. 1139–1168.

Tobin, James (1972). *The New Economics One Decade Older: The Eliot Janeway Lectureson Historical Economics in Honour of Joseph Schumpeter*. Princeton University Press.

Van Nieuwerburgh, Stijn and Laura Veldkamp (2006). "Learning asymmetries in real business cycles". *Journal of Monetary Economics* 53.4, pp. 753–772.

Veldkamp, Laura (2005). "Slow Boom, Sudden Crash". *Journal of Economic Theory* 124.2, pp. 230–257.

Vives, Xavier (2008). *Information and Learning in Markets: The Impact of Market Microstructure*. Princeton University Press.

— (2017). "Endogenous Public Information and Welfare in Market Games". *The Review of Economic Studies* 84.2, pp. 935–963.

Xiong, Wei (2013). *Bubbles, Crises, and Heterogeneous Beliefs*. NBER Working Paper No. 18905. National Bureau of Economic Research.