Essays in Economic Dynamics with Heterogeneous Expectations

Tommaso Di Francesco

Universiteit van Amsterdam
Università Ca' Foscari Venezia

Essays in Economic Dynamics with Heterogeneous Expectations

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. P.P.C.C. Verbeek
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op woensdag 4 juni 2025, te 10.00 uur

door Tommaso di Francesco
geboren te Atri

Promotiecommissie

Promotores:	prof. dr. C.H. Hommes	Universiteit van Amsterdam
	prof. dr. P. Pellizzari	Università Ca' Foscari Venezia
Overige leden:	prof. dr. J. Tuinstra	Universiteit van Amsterdam
	prof. dr. C.G.H. Diks	Universiteit van Amsterdam
	dr. K. Mavromatis	Universiteit van Amsterdam
	prof. dr. G. Iori	Università Ca' Foscari Venezia
	prof. dr. P.D.E. Dindo	Università Ca' Foscari Venezia
	prof. dr. R.C.J. Zwinkels	Vrije Universiteit Amsterdam

Faculteit Economie en Bedrijfskunde

Dit proefschrift is tot stand gekomen in het kader van het Marie Skłodowska-Curie Actions Innovative Training Network 'Economic Policy in Complex Environments' (EPOC) met als doel het behalen van een gezamenlijk doctoraat. Het proefschrift is voorbereid aan de Faculteit Economie en Bedrijfskunde van de Universiteit van Amsterdam en aan de Dipartimento di Economia van de Università Ca' Foscari Venezia.

This thesis was prepared within the framework of the Marie Skłodowska-Curie Actions Innovative Training Network 'Economic Policy in Complex Environments' (EPOC) with the purpose of obtaining a joint doctorate degree. The thesis was prepared in the Faculty of Economics and Business of the University of Amsterdam and in the Department of Economics of Ca' Foscari University of Venice.

Printed on non-ageing paper – ISO 9706

Acknowledgements

List of Authors

- Chapter 2 is a joint work with Cars Hommes. A previous version is available as EPOC Working Paper no 10, April 2023. Together, we elaborated the idea of the paper. I was responsible for the analytical and numerical results of the model and built the code. I collected the data and curated the empirical analysis. The writing and final editing was equally divided.

- Chapter 3 is a joint work with Daniel Torren Peraire. A previous version is available as EPOC Working Paper no 22, January 2025. Together, we elaborated the idea of the paper, drafted the model and built the code. Throughout the development of the paper, I was primarily responsible for the analytical results and the majority of the numerical results. I was responsible for the writing.

- Chapter 4 is single authored.

Contents

Acknowledgements **i**

List of Authors **iii**

Contents **v**

1 Introduction **1**

2 Sentiment-Driven Speculation in Financial Markets with Heterogeneous Beliefs: a Machine Learning approach **7**

2.1 Introduction 7
2.1.1 Related Literature 9
2.2 The model 10
2.2.1 Expectations 14
2.3 Numerical Simulations 15
2.3.1 Trend chasers vs pure bias 15
2.3.2 Numerical Simulations: b = 1.0, g = 1.3, R = 1.01 16
2.3.3 Introducing Speculation 18
2.3.4 Trend chasers vs pure bias vs fundamentalists 19
2.3.5 Numerical Simulations: b = 1.0, g = 1.3, R = 1.01 19
2.3.6 Trend chasers vs pure bias vs rational expectations 21
2.3.7 Numerical Simulations:b = 1.0, g = 1.3, R = 1.01 25
2.3.8 Speculation Effects 27
2.4 Taking the model to the data 28
2.4.1 Bitcoin Twitter Sentiment Index 28
2.4.2 Estimating Non-linear Rational Expectation models with a Neural Network 31
2.5 Estimation 34
2.6 Conclusion and Discussion 38

Appendix 2.A Proofs . . . 40
2.A.1 Proof of Lemma 1 . . . 40
2.A.2 Proof of Lemma 2 . . . 41
2.A.3 Proof of Lemma 3 . . . 44
2.A.4 Eigenvalues for the two type model . . . 47
2.A.5 Eigenvalues for the model with fundamentalist . . . 48
Appendix 2.B Steady States . . . 49
Appendix 2.C Neural Network . . . 50
Appendix 2.D Stationarity tests . . . 51
Appendix 2.E Robustness to different windows . . . 52
Appendix 2.F Bootstrapped standard errors . . . 53

3 (Mis)information diffusion and the financial market 55
3.1 Introduction . . . 55
3.1.1 Related literature and Contribution . . . 56
3.2 Model . . . 58
3.2.1 The financial market . . . 58
3.2.2 Actual and Perceived Law of Motion . . . 60
3.2.3 Information diffusion . . . 62
3.2.4 Payoffs and prices . . . 65
3.2.5 The role of social learning . . . 66
3.3 Empirical Calibration . . . 69
3.4 Numerical Simulations . . . 74
3.4.1 The effect of different Network Topologies . . . 76
3.5 Conclusion . . . 81
Appendix 3.A Forward Looking Price . . . 83
Appendix 3.B Relationship of updating to Kalman Filter . . . 84
Appendix 3.C Derivation of conditional variance . . . 85
Appendix 3.D Proofs . . . 85
3.D.1 Proof of Proposition 4 . . . 85
3.D.2 Proof of Proposition 5 . . . 86
Appendix 3.E Validation of the SNPE . . . 87

4 Sticky information across the wealth distribution 89
4.1 Introduction . . . 89
4.2 From uniform to wealth dependent stickiness . . . 92
4.3 Model . . . 96

4.3.1 Households . . . 97
4.3.2 Firms . . . 97
4.3.3 Labor market . . . 98
4.3.4 Government . . . 98
4.3.5 Monetary authority . . . 99
4.3.6 Equilibrium . . . 99
4.3.7 Beliefs . . . 99
4.3.8 Solution Method . . . 100
4.3.9 Calibration . . . 108
4.3.10 Monetary policy . . . 109
4.4 Estimation . . . 113
4.4.1 Model . . . 113
4.4.2 Empirical IRFs . . . 114
4.5 Conclusion . . . 117
Appendix 4.A Survey Data . . . 119
Appendix 4.B Method approximation . . . 119
Appendix 4.C Phillips Curve . . . 121
Appendix 4.D Empirical Impulse Responses . . . 122
Appendix 4.E Posterior Distribution . . . 123

5 Non technical summary **125**
5.1 Summary in English . . . 125
5.2 Samenvatting (Summary in Dutch) . . . 127

Bibliography **129**

Chapter 1

Introduction

Humans are among the few species able to form mental representations of the future. This ability allows us to simulate and plan accordingly to possible future scenarios.

Economics studies the interaction of economic agents. Economic agents, such as households, firms, and governments, inherit the forward-looking component from their human constituents. When these agents need to plan their savings, set prices, or implement policies, they do so by forming expectations about the consequences of their decisions. Expectation formation plays, therefore, a crucial role in economics.

Although in its inception, the most prominent theory was that of adaptive expectations[1], formalized by Cagan, Friedman, and Nerlove in the 1950s (Cagan, 1956; Nerlove, 1958), some scholars expressed their discontent with the permanent non-optimality of the agents' forecast. As Lucas puts it with adaptive expectations, *"price forecasts and actual prices will have different probability distributions, and this difference will be persistent, costly to forecasters, and readily correctible."*

This discontent was instrumental in paving the way for the theory that revolutionized the way expectations are modeled in economics and which is still the most widely used paradigm today, that of *rational expectations.*

The theory was brought forward in a seminal paper by John Muth (Muth, 1961). In his words: *"Expectations of firms (or, more generally, the subjective probability distribution of outcomes) tend to be distributed, for the same information set, about the prediction of the theory (or the 'objective' probability distributions of outcomes)."*

Rational expectations are, in other words, *model consistent.* To achieve such consistency, two assumptions are imposed on economic agents. The first one is that of *Full Information*; agents must have access to all relevant information to be included in

[1]This relies on the natural idea that individuals use experience from past events to form future projections. Some forms of adaptive expectations include purely naïve or some form of linear autoregressive model.

their forecasts. The second is that already mentioned, of rational expectations, that is, to be able to use the necessary information perfectly.[2] More recently, therefore, the Full Information part, implied in the earlier definition, was explicitly anteposed, and the theory is most known as Full Information Rational Expectations (FIRE). Lastly, although not explicitly mentioned, the theory applies to all agents so instead of considering individual entities, one can consider a representative agent for each group.

This assumption might seem unrealistic at the individual level, but two arguments were used to support it at the aggregate level. On the one hand, although agents might be heterogeneous in their expectations, the FIRE theory assumes that differences are due only to idiosyncratic noise. On the other hand, the learning argument argues that errors are only temporary and forecasts should converge to the FIRE model in the long run.

Along with the desirable property of subjective and objective expectations coinciding, two other factors contributed to making this theory so appealing to other scholars. First, rational expectations are a way of achieving *dimensionality reduction*: if agents' expectations coincide with mathematical expectations derived by the modeler, only the variables that the theory considers relevant are, indeed, relevant. Second, there is no need to explicitly concern oneself with expectations since the theory completely implies them. While Muth operated in a static environment, the FIRE framework was extended to dynamic equilibrium models by Lucas and Prescott (1971) and Lucas (1972), which paved the way for the future incorporation into Dynamic Stochastic General Equilibrium (DSGE). Models such as those of Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2007) rely on this assumption, and central banks' toolkits for policy evaluation are heavily built on them.

Following the financial crisis, however, policymakers have become increasingly dissatisfied with those toolkits and highlighted the necessity of operating with models with less restrictive assumptions and added realism. For example, Kenny and Morgan (2011) remarks that *"With regard to larger scale models, there is a need to continue developing more realistic models, especially as concerns the treatment of financial frictions and modelling of expectations formation."*

In response to this concern, the literature has started incorporating different expectation formation theories into macroeconomic models. These theories can be broadly divided into two categories.

The first one concerns those models that release the Full Information part of the

[2]In deterministic settings, this implies that agents have perfect foresight and are able to predict the future realization of a variable of interest perfectly. In stochastic environments, expectations are perfect in the sense that they are based on the correct probability distribution.

FIRE theory but keep rationality, like the model of noisy rational expectations Lucas (1972) and rational inattention (Maćkowiak and Wiederholt, 2009; Sims, 2003), and rely on methods like the Kalman Filter to derive the optimal forecast of agents conditional on their limited information set.

The second category focuses on the relaxation of perfect rationality. As Sargent (2008) puts it, *"There is an infinite number of ways to be wrong, but only one way to be correct."* Therefore, the literature has grounded alternative expectation formation processes in logically consistent theories or empirical and behavioral evidence. On one hand, the long-term argument that agents can eventually learn to behave according to FIRE has been challenged since, in general, there is no guarantee that the learning process will converge to a Rational Expectation Equilibrium (REE). Indeed, a prolific branch of this literature has focused on least-squared learning, when agents assume a variable follows a linear process and recursively use least squares to update the parameters of their model. In the case of linear models, extensive work has been done, and convergence conditions derived for the use of recursive least squares algorithms, (Bray, 1982; Bray and Savin, 1986; Marcet and Sargent, 1989*a*,*b*) or Bayesian learning (Townsend, 1978).

Evans and Honkapohja (2001) contributed to generalizing conditions for convergence, referred to as the "expectational stability." This concerns the map between the Perceived Law of Motion (PLM) that agents may have and the Actual Law of Motion that governs the economic model and defines a sufficient condition for the asymptotic stability of the REE. There are cases, however, in which the learning process, although not converging to the REE, could converge to an equilibrium that satisfies the consistency requirement highlighted by Lucas about a sensible expectation formation model. That is consistency between the agents' forecast and the outcome of a given variable.

Restricted Perceptions Equilibrium (RPE) (Branch, 2006; Evans and Honkapohja, 2001; Sargent, 1991) is an alternative to REE in which agents might have misspecified[3] beliefs but are unable to detect their misspecification because the orthogonality condition between forecast errors and the forecast model holds. A particular case of RPE was developed by Hommes and Sorger (1998), who introduced the idea of Consistent Expectations Equilibrium (CEE), in which the PLM is a misspecified autoregressive process and the ALM is a nonlinear chaotic process. In a CEE, however, agents cannot detect their misspecification because the mean and autocorrelation of the PLM and ALM coincide. Hommes, Sorger and Wagener (2013) then generalize this concept to

[3]In this context, a misspecified model is an under-parametrized model. For example, if two variables $y_{1,t}$ and $y_{2,t}$ are determined by a VAR(1) process, but the agent uses an AR(1) process for $y_{1,t}$, the agent is using a misspecified model.

stochastic (SCEE) environments, while Hommes and Zhu (2014) specialize it, limiting the PLM to be a first-order autoregressive process, to a so-called Behavioral Learning Equilibrium (BLE). As mentioned, these expectation formation theories share the same goal of the FIRE framework to build models that are theoretically consistent.

An alternative approach is measuring expectations directly by looking at empirical or experimental evidence and then building models to rationalize the observed behavior. Both REE and RPE impose strong assumptions on the agents' forecast, and in particular, that forecast errors be uncorrelated with an agent information set at the time in which the forecast was made. This assumption can be empirically tested and has been so by relying on survey data in, among the others, Manski (2004), Mankiw, Reis and Wolfers (2004), Armantier, Bruine de Bruin, Potter, Topa, van der Klaauw and Zafar (2013), Manski (2018), Weber, D'Acunto, Gorodnichenko and Coibion (2022), D'Acunto and Weber (2024). The common finding is that the FIRE assumption is not met in survey expectations. However, while the test rejects the null hypothesis, it remains silent on which alternative should be considered.

To close this gap, the literature has sought to explicitly qualify possible alternatives by relying on experiments where subjects are asked to form expectations in a financial or macroeconomic environment. Learning-to-Forecast (LtF) experiments can reconcile patterns such as herd behavior, and bubbles and bursts, observed in financial markets (Hommes, Sonnemans, Tuinstra and Velden, 2005) and offer a means of testing which forecasting models are most commonly employed by individual subjects. Two main features characterize the results. First, subjects have systematically heterogeneous[4] expectations. Second, their beliefs are based on parsimonious, backward-looking heuristics, such as naïve, adaptive expectations and trend-following rules. These results are in line with a theoretical body of work, building on the so-called Heuristic Switching Model (HSM) pioneered by Brock and Hommes (Brock and Hommes, 1997, 1998). In this setting, agents can form expectations with different heuristics and choose the best one to use according to some fitness measure. While in the adaptive learning literature, agents learn to optimize their forecasting model, in this context, they try to learn which forecasting model is the best.

Expectation formation models based on simple heuristics are also a core component of alternative approaches to economic modeling, such as Agent-Based Models (ABM)[5]. While the traditional macroeconomic approach is to work directly with aggregates, ABMs build economies from the ground-up by codifying individual agents' behavior and

[4]In the sense that unlike what Muth postulated, even after aggregating individual forecasts, the economy as a whole does not behave according to FIRE.

[5]See Dawid and Delli Gatti (2018) and Axtell and Farmer (Forthcoming) for extensive surveys.

then aggregating using simulations to offer a more realistic representation of complex economic interactions.

This thesis contributes to the renewed effort of building economic models with a more realistic treatment of agents' expectations by analyzing the effect of different alternatives to the FIRE assumption on financial and macroeconomic outputs.

Chapter 2 releases the Rational Expectation part of FIRE and builds on the Heuristic Switching Model. It studies a heterogenous asset pricing model in which different classes of investors coexist and evolve, switching among strategies over time according to a fitness measure. In the presence of boundedly rational agents, with biased forecasts and trend-following rules, the chapter studies the effect of two types of speculation: one based on fundamentalist and the other on rational expectations. While the first is only based on knowledge of the asset's underlying dynamics, the second also takes into account the behavior of other investors. The model is estimated on the Bitcoin Market with two contributions, relying on methods from Machine Learning. First, I construct the Bitcoin Twitter Sentiment Index to proxy a time-varying bias. Second, I propose a new method based on a Neural Network for the estimation of the resulting heterogeneous agent model with rational speculators. I show that the switching finds support in the data and that while fundamentalist speculation amplifies volatility, rational speculation has a stabilizing effect on the market.

Chapter 3 and 4 both work with models releasing the Full Information side of FIRE. Chapter 3 extends the noisy information model Lucas (1972) by endogenizing the information diffusion process, while Chapter 4 applies the Sticky Information framework (Carroll, Crawley, Slacalek, Tokuoka and White, 2020; Carroll, 2003; Mankiw and Reis, 2002) to a Heterogeneous Agent New Keynesian model (HANK).

Specifically, Chapter 3 investigates the interplay between information diffusion in social networks and its impact on financial markets using an ABM. Agents receive and exchange information about an observable stochastic component of the dividend process of a risky asset. A small proportion of the network has access to a private signal about the component, which can be either clean (information) or distorted (misinformation). Other agents are uninformed and can receive information only from their peers. All agents update their beliefs in a Bayesian manner, but they do so in a behavioral way, where they replace true precision with an individual parameter that depends on an endogenous and time-evolving measure of the agent's confidence in the source of the information. I examine, through simulations, how information diffuses in the network and provide a framework to account for the delayed absorption of shocks that are not immediately priced, as predicted by classical financial models. I show the

effect of network topology on the resulting asset price and offer an interpretation for excess volatility relative to fundamentals, persistence amplification, and leptokurtosis of returns.

Chapter 4 investigates the role of wealth-dependent information stickiness in the transmission of monetary policy in a Heterogeneous Agent New Keynesian (HANK) model. Using survey data, I provide empirical evidence that households do not form expectations according to the FIRE hypothesis but instead exhibit stickiness in updating their information, with wealthier households updating twice as much on a monthly basis. I evaluate the effect of this evidence on macroeconomic dynamics using a quantitative model. My findings reveal that inequality significantly affects the aggregate responses to monetary shocks. Specifically, models that neglect heterogeneity in information updating underestimate both the magnitude and the delay of the peak response to monetary policy shocks. Estimating the model by matching simulated Impulse Response Functions (IRFs) to empirical ones shows that stickiness is crucial for accurately capturing the dynamics observed in the data.

Each of the three chapters can be read independently and includes an individual introduction, conclusion, and appendix. Chapter 5 contains summaries in English and Dutch. References are contained in a common bibliography at the end of the thesis.

Chapter 2

Sentiment-Driven Speculation in Financial Markets with Heterogeneous Beliefs: a Machine Learning approach

2.1 Introduction

Expectations play a crucial role in economics. Aggregate variables depend on the interaction and decisions of multiple individuals. Since these decisions depend on how agents expect the future to evolve, the way in which expectations are formed affects the aggregate variables themselves. Financial markets provide a prime example of this *expectation feedback*, as the price of an asset includes expectations regarding its performance or its discounted cash flow in the future. Classical asset pricing models usually assume agents with rational expectations (RE) a là Muth (1961). The RE hypothesis assumes that agents form expectations which are model consistent and based on all available information. Moreover, classical models assume a representative agent. While idiosyncratic errors across agents are permitted, on average, the economy acts as if populated by a representative rational agent. If the RE hypothesis were true, then the price of any financial asset should be equal to its *fundamental value*, the expected present value of dividends. The literature has long observed that this is not true for many assets, but one cannot easily disentangle the role of expectations from that of the discount factor[6].

[6]See for example Adam and Nagel (2023).

In recent years, however, we have had the emergence of an asset class, *cryptocurrencies*, with no fundamental value[7]. In particular, Bitcoin, the first and most well-known cryptocurrency is the 10th largest asset for market capitalization with a value of 1.15 trillion dollars as of September 2024. Since Bitcoin does not and will never pay dividends, there is no discount factor that can explain a non-zero price. Among the possible explanations that can still maintain the RE hypothesis, two are the most prominent. The first one is that of Bitcoin as a *rational bubble* (Waters, 2019), for which all investors agree that the price will indefinitely grow at a rate higher than the risk-free rate. The second is that although there are no dividends associated with it, Bitcoin provides investors with some *convenience yield* (Hilliard and Ngo, 2022). Tax evasion and the purchasing of illegal goods are classical examples. Both explanations have drawbacks: the first because any transversality condition would make the rational bubble collapse as long as agents are rational; the second because it would imply a convenience yield almost as volatile as the Bitcoin price.

In this chapter, we propose a different explanation: *sentiment-driven speculation*. We relax the RE assumption in one dimension, namely, homogeneity, and allow for heterogeneous investors in the market. As Shalen (1993) notes, *"it is well understood that speculative trade in financial markets depends on divergent beliefs."* Papers like De Long, Shleifer, Summers and Waldmann (1990) introduce the concept of *noise traders*, a group of investors systematically acting on some signal uncorrelated with an asset's fundamental. In a similar spirit, Harris and Raviv (1993) and Hong and Stein (2003) study the effect on the market of differences of opinion among investors. A common theme of all these papers is that the different beliefs are exogenously determined. Brock and Hommes (1998) is an attempt to endogenize the divergence of beliefs by proposing a Heuristic Switching Model (HSM). Building on this work, our Bitcoin model considers a market with biased investors and trend chasers whose fractions evolve endogenously according to their realized profits. Biased investors believe that the fundamental value of Bitcoin is some positive value. Since this belief is not dictated by any financial reason, we also refer to them as *sentiment followers*. Trend chasers add a short-term demand, driven by momentum. Given the presence of non-rational behaviors among investors, one would expect speculation to be possible in the market. Therefore, we study two types of speculation: one associated with fundamentalists and one with rational expectations. It is important to note that the two do not coincide because of the presence of heterogeneous agents. Perfect rationality requires knowledge not only of the asset's fundamental value but also of the strategies adopted by other individuals.

[7]Some other examples of assets without associated cash flow include gold, silver or art.

It coincides with *perfect foresight* in a deterministic setting. To solve for RE, we rely on the Extended Path (EP) method of Fair and Taylor (1983). We then bring the model to the data, introducing two contributions. First, we propose to capture the bias of agents in the market by constructing an index based on textual data from Twitter. Using a dataset of more than ten million tweets containing the word "Bitcoin", we construct the Bitcoin Twitter Sentiment Index (BiTSI) through sentiment analysis in the form of the Valence Aware Dictionary and sEntiment Reasoner (VADER). Second, while there are multiple estimations of HSM in the literature, none of them, to the best of our knowledge, include fully rational investors. One of the reasons is that in order to estimate the model according to the Fair and Taylor (1983) method, one has to first solve for RE using the EP for each combination of the parameters of interest. As the parameter space grows exponentially with every additional parameter, this method suffers from the *curse of dimensionality*. We propose to tackle this issue with a Machine Learning model. Specifically, we show that it is possible to approximate RE with a Neural Network (NN). In practice, we re-parameterize the problem such that we only need to estimate the rational expectations vector once. After obtaining the vector, the model can be estimated using normal non-linear least squares.

The rest of the chapter is structured as follows: Section 2.1.1 discusses some related literature, Section 2.2 introduces the HSM and Section 2.3 explores its global dynamics with numerical simulations, Section 2.4 introduces the BiTSI index, Section 2.4.2 introduces our method to estimate the non-linear RE model and shows its effectiveness on simulated data, Section 2.5 estimates the models on real data, Section 2.6 concludes.

2.1.1 Related Literature

On the theoretical side, this chapter relates to the literature on HSM. Brock and Hommes (1997) focuses on a cobweb model. Since the agents need to make a one-period-ahead prediction, they can solve for the heterogeneous agents price dynamics analytically in the presence of rational agents. In financial markets with a temporary equilibrium setup, the prediction is a two-step ahead one, which makes the analysis of the model with RE non trivial. Also for this reason, in Brock and Hommes (1998), after analyzing the steady states for the model with perfect foresight, the authors replace this heuristic with fundamentalist traders. The interplay among fundamentalist, chartists and biased agents has been studied extensively both in discrete (Hommes, Huang and Wang, 2005) and continuous (He and Li, 2012) time and the reader can refer to Hommes (2021) for an extensive survey of the literature. An exception is Boehl and Hommes (2025), in which the authors investigate the effect of perfect foresight in a theoretical

HSM. We use the Fair and Taylor (1983) EP algorithm to obtain the perfect foresight solution and provide a method to estimate the resulting non-linear heterogeneous agent rational model. Moreover while in our case the non-linearity is given by heterogeneity, our method is general and can be applied to any type of non-linear model with RE. This is important since while some papers like Fisher, Holly and Hallett (1986) and Armstrong, Black, Laxton and Rose (1998) have improved the original Fair and Taylor (1983) method in terms of robustness and efficiency, their focus has been only on solving for RE and not on estimation.

This method allows us to contribute to the literature on the empirical validity of HSM by estimating a model with the RE heuristic for the first time. In this area, most papers relied on Non-Linear Least Squares to estimate a model with mean-reverter (fundamentalist) and chartists (trend-followers) on the S&P 500 (Boswijk, Hommes and Manzan, 2007; Chiarella, ter Ellen, He and Wu, 2014; Hommes and in 't Veld, 2017), the option market (Frijns, Lehnert and Zwinkels, 2010), the gold market (Baur and Glover, 2014), the European Credit Default Swap market (Chiarella, He and Zwinkels, 2014), the housing market (Bolt, Demertzis, Diks, Hommes and van der Leij, 2019), US inflation dynamics (Cornea-Madeira, Hommes and Massaro, 2019) or a comparison across asset classes (ter Ellen, Hommes and Zwinkels, 2021). Lof (2015) also focuses on the S&P 500 but focuses on a market with short-term and long-term fundamentalists. Finally some papers like Franke and Westerhoff (2012) and Schmitt (2021) rely on the Simulated Method of Moments to estimate a model with fundamentalists and chartists.

Another difference with this literature is that we estimate a model that includes biased investors. For this, we bridge the HSM estimation literature with research focusing on the role of sentiment in the Bitcoin market, which we use as a proxy for time-varying bias. Theoretically, we build on research that has focused on bias or sentiment in the form of animal spirits (De Grauwe and Rovira Kaltwasser, 2012; Gardini, Radi, Schmitt, Sushko and Westerhoff, 2022, 2025) and has shown that the presence of sentiment traders hampers the stability of stock markets. This is particularly relevant for the Bitcoin market, which exhibits a lesser degree of efficiency compared to other financial markets and is more prone to speculative bubbles and other market inefficiencies. These inefficiencies can potentially lead to substantial investor losses and present numerous arbitrage opportunities, as evidenced in studies like Cheah and Fry (2015), Urquhart (2016), and Makarov and Schoar (2020). In this setting, multiple papers have focused on the effect of sentiment measured on the website StockTwits (Chen and Hafner, 2019; Guégan and Renault, 2021), on Google Trends (Aalborg, Molnár and de Vries, 2019; Baig, Blau and Sabah, 2019; Liu and Tsyvinski, 2020; Urquhart,

2018), X (J.Parra-Moyano, D.Partida, Gessl and S.Mazumdar, 2024), the Cryptocurrency Forum Sentiment (Gurdgiev and O'Loughlin, 2020) or the Fear and Greed Index (Bourghelle, Jawadi and Rozin, 2022). However, while this literature is empirical and relies on a reduced form estimation, we estimate a structural model that offers a clear mechanism for how sentiment and switching among different strategies can explain the highly volatile Bitcoin market environment.

2.2 The model

We begin by laying down the theoretical framework of the Brock and Hommes (1998) model. Consider an overlapping generation model. The economy is populated by N agents living two periods. When young agents receive w_0 units of consumption good. When old, they consume all of their wealth, with a utility function given by

$$U(c_{t+1}) = -e^{-ac_{t+1}},$$

where $a > 0$ is the coefficient of absolute risk aversion. Agents can choose between two types of securities to transfer wealth from the first period to the second. They can use a riskless asset, which pays fixed interest $R > 1$ for each unit of saved goods, or alternatively, they can use a risky asset. Agents pay price p_t to purchase the asset at time t, and when old they obtain the payoff $y_{t+1} = p_{t+1} + d_{t+1}$. This is given by the price at which they can sell the asset p_{t+1} plus a dividend claim d_{t+1} which is assumed to be constant plus normal white noise. In our specific setting we will assume it constantly equal to 0 so that going forward $y_{t+1} = p_{t+1}$. Agents need to chose their demand of the risky asset, defined by z_t^d, in order to maximize their utility, subject to the following budget constraint

$$c_{t+1} = R(w_0 - p_t z_t^d) + z_t^d p_{t+1}.$$

All agents assume that c_{t+1} is normally distributed, then their utility maximization problem is equivalent to

$$\max_{\{z_t^d\}} \left(-exp\left\{ -a\mathbb{E}_t(c_{t+1}) + \frac{a^2}{2}\mathbb{V}_t(c_{t+1}) \right\} \right). \tag{2.2.1}$$

Exploiting the budget constraint we have that

$$\mathbb{E}_t(c_{t+1}) = R(w_0 - p_t z_t^d) + z_t^d \mathbb{E}_t(p_{t+1}),$$

and

$$\mathbb{V}_t(c_{t+1}) = (z_t^d)^2 \mathbb{V}_t(p_{t+1}).$$

Then, the optimal choice of the risky asset must satisfy the first order condition of equation (2.2.1).

$$z_t^d = \frac{\mathbb{E}_t(p_{t+1}) - Rp_t}{a\mathbb{V}_t(p_{t+1})}.$$

To get a solution for the price p_t, we impose the equilibrium condition that demand must equal supply z_t^s. Finally we assume net supply of the risky asset is 0.[8]

$$\frac{\mathbb{E}_t(p_{t+1}) - Rp_t}{a\mathbb{V}_t(p_{t+1})} = 0. \tag{2.2.2}$$

From equation (2.2.2) we can note that if agents had homogenous beliefs, therefore sharing the same expected value and variance of the asset, the pricing equation would be

$$Rp_t = \mathbb{E}_t(p_{t+1}). \tag{2.2.3}$$

Equation (2.2.3) has two class of solutions, usually referred to as fundamentalist and rational bubble solution. The fundamental price is $p_t \equiv \bar{p} = 0$ in each time step, where $r = R - 1$ is the net free-rate. Whereas the rational bubble price is $\tilde{p_t} = \zeta_t R^{-t}$ where ζ_t is any martingale process. Now consider J different types of agents, with different expectation formation processes regarding the future price of the risky asset. The equilibrium pricing equation (2.2.2) becomes

$$\sum_{j=1}^{J} \left(n_{j,t} \cdot \frac{\mathbb{F}_{j,t}(p_{t+1}) - Rp_t}{a\mathbb{V}_{j,t}(p_{t+1})} \right) = 0,$$

where $\mathbb{F}_{j,t}(p_{t+1})$ and $\mathbb{V}_{j,t}(p_{t+1})$ indicate the subjective expectation and variance of agent j and $n_{j,t}$ represents the fraction of agents using strategy j. Finally we make the assumption of constant beliefs on variance: $\mathbb{V}_{j,t}(p_{t+1}) = \sigma_p^2, \forall j$. [9] With this

[8] One could argue that the external net supply of Bitcoin is not 0, due to the mining activity. We justify this assumption in two ways. First, in the long run, the total amount of Bitcoins is fixed to 21 million. It is, therefore, theoretically relevant to explore dynamics under zero supply. Second, Bitcoin undergoes a process called halving, for which after 210000 blocks are added to the Block Chain, the reward that miners receive is halved. One could model the supply between two halvings as constant and obtain the following price equation $Rp_t = \mathbb{E}_t(p_{t+1}) - a\mathbb{V}_t(p_{t+1})z_t^s$. Notice that this implies a negative fundamental value for the asset. Mathematically, we could work with a model in deviation from this fundamental value and obtain similar dynamics to the one presented, just translated by a constant term.

[9] Although we shall overlook this second-order effect, it should be noted that heterogeneity in

assumption, the equilibrium price is given by

$$Rp_t = \sum_{j=1}^{J} n_{j,t} \cdot \mathbb{F}_{j,t}(p_{t+1}). \tag{2.2.4}$$

where $\mathbb{F}_{j,t}(p_{t+1})$ is the subjective forecast implied by belief j and $R = 1 + r$ is the gross risk-free rate. Timing is important, and we assume that the current equilibrium price p_t is not realized and, therefore, not available to the agents when forming their beliefs. In other words, they are making a two-period-ahead forecast. Equation (2.2.4) simply states that the current price is the discounted weighted average of the different J beliefs, weighted by the fraction of the population $n_{j,t}$ that embraces the belief at time t.

Fractions are updated every period according to a fitness measure that is public knowledge and is given by profits or returns in excess of the risk-free rate and given by

$$\pi_{j,t} = (p_t - Rp_{t-1})\, z_{t-1}^{d},$$

finally assuming that $a\sigma_p^2 = 1$ we get

$$\pi_{j,t} = (p_t - Rp_{t-1})\left(\mathbb{F}_{j,t-1}(p_t) - Rp_{t-1}\right). \tag{2.2.5}$$

The fraction of agents choosing strategy j is given by the multinomial logit model

$$n_{j,t} = \frac{e^{\beta \pi_{j,t-1}}}{\sum_{j=1}^{J} e^{\beta \pi_{j,t-1}}}. \tag{2.2.6}$$

The parameter β is the *intensity of choice*. When $\beta = 0$, agents ignore the fitness measure and distribute evenly across the set of predictors. At the aggregate level, fractions are constant at $1/J$. If the intensity of choice is infinite, all agents switch to the most profitable strategy, and all but the fraction associated with the best strategy are zero. Equations (2.2.4) and (2.2.6) jointly determine the price and the evolution of fractions. It is important to note that realized profits and forecast accuracy are not perfectly proportional. When an individual forecasts exactly the future price, their profits are guaranteed to be positive and given by $(p_t - Rp_{t-1})^2$. However, for an individual with an incorrect forecast to earn more than this quantity it is sufficient to be inaccurate in the "right direction". Consider the quantity $(p_t - Rp_{t-1})\left(\mathbb{F}_{j,t-1}(p_t) - p_t\right)$ which represents the numerator of the difference in realized profits obtained by a perfectly

conditional expectations does, in fact, cause heterogeneity in conditional variance. Again, we refer to Brock and Hommes (1998) for a more thorough discussion.

accurate individual and a generic individual employing strategy j. For this quantity to be positive, thus consisting of more profits for the agent using the "incorrect forecast", it is sufficient that

$$\operatorname{sgn}\left(p_t - Rp_{t-1}\right) \cdot \operatorname{sgn}\left(\mathbb{F}_{j,t-1}(p_t) - p_t\right) = 1,$$

with $\operatorname{sgn}(x) = x/|x|$ being the sign function. The intuition is that the forecasting error is in the same direction as the price change. The reason why even with a perfect forecast an individual does not purchase or sell unlimited quantities of the risky asset is the bound imposed by their risk aversion and the variance of the risky asset. Incorrect investors are overly optimistic or pessimistic, therefore purchasing or selling more than they should. Sometimes, by chance, their overconfidence pays off.

2.2.1 Expectations

We assume that all forecasting strategies are given by a linear function of an observable i.i.d stochastic process w_t and of past and future prices[10] .

$$\mathbb{F}_{j,t}(p_{t+1}) = h_{j,t}(\{p_{t-l}\}_{l=1}^{L}, \{\mathbb{E}_t(p_{t+j})\}_{j=0}^{K}, \{w_{t-l}\}_{l=1}^{L}).$$

The functional form allows for simple but typical beliefs supported by experimental evidence as in Hommes, Sonnemans, Tuinstra and Velden (2005), or for more sophisticated ones. The formulation so far is general enough to accommodate a multiplicity of strategies as in Brock, Hommes and Wagener (2005), but we will focus on four strategies that are representative of behaviors we expect to see in the market and are summarized below.

Trend chasers. Sometimes also referred to as chartists. They form their forecast by an analysis of past prices. We consider the simplest form given by

$$h_{j,t}(\{p_{t-l}\}_{l=1}^{L}, \{\mathbb{E}_t(p_{t+j})\}_{j=0}^{K}, \{w_{t-l}\}_{l=1}^{L}) = gp_{t-1}, \quad g > 0. \tag{2.2.7}$$

In forming their expectations, they extrapolate from last period price deviation. The case $0 < g < 1$ corresponds to *mean reverting* beliefs[11]. When $g \geq 1$ we have

[10]We remark that expectations are formed at the beginning of the period, while the price realizes at the end. In this setting the notation $\mathbb{E}_t(p_t)$ stresses that p_t is not yet realized when agents are forming their expectations at time t.

[11]The literature, especially in empirical works, has sometimes referred to this case as fundamentalist, since agents believe in long run mean reversion to the fundamental value. In our case we include them in this category, for mathematical consistency, as their behavior is also regulated by the single

trend chasing beliefs, with the case $g > R = 1 + r$ usually referred to as *strong trend chasing*. The presence of this category of investors in the system is consistent with evidence provided by the literature about the existence of a momentum factor in stock markets, and originated by the seminal paper of Jegadeesh and Titman (1993).

Pure bias. This class of investors base their beliefs on some process which in the empirical part we will think of as *sentiment*, unrelated to the asset's fundamental value or price. Their forecast is

$$h_{j,t}(\{p_{t-l}\}_{l=1}^{L}, \{\mathbb{E}_t(p_{t+j})\}_{j=0}^{K}, \{w_{t-l}\}_{l=1}^{L}) = bw_{t-1}, \quad b > 0. \tag{2.2.8}$$

Fundamentalists. Investors of this type base their expectations on the fundamental value of the asset. They would be rational in a homogeneous rational world, in the sense that if the market were populated only by fundamentalists, then their forecast would be the correct one. However, they fail to consider the presence of boundedly rational agents with different non-rational beliefs in the market. Recalling that the fundamental value in this setting is equal to zero, we have

$$h_{j,t}(\{p_{t-l}\}_{l=1}^{L}, \{\mathbb{E}_t(p_{t+j})\}_{j=0}^{K}, \{w_{t-l}\}_{l=1}^{L}) = 0. \tag{2.2.9}$$

Rational expectations. Investors of this type are the most sophisticated, possessing knowledge of the underlying model regulating asset price evolution and the composition of the market. Their forecast is *model consistent*:

$$h_{j,t}(\{p_{t-l}\}_{l=1}^{L}, \{\mathbb{E}_t(p_{t+j})\}_{j=0}^{K}, \{w_{t-l}\}_{l=1}^{L}) = \mathbb{E}_t(p_{t+1}). \tag{2.2.10}$$

2.3 Numerical Simulations

We now investigate the dynamics that can emerge from the interaction of these simple strategies. The first case we consider is a two-type market, which is going to serve as the baseline case. We will then analyze the effect of the introduction of fundamentalists and rational speculators in the market.

parameter g.

2.3.1 Trend chasers vs pure bias

We take the first strategy to be trend chasing as in (2.2.7) and the second one to be pure bias (2.2.8). We start by considering deterministic simulations and assume that w_{t-1} is constantly equal to its mean value, which we fix to 1 without loss of generality.

The full model is therefore given by the following two equations:

$$Rp_t = n_{1,t} g p_{t-1} + (1 - n_{1,t}) b, \tag{2.3.1}$$

$$n_{1,t} = \{1 + exp(\beta(p_{t-1} - Rp_{t-2})(b - gp_{t-3}))\}^{-1}. \tag{2.3.2}$$

We provide an analytical characterization of the steady states of the system for extreme values of the intensity of choice parameter β.

Lemma 1 (Steady States for the two type model)**.** *For $\beta = 0$ the model has a unique positive steady state $p = \frac{b}{2R-g}$ which is locally stable for $g < 2R$. For $\beta \to +\infty$, there are the following possibilities. If $g = R$ then any $0 < p \leq \frac{b}{R}$ is a steady state with all steady states $0 < p < \frac{b}{R}$ being locally unstable and $p = \frac{b}{R}$ being locally stable. If $g > R$, there exists a unique steady state $p = \frac{b}{R}$ which is locally stable. If $g < R$, there are no (positive) steady states.*

Proof. See Appendix 2.A.

The lemma shows that even in the neoclassical limit in which agents immediately switch to the best performing strategy, pure bias agents are not pushed out of the market. Indeed the most interesting case is the one in which $g > R$ and trend chasing is *strong*. In this situation the steady state is a positive deviation from the fundamental value, the magnitude of which depends on b. When β is strictly positive, the steady state equation can only be derived implicitly. We do this in Appendix 2.B, while below we use numerical simulations to highlight the global dynamics that this system may generate.

2.3.2 Numerical Simulations: b = 1.0, g = 1.3, R = 1.01

We fix all parameters except for the intensity of choice β and then study the global dynamics of the model. First, we report a bifurcation diagram for increasing values of the intensity of choice parameter. For each value of β, we simulate 10000 iterations of the system and then visualize the last 2000 price realizations to eliminate the effect of initial conditions. In the top left panel of Figure 2.1, we can observe that the system has a stable steady state for values of the intensity of choice lower than approximately 6.725.

After the parameter crosses this value, a bifurcation occurs. In order to characterize this bifurcation and the system dynamics after bifurcation, we use maximum Lyapunov exponents[12] and a plot of the modulus of the system's eigenvalues. For a n-dimensional system evolving according to $x_{t+1} = F(x_t)$, the maximum Lyapunov exponent is defined as

$$\lambda_1 = \lim_{t\to\infty} \frac{1}{t} \log(\|(D_x F^t)\delta x\|), \tag{2.3.3}$$

where $D_x F^t$ is the Jacobian of the t-iterate of the map F evaluated at the point x and δx is an initial perturbation. Now for a n-dimensional space, there exist a Lyapunov spectrum of exponents ($\lambda_1 \geq \lambda_2 \geq ... \geq \lambda_n$), measuring the expansion or contraction in the n different directions of the phase space. However, a study of the maximum Lyapunov exponent is sufficient to analyze the behavior of the system since we have the following cases:

- $\lambda_1 > 0$: the system is chaotic;
- $\lambda_1 = 0$: the system has a quasi-periodic attractor;
- $\lambda_1 < 0$: the system has a stable steady state or a stable cycle.

As we can see in the top right panel of Figure 2.1 the exponent is negative for low values of β and becomes equal to 0 after the bifurcation. Therefore, the system exhibits stable and quasi-periodic orbits, never producing chaos. In the bottom left panel, we plot the modulus of the eigenvalue of the system, that we compute in Section (2.A.4) of the appendix. We can see that for values of $\beta \approx 6.725$, the two complex eigenvalues cross the unit circle. We conclude, therefore, that a Hopf bifurcation occurs. Finally in the bottom right panel, we show the creation of stable invariant circles after the bifurcation, allowing us to classify the bifurcation as super-critical.

[12]The computation is an in-built function of the Julia package DynamicalSystems.jl, that relies on the method by Benettin, Galgani, Giorgilli and Strelcyn (1980).

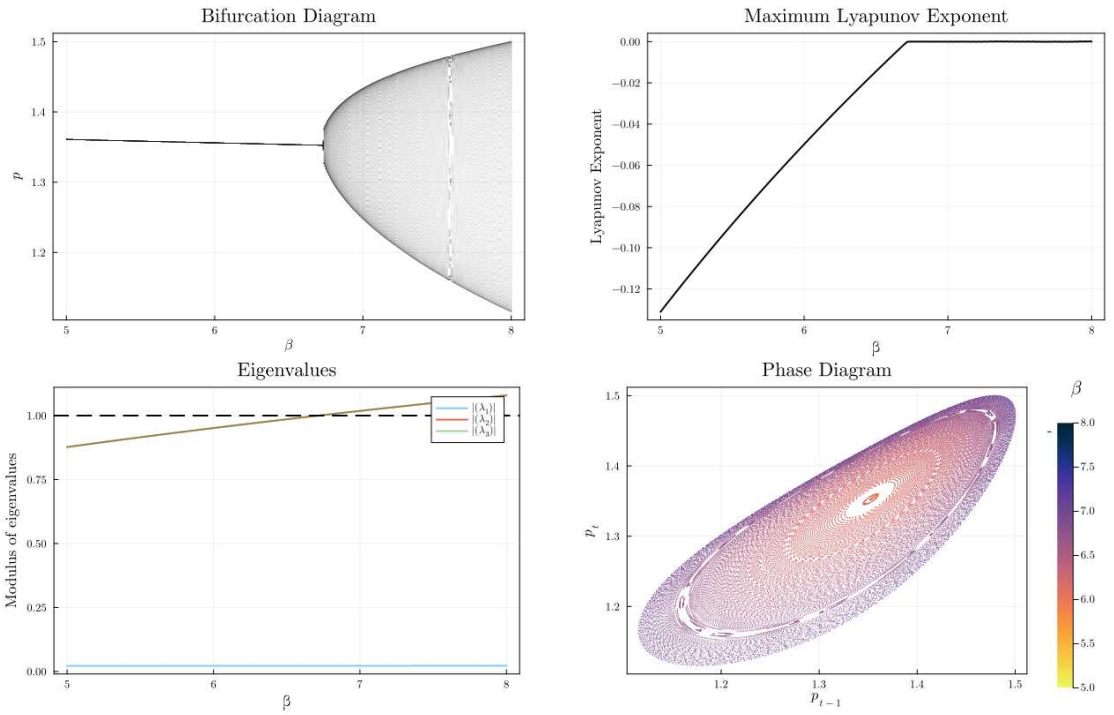

Figure 2.1: Numerical simulations for the two type system, other parameters are $g = 1.3$, $b = 1$, $R = 1.01$. A Hopf bifurcation occurs for $\beta \approx 6.725$, as the two complex eigenvalues have modulus $|\lambda_2| = |\lambda_3| = 1$. After the bifurcation, periodic and quasi-periodic orbits are created.

We then investigate the mechanism responsible for the orbit observed by plotting the trajectory of the system and the fractions' evolution for a value of the intensity of choice $\beta = 7.0$, in Figure 2.2. The bubble and burst behavior is caused by oscillating periods of optimism and pessimism. Pure bias agents act in a way that is qualitatively similar to fundamentalists in the original Brock and Hommes (1998) paper, but their bias causes the fluctuations to be translated upward. The bubble is sustained by a growing population of trend chasers in the market. As long as the next period has a higher concentration of trend chasers, their predictions are self-fulfilling. When their proportion reaches a value close to 55% of the population, however, the price starts to stagnate, implying a decline in the profitability of this strategy. The next period then sees a lower share of these agents in the market and so on, bursting the bubble.

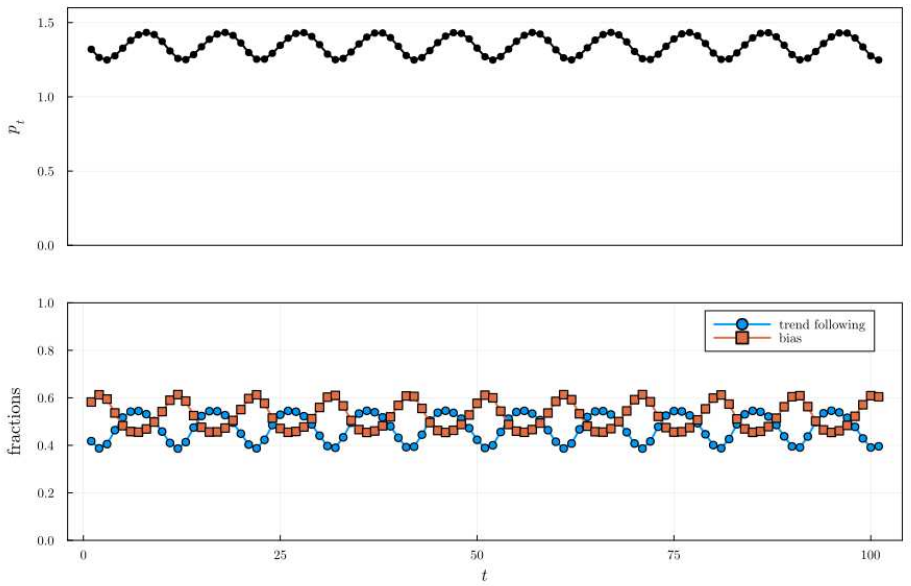

Figure 2.2: Time series of price deviation (top) and fractions (bottom) for $\beta = 7.0$

2.3.3 Introducing Speculation

Given the presence of non-rational behaviors among investors, economic theory would predict that a rational speculator entering the market will earn positive returns with their activity, until irrational investors are pushed out of the market. However, to put it in the words of Keynes: *"Markets can stay irrational longer than you can stay solvent"*. To test this hypothesis, we introduce two types of speculators in the two-types model we have just analyzed. We remark again that being rational in an irrational world is crucially different from being fundamentalist. Moving forward, therefore, we will compare the two cases: when speculators are fundamentalists and when they take into account the "irrationality" of others.

2.3.4 Trend chasers vs pure bias vs fundamentalists

We now study a three-type market. The first two strategies are the same as before, the third class of agents are fundamentalists as in (2.2.9). The resulting three-type model is then given by

$$Rp_t = n_{1,t} g p_{t-1} + n_{2,t} b,$$

with

$$n_{1,t} = \frac{exp\{\beta\pi_{1,t-1}\}}{\sum_{j=1}^{3} exp\{\beta\pi_{j,t-1}\}}, \quad n_{2,t} = \frac{exp\{\beta\pi_{2,t-1}\}}{\sum_{j=1}^{3} exp\{\beta\pi_{j,t-1}\}}.$$

As before, we offer a lemma about the existence of steady states of the model.

Lemma 2 (Steady States of the model with fundamental speculators)**.** *For $\beta = 0$, the model with the fundamentalist speculator has a positive and unique steady state $p = \frac{b}{3R-g}$, which is locally stable for $g < 3R$. For $\beta \to +\infty$ there are no positive steady states.*

Proof. See Appendix 2.A.

We observe that when no switching is present in the model, that is, the case $\beta = 0$, the resulting steady state is lower than the one in the benchmark two-types model. This is consistent with the price reflecting the fundamentalist expectations of the third category of investors. Just as before, when β is strictly positive, we plot the steady state, derived implicitly, in Appendix 2.B. To then analyze the impact of these agents on global dynamics, we turn again to numerical simulations.

2.3.5 Numerical Simulations: b = 1.0, g = 1.3, R = 1.01

Just as before, we fix all parameters except for the intensity of choice β. In the top left panel of Figure 2.3, we can observe that the system has a stable steady state for values of the intensity of choice lower than approximately 5.6. After the parameter crosses this value, a bifurcation occurs. Again, in order to characterize this bifurcation and the system dynamics after bifurcation, we use maximum Lyapunov exponents and a plot of the modulus of the system's eigenvalues. As we can see in the top right panel of Figure 2.3 the exponent is negative for low values of β, corresponding to a stable steadystate, becomes 0 after the first bifurcation and then becomes negative again, producing stable periodic orbits. In the bottom left panel we plot the modulus of the eigenvalues of the system, that we compute in Section (2.A.4) of the appendix. We can see that just as before for values of $\beta \approx 5.6$ the two complex eigenvalues cross the unit circle and a Hopf bifurcation occurs. Finally, in the bottom right panel, we show the creation of stable orbits after the bifurcation, again allowing us to classify the bifurcation as supercritical. Two considerations must be made at this point: first, the presence of the fundamentalist type pushes the overall price closer to the fundamental one. Second, disagreement is added to the system. This causes periodic and quasi-periodic orbits to occur for lower values of the intensity of choice, and orbits have higher amplitude compared to the two-type model.

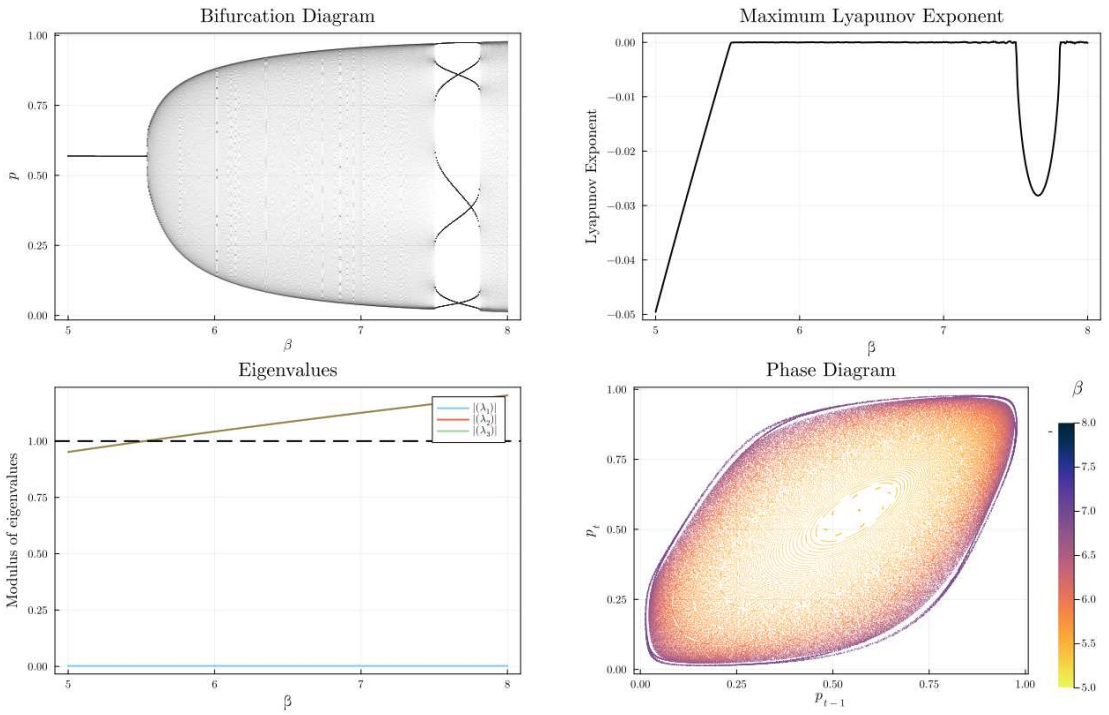

Figure 2.3: Numerical simulations for the fundamentalist speculation system, other parameters are $g = 1.3$, $b = 1$, $R = 1.01$. A Hopf bifurcation occurs for $\beta \approx 5.6$, as the two complex eigenvalues have modulus $|\lambda_2| = |\lambda_3| = 1$. After the bifurcation, periodic and quasi-periodic orbits are created.

Figure 2.4 shows that the dynamics are mainly led by the fundamentalist and the bias categories of investors, whith the trend-chasing strategy never being adopted by more than 40 % of the model's population. The bubbles are now formed after the price is close to the fundamental value. At this point, almost all traders use the fundamentalist strategy. In the next period however the price is going to slowly increase since the few pure bias agents in the market will demand large quantities at this price. This leads to the price slowly increasing and fueling the proportion of trend-chasing agents, sustaining the bubble. The burst has a similar but reversed pattern. When the price is high, a small fraction of fundamentalist agents want to sell high quantities of the asset. Pure bias agents are not ready to accommodate this pressure with their demand since the price is already close to what they believe is the proper valuation of the asset. As a result, the bubble collapses, and the price starts to go back to the fundamental value.

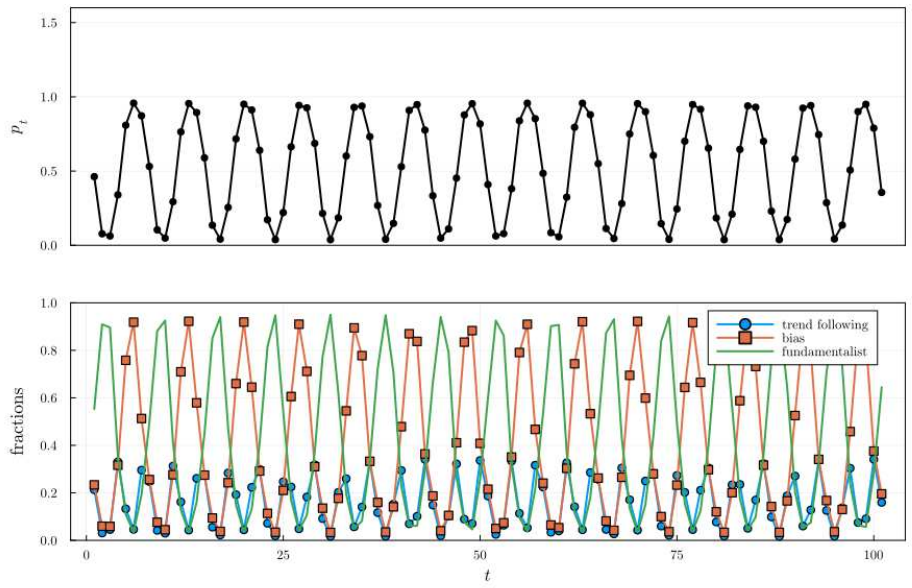

Figure 2.4: Time series of price deviation (top) and fractions (bottom) for $\beta = 7.0$

2.3.6 Trend chasers vs pure bias vs rational expectations

We now study the second three types of market. The first two strategies are the same but we replace fundamentalists with rational expectations[13] as in equation (2.2.10) which gives the following full model (in absence of noise)

$$Rp_t = n_{1,t}gp_{t-1} + n_{2,t}b + n_{3,t}p_{t+1}, \tag{2.3.4}$$

with

$$n_{1,t} = \frac{exp\{\beta\pi_{1,t-1}\}}{\sum_{j=1}^{3} exp\{\beta\pi_{j,t-1}\}}, \quad n_{2,t} = \frac{exp\{\beta\pi_{2,t-1}\}}{\sum_{j=1}^{3} exp\{\beta\pi_{j,t-1}\}}, \quad n_{3,t} = \frac{exp\{\beta\pi_{3,t-1}\}}{\sum_{j=1}^{3} exp\{\beta\pi_{j,t-1}\}}.$$

We offer the usual lemma for the existence of steady states for extreme values of the intensity of choice.

Lemma 3 (Steady States of the model with rational speculators)**.** *Assume $\beta = 0$, then the model with the rational speculator has a positive and unique steady state $p = \frac{b}{3R-g-1}$. When $\beta \to \infty$, then the model has no positive steady state.*

Proof. See Appendix 2.A.

As before, the first observation comes from the case of no switching among strategies,

[13] In this deterministic setting rational expectations coincide with *perfect foresight.*

corresponding to the case $\beta = 0$. The presence of rational speculators implies a steady state value slightly lower than in the baseline model. However, compared to the three-type model with fundamentalist speculators, we observe that the steady state value is much larger. This is the effect of rational speculators knowing that biased agents are going to push the price away from the fundamental value. While in the previous models with only backward looking agents, we were able to provide numerical simulations effortlessly, this is not possible in this case with forward looking rational agents. The presence of p_{t+1} on the right-hand side of equation (2.3.4) makes it difficult to simulate the model. To solve the HSM with rational agents, we then rely on the Extended Path (EP) method proposed by Fair and Taylor (1983)[14]. Before presenting the algorithm, it is worth discussing why it is not possible to simulate the model straightforwardly. One could be tempted to rewrite equation (2.3.4) in such a way as to make the future price explicit

$$p_{t+1} = \frac{Rp_t - n_{1,t}gp_{t-1} - n_{2,t}b}{n_{3,t}}, \tag{2.3.5}$$

hence transforming the problem in a non-linear difference equation of the fourth order. Then, one could simulate the model by starting from some initial conditions. Unfortunately, however, in the unstable region, this approach would converge to the bounded trajectory obtained with the EP method only if the initial conditions are chosen to be exactly on the trajectory. To illustrate the point in Figure 2.5, we plot the path obtained by the EP method and one obtained with the backward representation. We choose parameter values corresponding to a region with quasi-periodic orbits, specifically $g = 1.3, b = 1, \beta = 7, R = 1.01$. For the backward representation, we use initial conditions that are exactly on the EP path (blue) or perturbed by 100 different realizations of normal noise (black). As we can see, all the perturbed paths diverge rather quickly, and the path obtained with the same initial conditions also eventually diverges. This is because small numerical errors eventually accumulate, pushing the system out of the stable path.

[14]Boehl and Hommes (2025) recently proposed a new algorithm that achieves higher accuracy compared to the EP in a similar setting. In our specific case, we find that the EP obtains a sufficient accuracy, and as such, using a more sophisticated method is not needed.

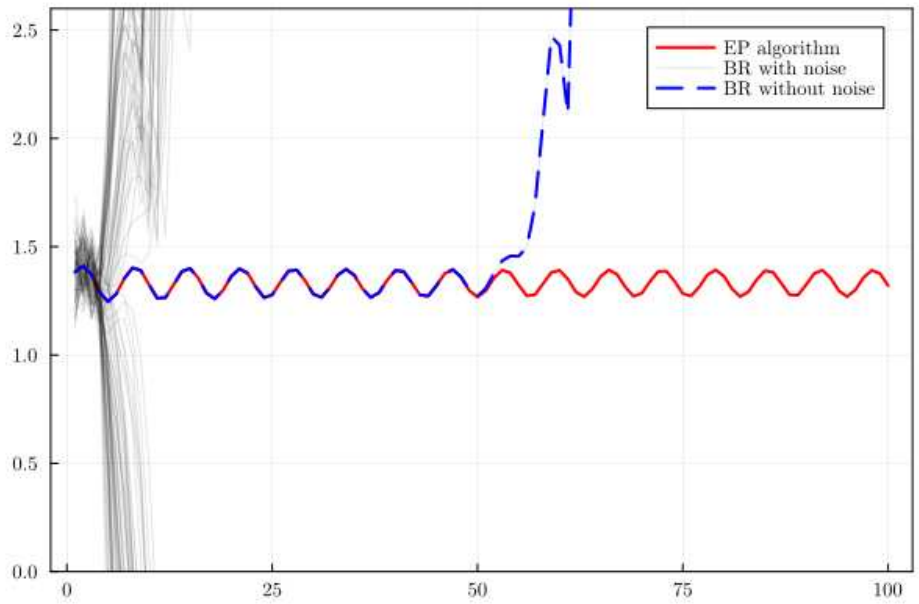

Figure 2.5: Comparison of trajectories obtained by the EP method and the backward representation.

Obtaining Rational Expectations with the Extended Path Method

It is useful to highlight the object of our interest and explain why we need the algorithm in the first place. Start by assuming that expectations formed by the rational agent in the past are indeed rational, that is, in the deterministic model, $\mathbb{F}_{t-3}(p_{t-1}) = p_{t-1}$. The rational agents then needs to choose $\mathbb{F}_{t-1}(p_{t+1})$, considering the following: after the choice is made, a realization of p_t is going to be determined by equation (2.3.4). Given this and $\mathbb{F}_t(p_{t+2})$ one can shift (2.3.4) forward one period to obtain the realization p_{t+1}. In order for an agent to have rational expectations we then require $\mathbb{F}_{t-1}(p_{t+1}) = p_{t+1}$. Hence, when forming expectations at time $t-1$, the agents must consider how expectations are going to be formed at time t. At time t this is still true, and the agent will have to consider what their expectations are going to be at time $t+1$ and so on. Essentially, the whole future *path* of expectations is relevant for the choice of expectations today and, hence, for the actual realization of the state variable. We are now ready to describe the EP algorithm. We are interested in a path $\mathcal{P}_t = (p_t, p_{t+1}, \ldots, p_{t+T})$ of rational expectations, to obtain it:

(i) Choose an integer k, an initial guess of the number of periods beyond the horizon T, to obtain a solution which differs from rational expectations below a tolerance level ϵ. Generate an initial expectations vector, including the necessary lags ($L = 3$ in the case of equation (2.3.4))
$\mathcal{E}^0 = (p_{t-L}, \ldots, p_{t-1}, p_t, p_{t+1}, \ldots, p_{t+T}, \ldots, p_{t+T+k})$;

(ii) Use equation (2.3.4) and then shift the resulting vector to obtain $\mathcal{E}^1 = (p_{t-L}, \ldots, p_{t-1}, p_t, p_{t+1}, \ldots, p_{t+T}, \ldots, p_{t+T+k})$;

(iii) Compute the sum of the absolute deviations between $\mathcal{E}^0$ and $\mathcal{E}^1$. If this is less than ϵ set $\mathcal{E}^0 = \mathcal{E}^1$ and return to step (i). These iterations are called *Type 1.* Call $\mathcal{E}_k$ the vector obtained after convergence;

(iv) Repeat steps (i) to (iii) by replacing k with $k+1$. Call $\mathcal{E}_{k+1}$ the vector obtained after convergence. Compute the sum of the absolute deviations between the corresponding elements in $\mathcal{E}_k$ and $\mathcal{E}_{k+1}$. Iterate until it holds that $|\mathcal{E}_{k+i} - \mathcal{E}_{k+i+1}| < \epsilon$ for some i. These iterations are called *type 2*;

(v) The rational *path* vector is given by the corresposding $T+1$ elements of the vector $\mathcal{E}_{k+i}$.

As Fair and Taylor (1983) noted in the original paper "for a general non-linear model there is no guarantee that any of the iterations will converge. If convergence is a problem, it is sometimes helpful to "damp" the successive solution values." In practice, this will be true in our case when the system enters the unstable region. To deal with this we change the updating mechanism such that the new vector is a convex combination, with weights equal to 0.5, of the original vector and of the new one. Finally, to avoid the algorithm running indefinitely, we set a maximum of 1000 *type 1* iterations and 100 *type 2* iterations. The parameters of the algorithm are then set to $T = 2000$, $\epsilon = e^{-14}$ and $k = 100$. In Figure 2.6 we report the sum of absolute deviations between the rational expectation vector and the realizations of state variables for increasing values of the intensity of choice β on the x-axis. We can see that although the error is slightly higher in the unstable region, its magnitude is almost always lower than 8×10^{-15}.

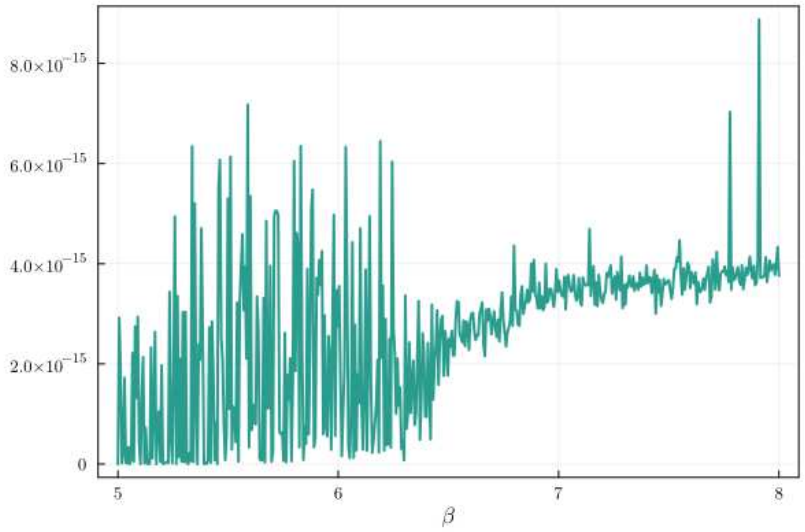

Figure 2.6: Sum of absolute errors for the Extended Path algorithm.

2.3.7 Numerical Simulations:b = 1.0, g = 1.3, R = 1.01

As with the previous two models, we study global dynamics of the system for fixed values of the parameters b, g and R, and varying the intensity of choice β. In the left panel of Figure 2.7, we plot the bifurcation diagram for increasing values of the intensity of choice β. We can observe that the system exhibits a stable steady state for values of the intensity of choice smaller than 6.7. After this value, a bifurcation occurs, and the system enters an unstable region. Unfortunately the presence of the rational agents prevents us from analytically computing the eigenvalues of the system, and we can not fully characterize the bifurcation.[15] However, by means of a phase plot in the right panel, we observe that the system exhibits quasi-periodic orbits with longer periodicity compared to the system with the fundamentalist speculators. In fact, the dynamics are similar to those of the two-type model.

[15]This feature also prevents us to numercially compute the Lyapunov exponents as in the previous cases, since the forward-time evolution cannot be computed directly without already knowing future states.

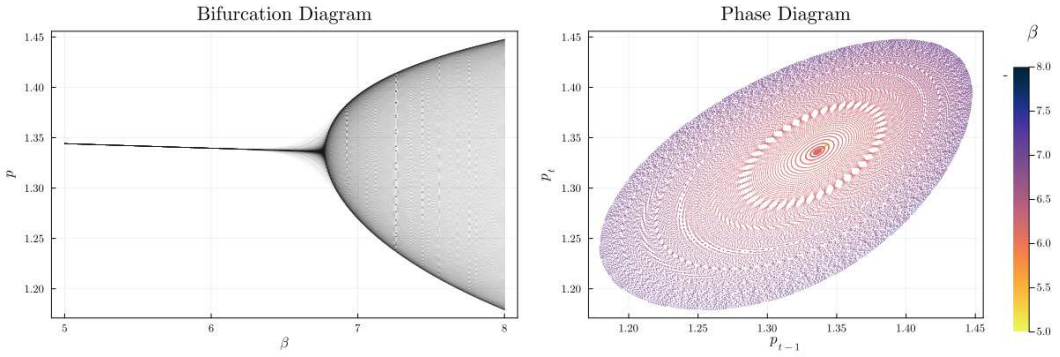

Figure 2.7: Numerical simulations for the model with rational speculators, other parameters are g = 1.3, b = 1, R = 1.01.

The main consideration to make is that, unlike the case with the fundamentalist speculators, the system does not approach the fundamental value, even after the bifurcation occurs. To get a better sense of this phenomenon, in Figure 2.8 we plot the fractions evolution for a value of the intensity of choice $\beta = 7.0$.

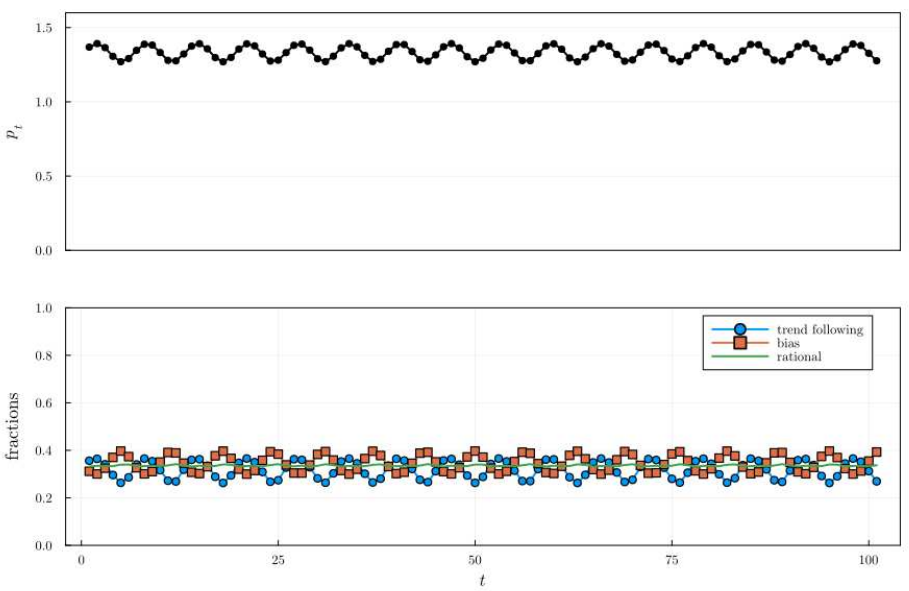

Figure 2.8: Time series of price deviation (top) and fractions (bottom) for $\beta = 7.0$

We can see that the fraction of agents using the rational forecast fluctuates significantly less than the other two. Therefore, there are two reasons why the resulting price is further away from the fundamental value. The first reason is that to produce an accurate forecast, rational speculators must take into account the behavior of the other two agents. This results in a forecast significantly higher than the one a fundamentalist agent would make. The second one is that, as we already remarked, realized profits and

forecast accuracy are not perfectly correlated. Speculators using the rational forecast are always present in the system, but they are not able to drive other strategies out of the market. The bubble and burst motif is similar to the one in the two-type model. The bubble is sustained by a growing fraction of trend chasers. Notice, however, how the fraction of rational agents increases right after the beginning of the bubble and after its collapse. This is the result of this strategy being able to forecast the bursts and booms perfectly. An additional result of this behavior is that the oscillations are dampened, as we discuss in the next Section.

2.3.8 Speculation Effects

As we can see in Figure 2.9, the two types of speculation have different effects on the market. When speculators are fundamentalist, they drive prices toward the fundamental benchmark. However, their presence results in increased disagreement, which in turn causes the model to enter the unstable region for lower values of the intensity of choice. Moreover, in the unstable region, the resulting volatility is substantially increased. When speculators are rational, we can observe that prices are, on average, slightly lower than those in the no speculation case but still far from the fundamentalist benchmark. An interesting result is that rational speculation decreases volatility in the unstable region with respect to the two-type model. This is because speculators can predict the bursting of the bubble and take positions before its realization. By doing so, they dampen the oscillations in the market. Rational speculation can have a positive effect on reducing volatility at the cost of keeping prices further away from the fundamental price that would be achieved in a homogeneous rational world.

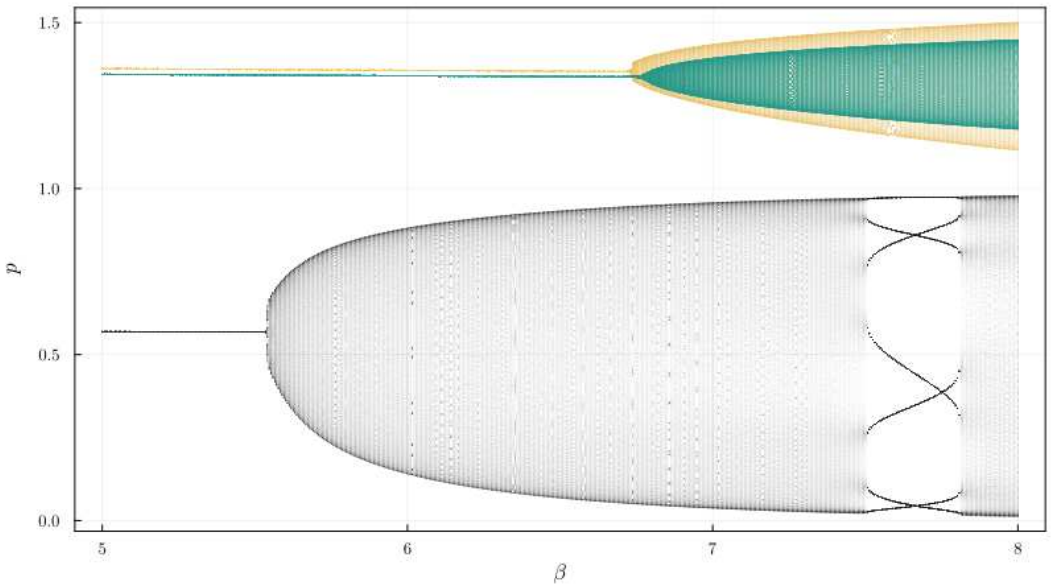

Figure 2.9: Bifurcation diagrams for the two-types model (orange), the model with rational speculators (green) and the one with fundamentalist speculator (black), other parameters are g = 1.3, b = 1, R = 1.01.

2.4 Taking the model to the data

We now estimate the theoretical model investigated so far on real Bitcoin data. The literature has successfully estimated heterogeneous agents switching models on multiple markets, as discussed in Section 2.1.1. Their appeal is that they offer an endogenous explanation, driven by the switching among different strategies, to observed stylized facts of financial markets, like bubble and burst sequences in prices and conditional heteroskedasticity of returns. The literature has mainly focused on estimating models with two strategies available to investors: mean-reversion and trend-following. This feature makes the empirical model a non-linear autoregressive process, which is then usually estimated by NLS. We contribute to this body of work with two novelties. First, by estimating models with biased investors. To do this, we quantify the bias with a sentiment index obtained from Twitter textual data. Second, by proposing a fast methodology to estimate the model, including rational, forward-looking agents. To our knowledge, we are the first to estimate an HSM with fully rational agents.

2.4.1 Bitcoin Twitter Sentiment Index

To construct a proxy for market sentiment toward Bitcoin, we utilize data from Twitter (currently known as X), assuming that a Tweet conveys part of what a specific investor

feels regarding the asset. We scraped the website for all posts containing the term "Bitcoin" from November 1, 2019, to December 30, 2022. We used standard preprocessing steps to derive a meaningful sentiment index from this textual data by removing stopwords and non-alphabetic characters. We then employed a Machine Learning algorithm for Natural Language Processing with the VADER sentiment analysis library, introduced by Hutto and Gilbert (2014). VADER determines the sentiment of text using a mix of lexical heuristics and a sentiment lexicon, and it can handle idiomatic language and sarcasm. Sentiment scores, ranging from -1 (highly negative) to +1 (highly positive), with 0 indicating neutrality, were computed for each tweet based on the presence of positive, negative, and neutral words. We then aggregated the daily scores to form the Bitcoin Twitter Sentiment Index (BiTSI). Figure 2.10 shows the evolution of the index and the daily closing price of Bitcoin in US Dollars. To contextualize the index's peaks and troughs, we highlight the headlines of articles associated with the three highest and lowest BiTSI values, offering insights into the news events likely driving these sentiment shifts.

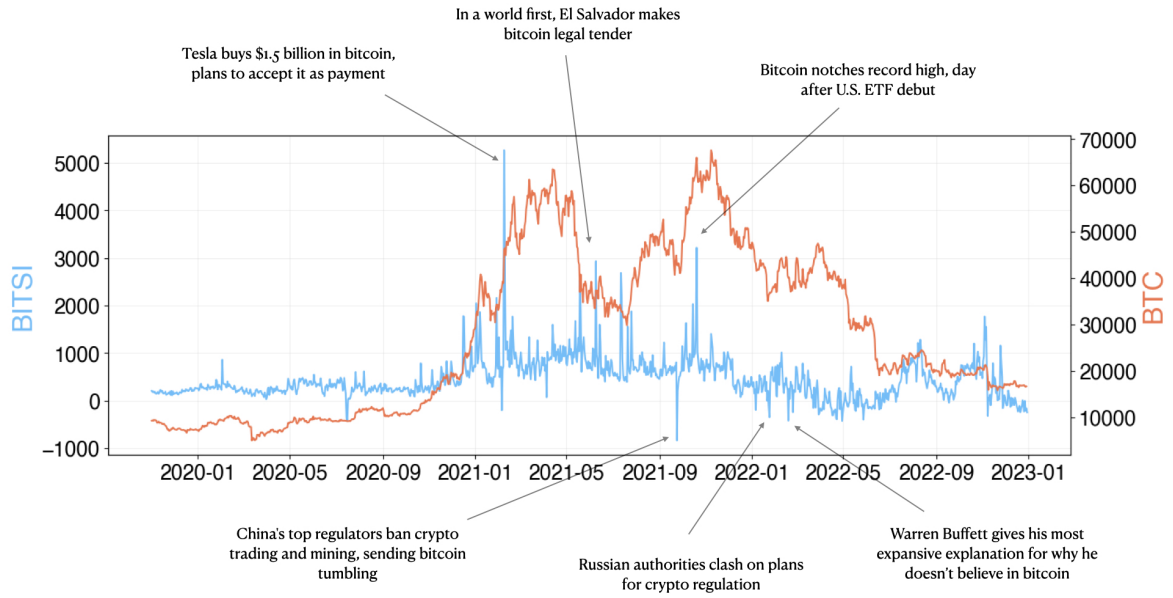

Figure 2.10: Bitcoin and BITSI time series, with relevant events

We report descriptive statistics for the index and the number of tweets per day in Table 2.1.

Table 2.1: Descriptive statistics for the BiTSI and number of Tweets per day

	BiTSI	**Volume**
mean	455.85	8626.90
std	426.51	5160.04
min	-841.15	650
max	5264.47	56179
skew	2.58	1.66
kurtosis	18.58	8.06
obs	1117	1117

As a first analysis, we demonstrate that our index is not capturing information from other sources highlighted by the literature as having explanatory power on the cryptocurrency return. Similar to traditional factor models, Shen, Urquhart and Wang (2019) and Liu, Tsyvinski and Wu (2022) show that three factors: market, size, and momentum can substantially explain the cross-section of cryptocurrencies' excess returns. In Figure 2.11, we report the correlation between the index and other variables that should proxy, at the daily level, the role of the three factors mentioned above. Excess return is measured as the difference between daily Bitcoin returns and the daily risk-free rate, represented by the Market Yield on U.S. Treasury Securities at 1-Year Constant Maturity and available on FRED.[16] We then have two proxies for momentum, given by the first and second difference of the daily closing price. Finally, 'Volume' represents the daily volume of Bitcoin. All Bitcoin data are gathered from Yahoo Finance. In order to check for the possibility of the delayed impact of these factors on the index, we also reported lagged values of all the variables, represented by the '(t-1)' label at the end of the symbol.

[16]The yearly rate is then transformed into the corresponding daily one by $r_{daily} = r_{yearly}/365$. The choice is motivated by the yearly capitalization of the instrument.

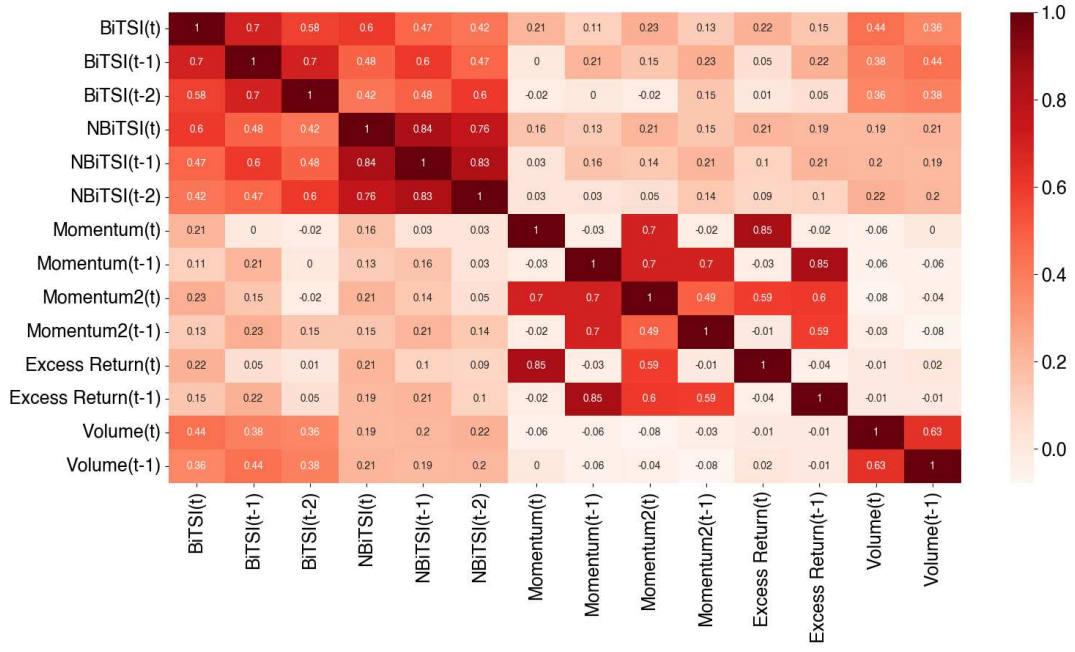

Figure 2.11: Correlation of the BiTSI and other variables

The correlation between the index and all the other variables is mostly positive and relatively small, except the one with Volume. The result is, however, explained by the correlation, already documented in the literature, between Bitcoin volume and Tweet volume. This is why we also include the Normalised BiTSI, where the normalization factor is the volume of Tweets per day. After removing the magnitude effect, the index appears to be only mildly correlated to the variables, therefore making it a good regressor for the model.

2.4.2 Estimating Non-linear Rational Expectation models with a Neural Network

In order to estimate the model with rational speculators, we first need to obtain the path of rational expectations[17]. In their paper, Fair and Taylor (1983) propose a maximum likelihood estimation built on the EP algorithm. The procedure involves solving the model for rational expectations given multiple combinations of parameters and then computing the log-likelihood associated with each combination.

For our case, this is computationally expensive as we are interested in the three behavioral parameters, g, b, and β, and would require a dense grid to evaluate the likelihood, given the highly non-linear structure of the problem. To summarize, their

[17]This can not simply be obtained by taking leads of the independent variables, as the term would also include the noise term.

procedure suffers from the *curse of dimensionality* as every extra parameter to be estimated increases computational time exponentially. To address this issue, we propose an estimation procedure that relies on a Neural Network and consists of two key components. The first is to ensure that we can express the state variable as a function of past states only. The second is that we can approximate this unknown function satisfactorily. A general formulation for our specific problem in (2.3.4) is of the form

$$p_t = \mathcal{F}(\{p_{t-l}\}_{l=0}^{L}, \{\mathbb{E}_t(p_{t+k})\}_{k=1}^{K}, \{w_{t-l}\}_{l=0}^{L}; \alpha) + \varepsilon_t$$

where now ε_t is i.i.d. noise, and we want to estimate the parameter vector α. Defining by $\mathcal{F}^{(n)}$ the function obtained from $\mathcal{F}$ by iterating all future expectations of the state variables n-1 times, we have that if

$$\lim_{n\to+\infty} \frac{\partial \mathcal{F}^{(n)}}{\partial \mathbb{E}_t(x_{t+k})} = 0, \text{for all } k > 0, x \in \{p, w, \varepsilon\},$$

and there exists a finite M such that

$$p_{t+k} \in [-M, M], \text{for all } k \geq 0,$$

then there exist a generic function $\mathcal{H}$ of past states only, a N and an $\epsilon > 0$ such that for all $n > N$

$$\left|\mathcal{F}^{(n)}(\{p_{t-l}\}_{l=0}^{L}, \mathbb{E}_t\{(p_{t+k})\}_{k=0}^{K}, \{w_{t-l}\}_{l=0}^{L}; \alpha) - \mathcal{H}(\{p_{t-l}\}_{l=0}^{L}, \{w_{t-l}\}_{l=0}^{L}; \alpha)\right| < \epsilon$$

This ensures that by iterating the function with respect to future variables, we eventually reach a point where the impact of the future on today becomes negligible. Unfortunately, demonstrating that a general function $\mathcal{F}$ satisfies the assumption is not trivial in non-linear cases.[18] We can only assess that this condition is met for realistic values of the parameters, by ex-post veryifing that the method offers a good approximation of the rational expectation orbit.

Having such a function $\mathcal{H}$ expressed only in terms of past variables is crucial for the estimation as it removes any iterative procedure needed to obtain the rational expectation forecast. Unfortunately, the functional form is unknown and cannot be derived analytically. Dealing with this unknown function is where the second component of our estimation procedure comes into play. As anticipated, this step involves obtaining an approximation to the unknown function $\mathcal{H}$, and we rely again on a Machine Learning

[18] For a detailed discussion, see Boehl and Hommes (2025).

model by using a Neural Network (NN) to obtain

$$\mathcal{H}(\{p_{t-l}\}_{l=0}^{L}, \{w_{t-l}\}_{l=0}^{L}; \alpha) \approx \hat{\mathcal{H}}(\{p_{t-l}\}_{l=0}^{L}, \{w_{t-l}\}_{l=0}^{L}; \theta),$$

where now θ are the parameters of the NN. We deem this appropriate as it was shown in Hornik, Stinchcombe and White (1989) that multilayer feedforward networks are universal approximators.

We now specialize to the case at hand. Start from

$$p_t = \mathcal{F}(p_{t-1}, p_{t-2}, p_{t-3}, p_{t+1}, w_{t-1}, w_{t-3}; \alpha) + \varepsilon_t.$$

we get

$$\mathbb{E}_t(p_{t+1}) = \mathcal{H}(p_{t-1}, p_{t-2}, p_{t-3}, w_{t-1}, w_{t-2}, w_{t-3}; \alpha),$$

and approximate it via a NN

$$\mathbb{E}_t(p_{t+1}) = \hat{\mathcal{H}}(p_{t-1}, p_{t-2}, p_{t-3}, w_{t-1}, w_{t-2}, w_{t-3}; \theta),$$

with architecture described in Appendix 2.C. After obtaining a vector of rational expectations $\{\mathbb{E}_{t-l}(p_{t+1-l})\}_{l=1}^{L}$ we can proceed with the estimation by including it as a regressor. To demonstrate that our method can indeed provide a satisfactory approximation of rational expectations, we conduct the following exercise. We generate synthetic data for the model with trend chasers vs pure bias vs rational speculators, for two values of the intensity of choice $\beta \in \{3.0, 7.0\}$, corresponding to the stable and unstable regions, while keeping all other parameters fixed at their levels of $g = 1.3$, $b = 1.0$, $R = 1.01$. For each of these two regions, we generate a model with a Noise-to-Signal ratio of 0 (corresponding to the deterministic case), 0.01 and 0.1. We then train the NN on the synthetic data and compute the Mean Squared Error (MSE) between the NN estimates and the true rational expectations vector. Using the NN estimates, we estimate the three parameters g, b and β with NLS and report the point estimate and the associated standard errors in parentheses. Results are shown in Table 2.2.

Table 2.2: Estimation on synthetic data

	Stable region, $\beta = 3$			**Unstable region,** $\beta = 7$		
N/S	0	0.01	0.1	0	0.01	0.1
MSE	0.0000	0.0009	0.0009	0.0037	0.0050	0.0235
g	1.17 (0.00)	1.26 (0.03)	1.26 (0.03)	1.29 (0.00)	1.25 (0.02)	1.10 (0.01)
b	1.17 (0.00)	1.07 (0.04)	1.07 (0.04)	1.01 (0.00)	1.05 (0.02)	1.11 (0.02)
β	0.88 (0.00)	3.10 (0.89)	3.07 (0.88)	7.40 (0.05)	6.79 (1.02)	8.85 (1.28)

The NN is able to approximate the rational expectations vector with an extremely low Mean Squared Error. Clearly, the approximation deteriorates with increasing Noise-to-Signal ratio. This, however, does not seem to impair the estimation. Estimates for the g and b parameters are close to their true value, with the notable exception of the first column. When the dynamics are stable and there is no noise, there is simply no variance to correctly estimate the parameter values. The estimation of the intensity of choice parameter β is associated with the higher standard errors. This is in line with findings in the literature Boswijk et al. (2007), ter Ellen et al. (2021) documenting that pinning down this parameter is challenging. The reason is that large changes in this parameter have only mild effects on the overall fractions, given the smoothness of the multinomial logit function. In any case we find that the point estimate is always within the confidence interval, except when no noise is present.

2.5 Estimation

We now turn to the estimation on real data and accommodate some changes to incorporate financial data into the theoretical model. First, we use a time-varying risk-free rate, proxied by the Market Yield on U.S. Treasury Securities at 1-Year Constant Maturity, as discussed above, so that instead of having a fixed R, we have a time-varying R_t. Second, we address the non-stationarity of the data. We follow Kukacka and Barunik (2017) and ter Ellen et al. (2021) in rewriting the model in deviation from a moving average[19]

$$\mathrm{MA}_t = \frac{\sum_{i=W}^{1} p_{t-i}}{W},$$

where W is the window size, so that our state variable is $x_t = p_t - \mathrm{MA}_t$. In the baseline, we fix it to 40 days. We discuss the sensitivity of the estimation to this choice in

[19]This implies a fundamental value for the asset different than 0. For a discussion on the fundamental value of cryptocurrencies we can refer to Kukacka and Kristoufek (2023).

Appendix 2.E. Third, we relax the assumption of constant beliefs about the variance of returns and allow for time varying volatility, proxied by a non-centered second moment. We assume it is homogeneously equal to the square of the last available moving average $(\mathrm{MA}_{t-3})^2$ and fix the coefficient of absolute risk aversion a to 1. Lastly, we address the heteroskedasticity of residuals by rewriting the model in percentage deviation from the moving average fundamental, by simply dividing both sides of the pricing equation by the fundamental value itself.

The model we take to the data is then the following

$$R_t \frac{x_t}{\mathrm{MA}_t} = n_{1,t} g \frac{x_{t-1}}{\mathrm{MA}_t} + n_{2,t} b \frac{w_{t-1}}{\mathrm{MA}_t} + n_{3,t} \frac{\mathbb{F}_{3,t}(x_{t+1})}{\mathrm{MA}_t} + \varepsilon_t, \tag{2.5.1}$$

where $n_{3,t}$ will be 0 in the two-types model, and $\mathbb{F}_{3,t}(x_{t+1})$ will either be the fundamental or the output of the Neural Network model in the rational expectation model.

Fractions are given as the usual multinomial logit, where now the fitness measure is given by

$$\pi_{1,t} = \frac{1}{\mathrm{MA}_{t-2}^2} \left(x_t - R_{t-2} x_{t-1}\right) \left(g x_{t-2} - R_{t-2} x_{t-1}\right),$$

$$\pi_{2,t} = \frac{1}{\mathrm{MA}_{t-2}^2} \left(x_t - R_{t-2} x_{t-1}\right) \left(b w_{t-2} - R_{t-2} x_{t-1}\right),$$

$$\pi_{3,t} = \frac{1}{\mathrm{MA}_{t-2}^2} \left(x_t - R_{t-2} x_{t-1}\right) \left(\mathbb{F}_{3,t-1}(x_t) - R_{t-2} x_{t-1}\right),$$

which are the usual profits divided by the square of the moving average fundamental. We then plot and test the variables for stationarity in Appendix 2.D.

We are finally ready to estimate the parameters g, b, and β. Thanks to the methodology we have discussed above, we can express all the regressors in terms of past variables only, and non-linear least squares can estimate the three models presented. We are particularly interested in checking if there is evidence of switching between strategies and if the switching is robust to different strategies that might be present in the market[20]. For this reason, we estimate all three models presented above and report and compare the results in Table 2.3.

[20] As we only use aggregate data, we can not use this exercise to say which strategies are more likely to be used in the market. It is, however, useful to study the empirical implications of the models we have analyzed only theoretically so far.

Table 2.3: Estimation results

	Trend chasers vs pure bias	**Tc vs B vs Fundamentalist**	**Tc vs B vs Rational Expectations**
g	1.91 (121.66)	2.86 (121.51)	1.92 (84.14)
b	0.58 (6.5)	0.87 (6.5)	0.83 (6.41)
β	1.03 (2.12)	0.51 (2.14)	1.45 (2.05)
Adj r^2	0.94	0.94	0.94
F-stat	4.61 (0.03)	4.74 (0.03)	4.34 (0.04)
het	0.17	0.17	0.11

We report the point estimate of the parameters and the associated t-statistic in parentheses. We also report the adjusted R-squared value, along with the value and associated p-value of the F-test for the significance of the non-linear model relative to a linear one. The linear model can be seen as nested within the non-linear model, corresponding to a value of the intensity of choice β equal to 0 and a market in which agents do not switch between strategies. So if the estimated value of β is significantly different from 0, we can confirm evidence of switching. Lastly, we report the p-value associated with Engle's Test for Autoregressive Conditional Heteroscedasticity on the estimation residuals. The null hypothesis in the test is that of homoskedasticity and we can not reject it in any of the specifications given the p-values. Nonetheless, we re-run the estimation by obtaining confidence intervals for the parameters using a heteroskedasticity-robust bootstrap procedure, as detailed Appendix 2.F. The first type of investors are strong trend chasers[21], expecting deviations from fundamentals to continue with increased magnitude in subsequent periods. This behavior aligns with psychological factors such as the "fear of missing out" (FOMO) which is particularly relevant in the cryptocurrency market and is documented, for example, by Baur and Dimpfl (2018). Their extrapolation becomes most pronounced in the model with a fundamentalist agent. This can be explained by the presence of a type of investor in the model that always forecasts the fundamental price, thereby absorbing what might have previously been classified as mild trend chasers. The second category of investors can be classified as weak sentiment followers, as they constantly expect a small deviation from the fundamental moving average, which is positively proportional to the sentiment index. The F-test confirms evidence of non-linearities in every case, at least at the 5% level.

The small but significant estimate for β implies switching among the different strategies. A typical result in the estimation of HSM is that the parameters are close to the

[21]To be more precise, in this version of the model, trend chaser agents are expecting *percentage deviations* to persist, which also justifies the high estimates for the parameter regulating their behavior.

bifurcation value, (Lux, 2009). This result also holds in our case, and we use the following strategy to verify this claim. We simulate equation (2.5.1) starting from initial conditions based on real data to obtain a path of x_t. We use actual endogenous variables w_t, R_t, and we use the residuals of the estimation for the noise term ε_t. After exploring the parameter space, we find that the crucial parameter in this setting is the strength of the trend-chasing strategy g. We, therefore, vary the parameter g and report a bifurcation diagram for the two-type and the three-type model with fundamentalist speculators[22] in Figure 2.12. The vertical red line marks the estimated value of g, and the associated standard error is reported with a shaded area. As we can see, the estimated value of g is right before a bifurcation point. For a given g^*, within the estimated confidence interval, the system is stable. At the same time, for a value slightly lower than this critical value, shocks can trigger waves of optimistic trend-chasing, which are self-reinforcing and can sustain themselves for several periods. Interestingly, if the value of g is higher than the critical value, it appears that the strong trend-chasing expectations are not realized, and the fraction of agents employing this strategy decreases substantially.

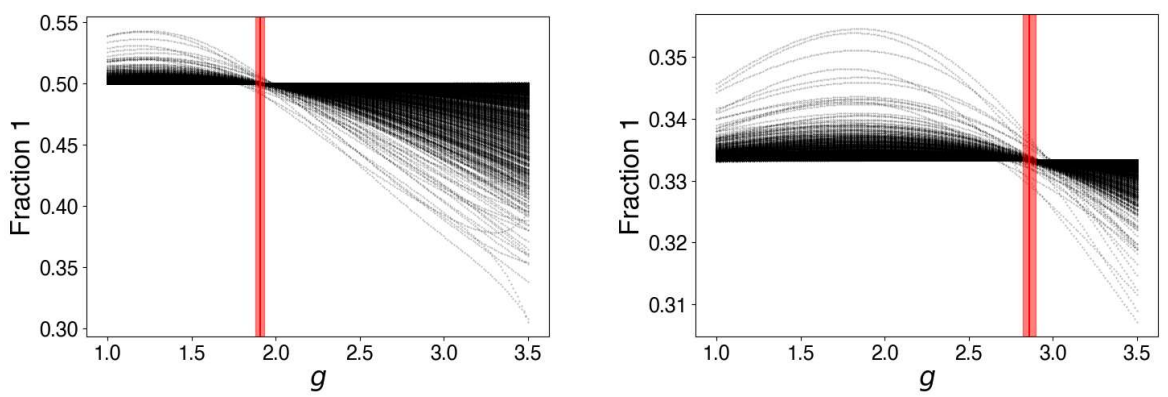

(a) Trend chasers vs bias

(b) Trend chasers vs bias vs fundamentalists

Figure 2.12: Bifurcation diagram for the empirical models

The implication of this finding on the model dynamics is that large enough shocks can temporarily affect the stability of the steady state and drive the model into the unstable region, resulting in volatility amplification that can last for several periods (days). In Figure 2.13, we plot the implied fractions for the three models at the estimated parameter values to further investigate this effect. We can observe a long

[22]We cannot provide a similar analysis for the three-type model with rational speculators. The reason is that without having counterfactual data for different values of the parameters, we can not use the methodology described above to obtain rational expectations. However, by the similarity of the theoretical and empirical results, we can conjecture that this is true also in this case.

window of coexistence of strategies alternated to periods of substantial switching. All three models provide remarkably similar results, offering robustness of the switching mechanism to the different strategies and confirming the underlying mechanism. Sentiment-driven and exogenous shocks affect agents' profits. This, in turn, makes investors more likely to switch to more profitable strategies, adding an endogenous motif that amplifies the underlying volatility.

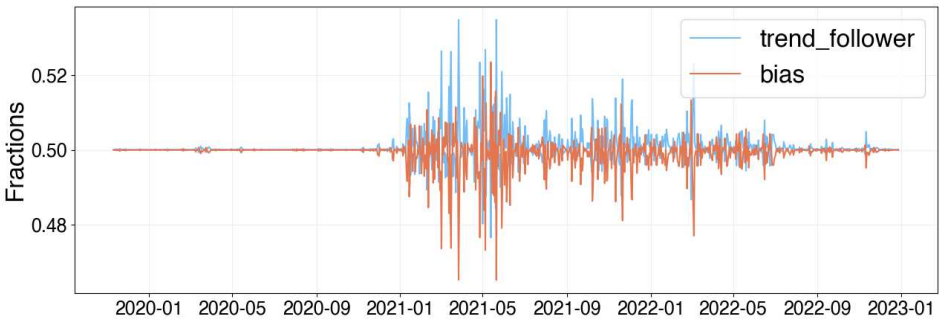

(a) Trend chasers vs bias

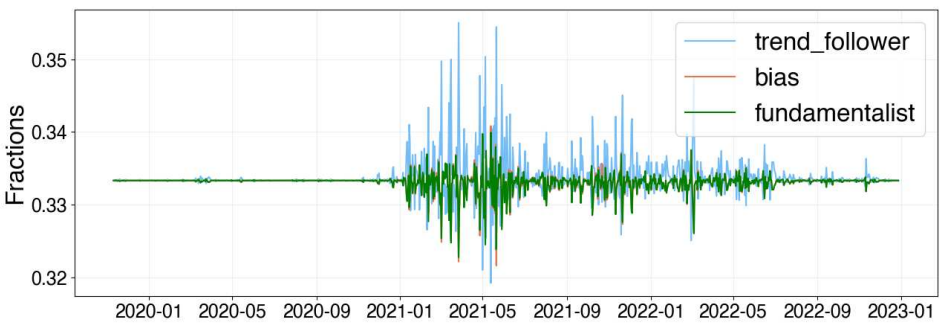

(b) Trend chasers vs bias vs fundamentalists

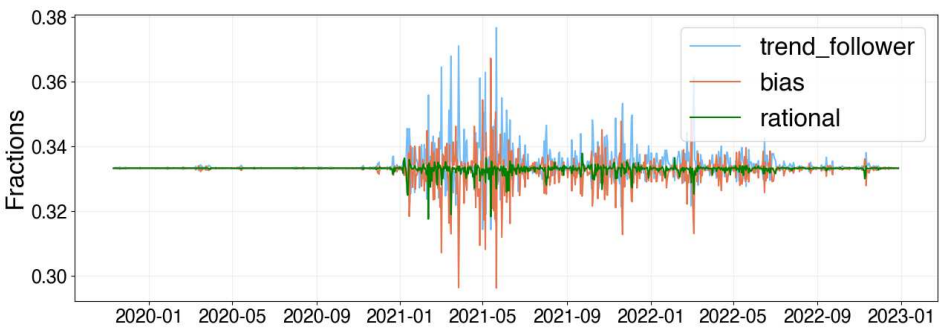

(c) Trend chasers vs bias vs rational expectations

Figure 2.13: Fractions evolution for estimated parameters

2.6 Conclusion and Discussion

In this chapter we have proposed a heterogeneous agent asset pricing model in which different categories of investors evolve endogenously over time. The presence of boundedly rational investors in the market, employing trend-chasing and biased forecasting rules, raises the question of speculation opportunities. We introduced two different types of speculation. The first is associated with fundamentalist traders, who know the true underlying process of the asset and behave accordingly. However, we noted that to have rational expectations in our model, one must consider the (possibly incorrect) strategies of other investors. We showed that when speculators are of the second type, the resulting price deviates further from the fundamental price, but volatility is reduced compared to the fundamentalist speculation benchmark. Hence, rational speculators have a stabilizing effect on the market.

The second part of the chapter is devoted to providing empirical support to our model. We used data from the Bitcoin market and constructed the BiTSI index, which we used as a proxy of a time-evolving bias. The BiTSI index is shown to capture an aspect of the market that is uncorrelated with the main factors highlighted by the literature explaining Bitcoin returns. To estimate the non-linear dynamic rational expectations model, we proposed a methodology using a Neural Network that overcomes the curse of dimensionality. We demonstrated that this method provides a satisfactory approximation of rational expectations on synthetic data.

We then estimate the models and find values of the parameters close to a bifurcation. The implication for the market is that shocks can push the system outside the stable region, trigger the endogenous switching process, and induce windows of persisting volatility.

Appendix 2.A Proofs

2.A.1 Proof of Lemma 1

Positive steady states of the system must solve

$$Rp = n_1 gp + (1 - n_1)b,$$

where

$$n_1 = \frac{1}{1 + exp\{\beta(p - Rp)(b - gp)\}},$$

and satisfy the constraint on the parameters $b > 0$, $g > 0$, $R > 1$ and the positivity requirement $p \geq 0$ When $\beta = 0$ then $n_1 = \frac{1}{2}$ for all p which implies that $p = \frac{b}{2R-g}$. The stability of the system can be determined by studying the first order difference equation

$$p_t = \frac{g}{2R} p_{t-1} + \frac{b}{2R},$$

and imposing

$$\left| \frac{dp_t}{p_{t-1}} \right| = \left| \frac{g}{2R} \right| < 1.$$

Since we are working under the assumption that $g > 0$ and $R > 1$ the steady state is then locally stable for $g < 2R$.

When $\beta = +\infty$, the price function is piecewise-defined as

$$p_t = \begin{cases} \frac{g}{R} p_{t-1} & \text{for} \quad \Delta\pi_{t-1} < 0 \\ \frac{g}{2R} p_{t-1} + \frac{b}{2R} & \text{for} \quad \Delta\pi_{t-1} = 0 \\ \frac{b}{R} & \text{for} \quad \Delta\pi_{t-1} > 0 \end{cases}$$

where $\Delta\pi_{t-1} = (p_{t-1} - Rp_{t-2})(b - gp_{t-3})$, is the difference in realized profits between the bias traders and the trend chasers. To obtain the steady states and their stability, we can then proceed by cases

(i) $\Delta\pi_{t-1} < 0$. Then a steady state p must satisfy the following pair of equations

$$\begin{cases} p = \frac{g}{R} p \\ (p - Rp)(b - gp) < 0 \end{cases}$$

which has solution if and only if $g = R$ with $p < \frac{b}{R}$. The steady-state stability is

determined by

$$\left|\frac{dp_t}{p_{t-1}}\right| = \left|\frac{g}{R}\right| = 1,$$

therefore the steady state is locally unstable.

(ii) $\Delta\pi_{t-1} = 0$. Then a steady state p must satisfy the following pair of equations

$$\begin{cases} p = \frac{g}{2R}p + \frac{b}{2R} \\ (p - Rp)(b - gp) = 0 \end{cases}$$

Since $b > 0$, a solution in this case exists if and only if $g = R$ which then leads to $p = b/R = b/g$. The stability of this solution is given by

$$\left|\frac{dp_t}{p_{t-1}}\right| = \left|\frac{g}{2R}\right| = \left|\frac{1}{2}\right| < 1,$$

which implies that the steady state is locally stable.

(iii) $\Delta\pi_{t-1} > 0$. Then a steady state p must satisfy the following pair of equations

$$\begin{cases} p = \frac{b}{R} \\ (p - Rp)(b - gp) > 0 \end{cases}$$

The solution $p = \frac{b}{R}$ can be sustained only if $g > R$ since $b > 0$. The stability of the steady state is regulated by

$$\left|\frac{dp_t}{p_{t-1}}\right| = 0,$$

implying a locally stable steady state.

2.A.2 Proof of Lemma 2

Positive steady states of the system must solve

$$Rp = n_1 gp + n_2 b,$$

where

$$n_1 = \frac{exp\{\beta\pi_1\}}{exp\{\beta\pi_1\} + exp\{\beta\pi_2\} + exp\{\beta\pi_3\}}, \quad n_2 = \frac{exp\{\beta\pi_2\}}{exp\{\beta\pi_1\} + exp\{\beta\pi_2\} + exp\{\beta\pi_3\}},$$

and

$$\pi_1 = (p - Rp)(gp - Rp), \quad \pi_2 = (p - Rp)(b - Rp), \quad \pi_3 = (p - Rp)(-Rp),$$

aswell as the constraint on the parameters $b > 0$, $g > 0$, $R > 1$ and the positivity constraint.

For $\beta = 0$ then $n_1 = n_2 = \frac{1}{3}$, which implies $p = \frac{b}{3R-g}$. The stability of the system can be determined by studying the first order difference equation

$$p_t = \frac{g}{3R} p_{t+1} + \frac{b}{3R},$$

and imposing

$$\left| \frac{dp_t}{p_{t-1}} \right| = \left| \frac{g}{3R} \right| < 1.$$

Since we are working under the assumption that $g > 0$ and $R > 1$ the steady state is then locally stable for $g < 3R$.

When $\beta = +\infty$, the price function is piecewise-defined as

$$p_t = \begin{cases} \frac{g}{R} p_{t-1} & \text{for} \quad \pi_{1,t-1} > \pi_{2,t-1} \quad \text{and} \quad \pi_{1,t-1} > \pi_{3,t-1} \\ \frac{b}{R} & \text{for} \quad \pi_{2,t-1} > \pi_{1,t-1} \quad \text{and} \quad \pi_{2,t-1} > \pi_{3,t-1} \\ 0 & \text{for} \quad \pi_{3,t-1} > \pi_{1,t-1} \quad \text{and} \quad \pi_{3,t-1} > \pi_{2,t-1} \\ \frac{g}{2R} p_{t-1} + \frac{b}{2R} & \text{for} \quad \pi_{1,t-1} = \pi_{2,t-1} > \pi_{3,t-1} \\ \frac{g}{2R} p_{t-1} & \text{for} \quad \pi_{1,t-1} = \pi_{3,t-1} > \pi_{2,t-1} \\ \frac{b}{2R} & \text{for} \quad \pi_{2,t-1} = \pi_{3,t-1} > \pi_{1,t-1} \\ \frac{g}{3R} p_{t-1} + \frac{b}{3R} & \text{for} \quad \pi_{1,t-1} = \pi_{2,t-1} = \pi_{3,t-1} \end{cases}$$

As before, we proceed by analysing the multiple cases:

(i) $\pi_{1,t-1} > \pi_{2,t-1}$ and $\pi_{1,t-1} > \pi_{3,t-1}$. Then a non-negative steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{g}{R} p \\ (p - Rp)(gp - Rp) > (p - Rp)(b - Rp) \\ (p - Rp)(gp - Rp) > -Rp(p - Rp) \\ p \geq 0 \end{cases}$$

The system, however, has no solution since the first equation implies either $p = 0$ or $g = R$. In the first case, equations two and three are not satisfied. In the latter, the third is not.

(ii) $\pi_{2,t-1} > \pi_{1,t-1}$ and $\pi_{2,t-1} > \pi_{3,t-1}$. Then a positive steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{b}{R} \\ (p - Rp)(b - Rp) > (p - Rp)(gp - Rp) \\ (p - Rp)(b - Rp) > -Rp(p - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the third equation is not satisfied for $b > 0$ and $p \geq 0$.

(iii) $\pi_{3,t-1} > \pi_{1,t-1}$ and $\pi_{3,t-1} > \pi_{2,t-1}$. Then a positive steady state p must satisfy the following system of equations

$$\begin{cases} p = 0 \\ -Rp(p - Rp) > (p - Rp)(gp - Rp) \\ -Rp(p - Rp) > (p - Rp)(b - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the second and third equations are not satisfied.

(iv) $\pi_{1,t-1} = \pi_{2,t-1} > \pi_{3,t-1}$. Then a positive steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{g}{2R}p + \frac{b}{2R} \\ (p - Rp)(gp - Rp) = (p - Rp)(b - Rp) \\ (p - Rp)(gp - Rp) > -Rp(p - Rp) \\ (p - Rp)(b - Rp) > -Rp(p - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the fourth constrain implies $b < 0$.

(v) $\pi_{1,t-1} = \pi_{3,t-1} > \pi_{2,t-1}$. Then a positive steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{g}{2R}p \\ -Rp(p - Rp) = (p - Rp)(gp - Rp) \\ -Rp(p - Rp) > (p - Rp)(b - Rp) \\ (p - Rp)(gp - Rp) > (p - Rp)(b - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the first equation implies either $p = 0$ which con-

tradicts the third and fourth equation or $g = 2R$ which contradicts the second equation.

(vi) $\pi_{2,t-1} = \pi_{3,t-1} > \pi_{1,t-1}$. Then a non-negative steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{b}{2R} \\ -Rp(p - Rp) = (p - Rp)(b - Rp) \\ -Rp(p - Rp) > (p - Rp)(gp - Rp) \\ (p - Rp)(b - Rp) > (p - Rp)(gp - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the first equation contradicts the second.

(vii) $\pi_{1,t-1} = \pi_{2,t-1} = \pi_{3,t-1}$. Then a non-negative steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{g}{3R}p + \frac{b}{3R} \\ -Rp(p - Rp) = (p - Rp)(b - Rp) \\ -Rp(p - Rp) = (p - Rp)(gp - Rp) \\ (p - Rp)(b - Rp) = (p - Rp)(gp - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the first equation implies $p = \frac{b}{3R-g}$ and the second would imply $b = 0$.

2.A.3 Proof of Lemma 3

Steady states of the system must solve

$$Rp = n_1 gp + n_2 b + n_3 p,$$

where

$$n_1 = \frac{exp\{\beta\pi_1\}}{exp\{\beta\pi_1\} + exp\{\beta\pi_2\} + exp\{\beta\pi_3\}}, \quad n_2 = \frac{exp\{\beta\pi_2\}}{exp\{\beta\pi_1\} + exp\{\beta\pi_2\} + exp\{\beta\pi_3\}},$$

and

$$\pi_1 = (p - Rp)(gp - Rp), \quad \pi_2 = (p - Rp)(b - Rp), \quad \pi_3 = (p - Rp)(p - Rp).$$

For $\beta = 0$ then $n_1 = n_2 = n_3\frac{1}{3}$, which implies $p = \frac{b}{3R-g-1}$. The stability of the system can not be directly analyzed by the application of dynamical systems theory because of the 'forward looking' element p_{t+1}.

When $\beta = +\infty$, the price function is piecewise-defined as

$$p_t = \begin{cases} \frac{g}{R}p_{t-1} & \text{for} \quad \pi_{1,t-1} > \pi_{2,t-1} \quad \text{and} \quad \pi_{1,t-1} > \pi_{3,t-1} \\ \frac{b}{R} & \text{for} \quad \pi_{2,t-1} > \pi_{1,t-1} \quad \text{and} \quad \pi_{2,t-1} > \pi_{3,t-1} \\ \frac{p_{t+1}}{R} & \text{for} \quad \pi_{3,t-1} > \pi_{1,t-1} \quad \text{and} \quad \pi_{3,t-1} > \pi_{2,t-1} \\ \frac{g}{2R}p_{t-1} + \frac{b}{2R} & \text{for} \quad \pi_{1,t-1} = \pi_{2,t-1} > \pi_{3,t-1} \\ \frac{g}{2R}p_{t-1} + \frac{p_{t+1}}{2R} & \text{for} \quad \pi_{1,t-1} = \pi_{3,t-1} > \pi_{2,t-1} \\ \frac{b}{2R} + \frac{p_{t+1}}{2R} & \text{for} \quad \pi_{2,t-1} = \pi_{3,t-1} > \pi_{1,t-1} \\ \frac{g}{3R}p_{t-1} + \frac{b}{3R} + + \frac{p_{t+1}}{3R} & \text{for} \quad \pi_{1,t-1} = \pi_{2,t-1} = \pi_{3,t-1} \end{cases}$$

As before, we proceed by analysing the multiple cases:

(i) $\pi_{1,t-1} > \pi_{2,t-1}$ and $\pi_{1,t-1} > \pi_{3,t-1}$. Then a non-negative steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{g}{R}p \\ (p - Rp)(gp - Rp) > (p - Rp)(b - Rp) \\ (p - Rp)(gp - Rp) > (p - Rp)(p - Rp) \\ p \geq 0 \end{cases}$$

The system, however, has no solution since the first equation implies either $p = 0$ or $g = R$. In the first case, equations two and three are not satisfied. In the latter, the third is not.

(ii) $\pi_{2,t-1} > \pi_{1,t-1}$ and $\pi_{2,t-1} > \pi_{3,t-1}$. Then a non-negative steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{b}{R} \\ (p - Rp)(b - Rp) > (p - Rp)(gp - Rp) \\ (p - Rp)(b - Rp) > (p - Rp)(p - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the third equation is not satisfied for $b > 0$ and $p \geq 0$.

(iii) $\pi_{3,t-1} > \pi_{1,t-1}$ and $\pi_{3,t-1} > \pi_{2,t-1}$. Then a non-negative steady state p must

satisfy the following system of equations

$$\begin{cases} p = \frac{p}{R} \\ (p - Rp)(p - Rp) > (p - Rp)(gp - Rp) \\ (p - Rp)(p - Rp) > (p - Rp)(b - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the first equation implies either $p = 0$, which contradicts equations two and three, or $R = 1$.

(iv) $\pi_{1,t-1} = \pi_{2,t-1} > \pi_{3,t-1}$. Then a non-negative steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{g}{2R}p + \frac{b}{2R} \\ (p - Rp)(gp - Rp) = (p - Rp)(b - Rp) \\ (p - Rp)(gp - Rp) > (p - Rp)(p - Rp) \\ (p - Rp)(b - Rp) > (p - Rp)(p - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the first and second equations imply $g = R$, which then contradicts the third constraint.

(v) $\pi_{1,t-1} = \pi_{3,t-1} > \pi_{2,t-1}$. Then a non-negative steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{g}{2R}p + \frac{p}{2R} \\ (p - Rp)(p - Rp) = (p - Rp)(gp - Rp) \\ (p - Rp)(p - Rp) > (p - Rp)(b - Rp) \\ (p - Rp)(gp - Rp) > (p - Rp)(b - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the first equations imply either $p = 0$ which contradicts the third and fourth equation or $g = 2R - 1$ which cupled with the second equation then implies $g = R = 1$.

(vi) $\pi_{2,t-1} = \pi_{3,t-1} > \pi_{1,t-1}$. Then a positive steady state p must satisfy the following

system of equations

$$\begin{cases} p = \frac{b}{2R} + \frac{p}{2R} \\ (p - Rp)(p - Rp) = (p - Rp)(b - Rp) \\ (p - Rp)(p - Rp) > (p - Rp)(gp - Rp) \\ (p - Rp)(b - Rp) > (p - Rp)(gp - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since the first equation implies $p = \frac{b}{2R-1}$ and the second $p = b$ which is possible only for $R = 1$.

(vii) $\pi_{1,t-1} = \pi_{2,t-1} = \pi_{3,t-1}$. Then a positive steady state p must satisfy the following system of equations

$$\begin{cases} p = \frac{g}{3R}p + \frac{b}{3R} + \frac{p}{3R} \\ (p - Rp)(p - Rp) = (p - Rp)(b - Rp) \\ (p - Rp)(p - Rp) = (p - Rp)(gp - Rp) \\ (p - Rp)(b - Rp) = (p - Rp)(gp - Rp) \\ p \geq 0 \end{cases}$$

which has no solution since equations two to four imply either $p = 0$ which contradicts equation one or $p = b = gp$ implying $g = 1$ and $p = b$. Then the first equation would read $3Rb = 3b$ which implies $R = 1$.

2.A.4 Eigenvalues for the two type model

It is convenient to rewrite the model in (2.3.1) and (2.3.2) transforming it from a univariate third order difference equation to a first order difference equation with three states. We use the following change of variables $(p_t, p_{t-1}, p_{t-2}) = (x_{t+1}, w_{t+1}, z_{t+1})$ to rewrite the model as the following system:

$$\begin{cases} x_{t+1} = \frac{g}{R}x_t \left\{1 + exp[\beta(x_t - Rw_t)(b - gz_t)]\right\}^{-1} + \frac{b}{R}\left(1 - \left\{1 + exp[\beta(x_t - Rw_t)(b - gz_t)]\right\}^{-1}\right) \\ w_{t+1} = x_t \\ z_{t+1} = w_t \end{cases}$$

Taking first order derivatives, we get the following Jacobian matrix

$$\begin{bmatrix} A & B & C \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix},$$

with

$$A = \frac{\left(\left(\beta g^2 z_t - b\beta g\right) x_t - b\beta g z_t + g + b^2\beta\right) \mathrm{e}^{\beta\cdot(b-gz_t)(x_t-Rw_t)} + g}{R\cdot\left(\mathrm{e}^{\beta\cdot(b-gz_t)(x_t-Rw_t)}+1\right)^2},$$

$$B = -\frac{\beta\cdot\left(gx_t-b\right)\left(gz_t-b\right)\mathrm{e}^{\beta\cdot(b-gz_t)(x_t-Rw_t)}}{\left(\mathrm{e}^{\beta\cdot(b-gz_t)(x_t-Rw_t)}+1\right)^2},$$

$$C = \frac{\beta g\cdot\left(x_t-Rw_t\right)\left(gx_t-b\right)\mathrm{e}^{\beta\cdot(x_t-Rw_t)(b-gz_t)}}{R\cdot\left(\mathrm{e}^{\beta\cdot(x_t-Rw_t)(b-gz_t)}+1\right)^2}.$$

To evaluate the stability of the steady state, one can observe the eigenvalues of the Jacobian matrix, which in this case are given by:

$$\lambda_1 = \frac{\sqrt[3]{2A^3+3\sqrt{3}\sqrt{4A^3C-A^2B^2+18ABC-4B^3+27C^2}+9AB+27C}}{3\sqrt[3]{2}} -$$

$$-\frac{\sqrt[3]{2}\left(-A^2-3B\right)}{3\sqrt[3]{2A^3+3\sqrt{3}\sqrt{4A^3C-A^2B^2+18ABC-4B^3+27C^2}+9AB+27C}} + \frac{A}{3}$$

$$\lambda_2 = -\frac{1}{6\sqrt[3]{2}}(1-i\sqrt{3})\sqrt[3]{2A^3+3\sqrt{3}\sqrt{4A^3C-A^2B^2+18ABC-4B^3+27C^2}+9AB+27C}+$$

$$+\frac{\left(1+i\sqrt{3}\right)\left(-A^2-3B\right)}{3\cdot 2^{2/3}\sqrt[3]{2A^3+3\sqrt{3}\sqrt{4A^3C-A^2B^2+18ABC-4B^3+27C^2}+9AB+27C}} + \frac{A}{3}$$

$$\lambda_3 = -\frac{1}{6\sqrt[3]{2}}(1+i\sqrt{3})\sqrt[3]{2A^3+3\sqrt{3}\sqrt{4A^3C-A^2B^2+18ABC-4B^3+27C^2}+9AB+27C}+$$

$$+\frac{\left(1-i\sqrt{3}\right)\left(-A^2-3B\right)}{3\cdot 2^{2/3}\sqrt[3]{2A^3+3\sqrt{3}\sqrt{4A^3C-A^2B^2+18ABC-4B^3+27C^2}+9AB+27C}} + \frac{A}{3}$$

2.A.5 Eigenvalues for the model with fundamentalist

As before, we rewrite the model as the "following" system:

$$\begin{cases} x_{t+1} = \frac{g}{R} x_t n_{1,t} + \frac{b}{R} n_{2,t} \\ w_{t+1} = x_t \\ z_{t+1} = w_t \end{cases}$$

with

$$n_{1,t} = \frac{exp[\beta(x_t - Rw_t)(gz_t - Rw_t)]}{Z_t},$$

$$n_{2,t} = \frac{exp[\beta(x_t - Rw_t)(by_{t-3} - Rw_t)]}{Z_t},$$

$$n_{3,t} = \frac{exp[\beta(x_t - Rw_t)(-Rw_t)]}{Z_t},$$

$$Z_t = exp[\beta(x_t - Rw_t)(gz_t - Rw_t)] + exp[\beta(x_t - Rw_t)(by_{t-3} - Rw_t)] + exp[\beta(x_t - Rw_t)(-Rw_t)].$$

Taking first order derivatives, we get the following Jacobian matrix:

$$\begin{bmatrix} A & B & C \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$$

$$A = \frac{e^{R\beta w\cdot(x-Rw)} \cdot \left(g e^{2\beta\cdot(gz-Rw)(x-Rw)+R\beta w\cdot(x-Rw)} + e^{\beta\cdot(gz-Rw)(x-Rw)} \cdot \left(\left(\beta g^2 z - b\beta g\right) x - b\beta g z + g + b^2\beta\right) e^{\beta\cdot(b-Rw)(x-Rw)+R\beta w\cdot(x-Rw)} + \beta g^2 z x + g\right) + b^2\beta e^{\beta\cdot(b-Rw)(x-Rw)}\right)}{R \cdot \left(e^{\beta\cdot(gz-Rw)(x-Rw)+R\beta w\cdot(x-Rw)} + e^{\beta\cdot(b-Rw)(x-Rw)+R\beta w\cdot(x-Rw)} + 1\right)^2},$$

$$B = -\frac{\beta e^{R\beta w\cdot(x-Rw)} \cdot \left(e^{\beta\cdot(x-Rw)(gz-Rw)} \cdot \left(\left(\left(g^2 x - bg\right) z - bgx + b^2\right) e^{\beta\cdot(b-Rw)(x-Rw)+R\beta w\cdot(x-Rw)} + g^2 xz\right) + b^2 e^{\beta\cdot(b-Rw)(x-Rw)}\right)}{\left(e^{\beta\cdot(x-Rw)(gz-Rw)+R\beta w\cdot(x-Rw)} + e^{\beta\cdot(b-Rw)(x-Rw)+R\beta w\cdot(x-Rw)} + 1\right)^2},$$

$$C = \frac{\beta g \cdot (x - Rw)\left(\left(gx - b\right) e^{\beta\cdot(b-Rw)(x-Rw)+R\beta w\cdot(x-Rw)} + gx\right) e^{\beta\cdot(x-Rw)(gz-Rw)+R\beta w\cdot(x-Rw)}}{R \cdot \left(e^{\beta\cdot(x-Rw)(gz-Rw)+R\beta w\cdot(x-Rw)} + e^{\beta\cdot(b-Rw)(x-Rw)+R\beta w\cdot(x-Rw)} + 1\right)^2}.$$

The general form of the eigenvalues computed in Section (2.A.4) is still valid, so that we can obtain them by simply replacing the formulae for A, B and C.

Appendix 2.B Steady States

In Figure 2.14, we numerically solve for the implicit function defining the steady state of the systems. To obtain a plot in three dimension we fix the value of the parameters $R = 1.01$ and $b = 1$ and vary the parameters g and β. We can see that by fixing β or g, the steady state is monotonic in the other parameter. Therefore, in the region highlighted, there is always a single positive steady state. However, we showed analytically, that

while in the two-types market the steady state remains positive for increasing values of β, in the model with fundamentalists and rational agents the steady state becomes negative as β goes to infinity.

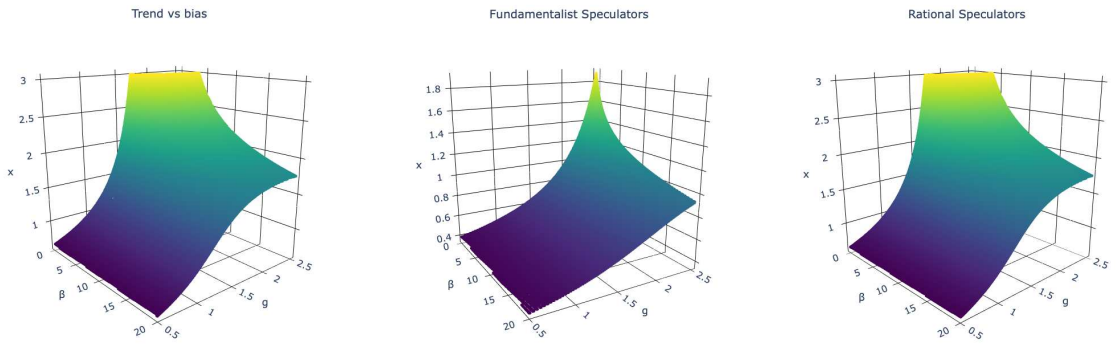

Figure 2.14: Steady states of the models, for different values of g and β. Other parameters are $R = 1.01$ and $b = 1$.

Appendix 2.C Neural Network

The network architecture we use is the following.

Hidden Layers	Hidden Layer size	Activation	Loss	Optimizer	epochs
1	16	tanh	mse	adam	1000

An overview of the Neural Network is given below, while we refer to Goodfellow, Bengio and Courville (2016) for full details on this kind of model.

Input Layer The input layer passes the input data to the next layer. Denote the input as $\mathbf{X} \in \mathbb{R}^k$, where k is the size of the input vector.

Hidden Layer The first dense layer transforms the input $\mathbf{X}$ using a weight matrix $\mathbf{W}_1 \in \mathbb{R}^{k \times 16}$, a bias vector $\mathbf{b}_1 \in \mathbb{R}^{16}$, and applies the tanh activation function. The operation can be described by the equation:

$$\mathbf{H} = \tanh(\mathbf{W}_1^\top \mathbf{X} + \mathbf{b}_1)$$

Here, $\mathbf{H} \in \mathbb{R}^{16}$ is the output of the hidden layer, which becomes the input to the next layer.

Output Layer The output layer further transforms the output of the hidden layer $\mathbf{H}$ using another weight matrix $\mathbf{W}_2 \in \mathbb{R}^{16 \times 1}$, a bias term $b_2 \in \mathbb{R}$, and applies a linear

activation function. The operation for the output layer is given by:

$$\mathbf{Y} = \mathbf{W}_2^\top \mathbf{H} + b_2$$

$\mathbf{Y} \in \mathbb{R}$ is the scalar output of the network that approximates the state variable. The input vector in Section (2.4.2) is of size $k = 3$ and contains p_{t-1}, p_{t-2} and p_{t-3} while the target is p_{t+1}. In Section (2.5) the size is equal to 9 and the input vector contains $x_{t-1}, x_{t-2}, x_{t-3}, w_{t-1}, w_{t-2}, w_{t-3}, R_t, R_{t-1}, R_{t-2}$, that is lags of the Bitcoin in deviation form, of the BiTSI index and of the time-varying risk-free rate.

Appendix 2.D Stationarity tests

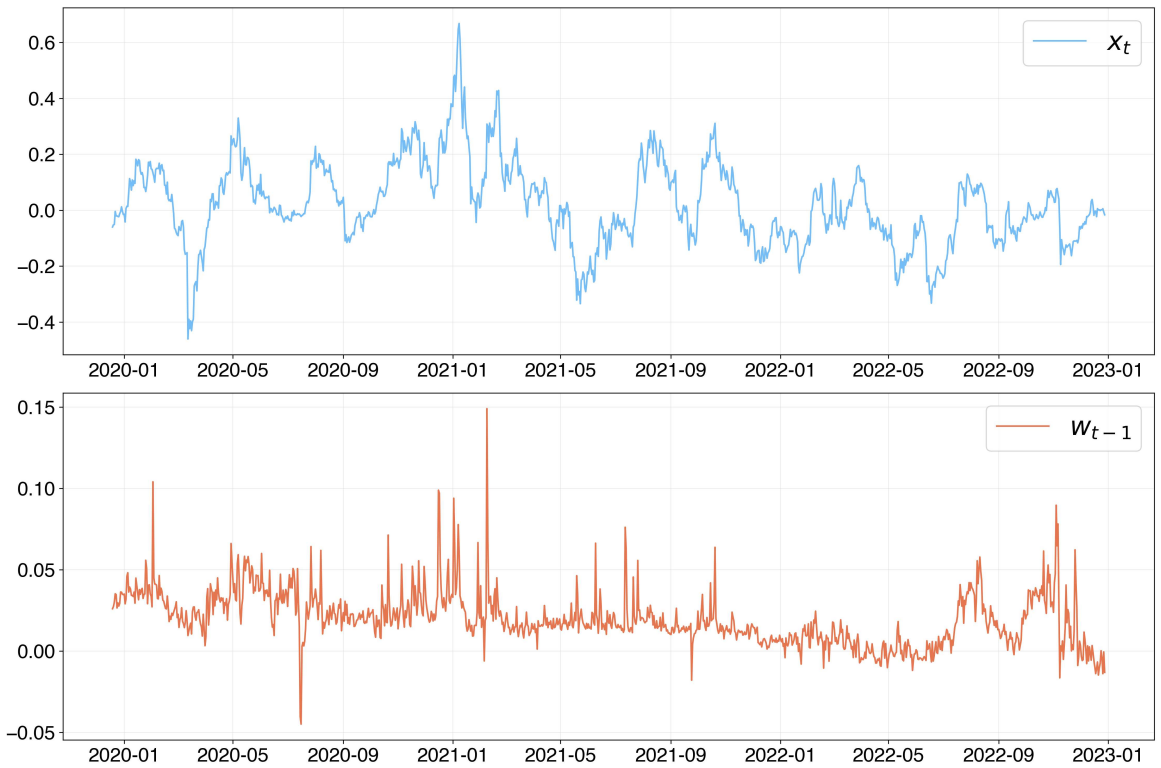

Figure 2.15: Variable for the estimation

We test the presence of a unit root for the two time series related to the percentage deviation form the moving average fundamental by means of the augmented Dickey–Fuller test (ADF). The associated p-values are: 1.27e-06 for x_t and 0.037 for w_{t-1}. We can, therefore, reject the null hypothesis of the unit root of the ADF.

Appendix 2.E Robustness to different windows

In this section, we check the robustness of the estimation to a different choice of the window length in the moving average fundamental value. We repeat the analysis for each value between 1 and 99. A window choice of 1 implies that we are actually estimating the model on daily returns. For each choice we re-estimate the two-types model with trend chasers and bias and plot the point estimate of the parameters with associated standard errors in the top panel of Figure 2.16, the dashed line is at 0. The bottom panel shows the p-value for the F-statistic of significance of the non-linear model with respect to the linear one and the adjusted r-squared. The dashed line marks the 5% significance level.

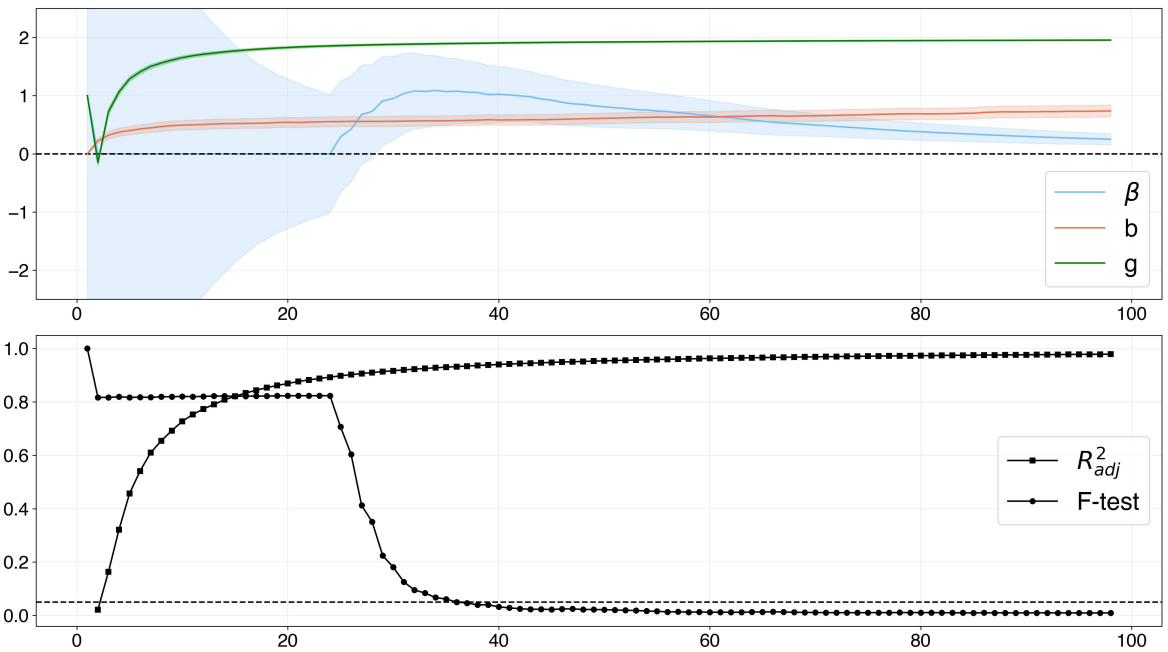

Figure 2.16: Effect of different windows

Daily returns (window length = 1) are, as one would expect, unpredictable, with an adjusted r-squared close to 0. For a window length smaller than 25, the intensity of choice β is not significantly different from 0. For values lower than 7, we get a confidence interval so large that we had to impose limits on the axis in order to obtain a meaningful visualization. For values higher than 30, however, we observe results that are qualitatively similar, with high R-squared, positive and significant βs and preference for the non-linear model as conveyed by the p-value of the F-test.

Appendix 2.F Bootstrapped standard errors

In this section, we show the robustness of the standard errors by computing them via bootstrap. The p-value of the test for Conditional Heteroscedasticity does not reject homoskedasticity in the residuals. Nonetheless, we follow the approach of Goncalves and Kilian (2004). That is for 2000 bootstrap replications we generate a series of pseudo residuals by multiplying the original residuals by a random number drawn from a Standard Normal Distribution We then create a pseudo time series using the associated non-linear model by replacing the actual residuals with the pseudo residuals. Finally we re-estimate the model using the pseudo time series to obtain a new set of parameter estimates and report the values of the associated confidence interval at the .99 percent level in Table 2.4.

Table 2.4: 99% bootstrapped confidence intervals

	g	b	β
Trend chasers vs pure bias:	[1.87 1.94]	[0.41 0.76]	[0.03 2.01]
Tc vs B vs Fundamentalist:	[2.81 2.9]	[0.61 1.14]	[0.03 1.0]
Tc vs B vs Rational Expectations:	[1.87 1.96]	[0.58 1.1]	[0.12 2.85]

Chapter 3

(Mis)information diffusion and the financial market

3.1 Introduction

Financial markets exhibit empirical regularities that are challenging to capture by relying on the canonical assumption of a representative agent endowed with Full Information Rational Expectation (FIRE) (Muth, 1961). Asset returns display skewed distribution with fat tails and while the Efficient Market Hypothesis (Fama, 1970) suggests that information is immediately reflected in prices, empirical evidence seems to point in the direction of some frictions in its incorporation (Huberman and Regev, 2001; Vozlyublennaia, 2014).

A natural starting point in addressing these findings is that of releasing the assumption of a representative agent, and a large body of literature has done so by incorporating agents' heterogeneity in otherwise standard asset pricing models. There are mainly two ways in which this heterogeneity can be modeled. The first is assuming that not all agents have rational expectations. Theoretical works in this setting have shown how *boundedly rational* agents can survive in the market and how their interaction with rational agents can lead to complex dynamics (Brock and Hommes, 1997, 1998; Chiarella, 1992; Lux, 1998).

The second relaxation focused on the full information part of the assumption and has been the focus of early research on information acquisition and processing in financial markets (Barberis, Shleifer and Vishny, 1998; Grossman and Stiglitz, 1980; Kyle, 1985). While this literature has focused on the effect of information frictions, it has not fully explored the way in which information is diffused in the market.

In this chapter, we construct an Agent-Based Model (ABM) of financial agents con-

nected through a social network to study the effect of information and misinformation diffusion on asset prices. ABMs have been widely used in finance and economics, since they allow modeling complex interactions among agents and capturing emergent phenomena that are difficult to predict from individual behavior alone (Axtell and Farmer, 2023; Dieci and He, 2018). We introduce parsimonious relaxations to the FIRE assumption in two dimensions: heterogeneous access to information and delayed transmission of news. Only a fraction of agents has access to direct information about the fundamental value of the asset, while the rest of the network can only receive information from their peers. We then introduce a new social learning mechanism that allows agents to update their beliefs in a Bayesian way while incorporating a behavioral component. These modifications, based on empirical evidence, allow us to replicate stylized facts of financial markets without resorting to stronger forms of bounded rationality such as zero-intelligence or purely backward-looking agents. Conditional on their heterogeneous information sets, agents are perfectly forward-looking. To discipline the model, we calibrate it to match moments of the empirical return distribution of TESLA stocks. We use the calibrated model to offer an explanation of the delayed absorption of shocks in the market and to investigate the effect of (mis)information diffusion, in different network topologies, on market efficiency.

3.1.1 Related literature and Contribution

Our chapter merges insights from two main literatures: one focusing on information diffusion in social networks and one on the effect of information frictions on financial markets. In the former, since the seminal work of Degroot (1974) on consensus reaching, scholars have proposed multiple mechanisms of belief updating. Gale and Kariv (2003) are among the first to introduce a network component in the social learning literature. The presence of an explicit structure leads to the question of whether agents should consider it when making decisions. DeMarzo, Vayanos and Zwiebel (2003) posit the notion that the information an agent obtains from the network may exhibit bias, since other agents within the network might derive their beliefs from the agent's own beliefs. After presenting our model, we argue that in our work this problem is not present. Moreover, we assume that agents do not think strategically about the network structure. Acemoglu, Ozdaglar and ParandehGheibi (2010) introduce a model addressing (mis)information diffusion in network structures where updates occur bilaterally, while Acemoglu and Ozdaglar (2011) and Kanoria and Tamuz (2013) both focus on Bayesian learning. Buechel, Hellmann and Klößner (2015) emphasizes the role of conformity and centrality in shaping the collective wisdom of the network. Con-

versely, Rusinowska and Taalaibekova (2019) propose a model in which individuals strategically aim to exert influence on others. In our context, strategic considerations are absent, and communication is truthful, which differentiates our work from previous literature that considered the effect of access to private information to manipulate markets Benabou and Laroque (1992).

We propose a mechanism of belief updating when receiving multiple sources of information simultaneously. The mechanism is quite general and nests other social learning models, like the one of Degroot (1974). It relies on Bayesian updating with a behavioral component, since agents construct a time-varying measure of precision for each source of information and use it to derive the posterior mean and variance of the signal of interest. The second object is particularly important in our case, as agents are risk-averse and will use this variance to compute their optimal demands.

Information frictions have been documented empirically by multiple sources. Huberman and Regev (2001) show that the stock prices of a company, CASI Pharmaceuticals, did not incorporate new information for five months. They point out that the news was initially released as a research article in the journal Nature, but investors reacted only when a Wall Street Journal article reposted the findings of the original study. Behavioral factors can also contribute to information delay. Dellavigna and Pollet (2009) provide evidence of limited attention, demonstrating reduced investor responsiveness on Fridays and identifying profitable strategies that exploit such underreactions. This delay has been studied theoretically, with a focus on insider trading (Benabou and Laroque, 1992; Collin-Dufresne and Fos, 2016; Kyle, 1985) or herding behavior in information acquisition (Banerjee, 1992; Cont and Bouchaud, 2000; Orléan, 1995).

Finally, a number of previous studies have examined the role of network dynamics in financial markets. Most of this work has used networks to model imitation among agents, such as in Iori (2002). Panchenko, Gerasymchuk and Pavlov (2013) build on the Brock and Hommes (1998) framework, allowing agents to choose among various trading strategies observed within their network. Khashanah and Alsulaiman (2016) develop an ABM in which agents share their optimal holdings with their neighbors. Wu, He and Li (2018) also explore different network topologies within an ABM of financial markets, focusing on how traders switch between fundamentalist and chartist strategies. In contrast, Biondo (2020) examines various network structures where agents imitate the discrete trading decisions of their peers, while Bertella, Silva, Correa and Sornette (2021) considers scenarios in which agents compare their wealth to that of their neighbors, leading to asset reallocation when their performance lags behind. The novelty of

our model lies in the fact that agents do not imitate each other. As already remarked, all agents are forward-looking, and the network is used only to model information flow.

The rest of the chapter is structured as follows: Section 3.2 introduces the model, focusing on the two blocks that constitutes the ABM, the financial market and the information diffusion process. Section 3.3 describes the process we use to obtain a realistic calibration of the model. Section 3.4 presents properties of the model by means of numerical simulations. Section 3.5 concludes.

3.2 Model

3.2.1 The financial market

Consider an economy with I consumers, indexed by $i = (1, 2, \ldots, \text{I})$. Consumers are infinitely lived, and at the beginning of every period, they receive the same endowment W_0. They have the same utility over the end of period wealth, given by

$$U(W_t) = -e^{-aW_t}, \tag{3.2.1}$$

where $a > 0$ is the coefficient of risk aversion. In order to transfer wealth from the beginning of the period to the end, there are two types of security: a risk-free and a risky asset. We define p_t as the price of the risky asset at a generic time t and normalize the price of the risk-free asset to 1. In every period consumers decide how to allocate their initial endowment, choosing between the two possible securities. At the end of the period they receive profits based on their portfolio, and immediately consume their wealth. Given their participation in the financial market, throughout the chapter we use interchangeably the terms consumers, investor and agent. Defining $X_{i,t}$ as the consumer's demand for the risky asset and $M_{i,t}$ as the demand for the safe asset, the allocation choice is subject to the budget constraint

$$p_t X_{i,t} + M_{i,t} = W_0. \tag{3.2.2}$$

The risk-free rate is $R > 1$ and the risky asset pays a stochastic payoff which is equal to a dividend claim plus the future price of the asset

$$y_{t+1} = p_{t+1} + d_{t+1}. \tag{3.2.3}$$

The presence of the future price on the right-hand side of equation (3.2.3) ensures a positive feedback mechanism of expectations. This is a well-established feature of

financial markets, as documented among others by Heemeijer, Hommes, Sonnemans and Tuinstra (2009). There are two stochastic components determining the realization of future dividends

$$d_{t+1} = d + \theta_{t+1} + \varepsilon_{t+1}, \tag{3.2.4}$$

with $\varepsilon_{t+1} \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ being pure unobservable noise. The stochastic component θ_{t+1} follows a stationary[23] AR(1) process, $\theta_{t+1} = \beta\theta_t + \eta_{t+1}$ with $\eta_{t+1} \sim \mathcal{N}(0, \sigma_\eta^2)$ and $\beta \in (0, 1)$ and is observable by some agents before its actual realisation. One can think of it as information, in a similar spirit to Grossman and Stiglitz (1980) and Gerotto, Pellizzari and Tolotti (2019). This implies that the underlying fundamental value of the risky asset is stochastic but that some information about it is revealed in advance. However, this information is not immediately incorporated in the asset price. The end of period t wealth for the i^{th} consumer is given by

$$W_{i,t} = RM_{i,t} + y_{t+1}X_{i,t} = R(W_0 - p_t X_{i,t}) + y_{t+1}X_{i,t}. \tag{3.2.5}$$

Agents optimize their end of period wealth, which given the normality of y_{t+1} is also normal. The optimization problem is therefore given by

$$\max_{\{X_{i,t}\}} \left(-exp\left\{ -a\tilde{\mathbb{E}}_{i,t}(W_{i,t}) + \frac{a^2}{2}\tilde{\mathbb{V}}_{i,t}(W_{i,t}) \right\} \right). \tag{3.2.6}$$

Using equation (3.2.5) and solving for the optimal choice of risky asset yields

$$X_{i,t} = \frac{\tilde{\mathbb{E}}_{i,t}(y_{t+1}) - Rp_t}{a\tilde{\mathbb{V}}_{i,t}(y_{t+1})}. \tag{3.2.7}$$

The notation $\tilde{\mathbb{E}}_{i,t}$ and $\tilde{\mathbb{V}}_{i,t}$ is used to represent subjective expected value and subjective variance for agent i. It is short notation $\mathbb{E}(\cdot|\mathcal{I}_{i,t})$ and $\mathbb{V}(\cdot|\mathcal{I}_{i,t})$, where $\mathcal{I}_{i,t}$ is the information set of agent i at time t, that is before the realization of θ_{t+1}. We set net supply of outside share of the risky asset equal to 0 and use the market clearing condition, imposing supply equal to aggregate demand, to obtain an implicit pricing equation

$$\sum_{i=1}^{I} X_{i,t} = \sum_{i=1}^{I} \left(\frac{\tilde{\mathbb{E}}_{i,t}(y_{t+1}) - Rp_t}{\tilde{\mathbb{V}}_{i,t}(y_{t+1})} \right) = 0. \tag{3.2.8}$$

[23] This is a simplifying assumption since it would imply stationary prices. A more accurate representation could be given by considering a growth model for dividends like in Diks and Dindo (2008) and having agents forecast the growth rate. However, while more realistic, this assumption would not change the main insight of the chapter, nor affect the calibration which is based on the model returns.

All agents in the model are assumed to be forward-looking in evaluating the asset price. They expect the asset price to be its fundamental value, which is determined by the present discounted value of the stream of future dividends. This implies that in the model there is only a minimal deviation from rationality, given by the asymmetry in information. This is a crucial aspect and one of the main features that distinguish our chapter from other prominent works in the literature. Chiarella (1992), Brock and Hommes (1998) and Lux (1998) among the others, focus on the coexistence of (rational) fundamentalists and some type of boundedly rational backward-looking agent. The most common type of bounded rationality for the latter is associated with technical trading rules and the literature has usually labeled them as chartists or trend followers (Anufriev, Gardini and Radi, 2020; Day and Huang, 1990; Gardini et al., 2025; Tramontana, Westerhoff and Gardini, 2010, 2013). This interplay has been shown to generate complex dynamics, and in some cases, even chaotic behavior. A third type of agents which are sometimes included in this framework are sentiment followers (Di Francesco and Hommes, 2024; Gardini et al., 2022) and their presence can contribute to destibilize the market even further. By contrast, in our model, all agents act optimally given the information available and if frictions were removed from the information diffusion process, we would end back in the rational expectations setting. When solving the forward-looking problem, agents in the model assume that other agents behave identically. This assumption leads them to solve the problem as if they were the representative agent. In other words, they consider the aggregate behavior of all agents to be equivalent to their own individual behavior. With this assumption the pricing equation for each individual is given by

$$p_t = R^{-1}\tilde{\mathbb{E}}_t\left(y_{t+1}\right), \tag{3.2.9}$$

in which we omit the subscript i and solving by iterating forward, which is done in section 3.A of the Appendix gives

$$p_t = \frac{d}{r} + \frac{\tilde{\mathbb{E}}_t(\theta_{t+1})}{R-\beta}. \tag{3.2.10}$$

The first component of equation (3.2.10) is the usual discounted value of future expected dividends: without information, the price would be constant for all time periods. The second component is specific of our model and imposed by the presence of the observable component of dividend θ. The next section is devoted to describe the mechanism for which agents receive information about this component.

3.2.2 Actual and Perceived Law of Motion

We now describe the individual beliefs, by distinguishing between the actual law of motion (ALM) of the observable component of dividends θ_{t+1} and the perceived law of motion (PLM) which agents believe to be true. We note that the actual process θ_{t+1} is not publicly available, and it can not be directly observed even after its realization. A subset of agents, whom we call informed and misinformed, have access to a private signal which they believe to be the true process. They have full confidence in their source of information and as soon as they receive this private signal, they update their beliefs to it. This behavior can be sustained by the fact that they only observe the realization of payoffs. Errors in their payoff forecasts can be attributed to the idiosyncratic noise ε_{t+1} and not necessarily to their source of information. The remaining agents, which we call uninformed, do not have access to the private signal and update their beliefs in a Bayesian way, which we specify below. The ALM is given by the AR(1) process

$$\theta_{t+1} = \beta\theta_t + \eta_{t+1}. \tag{3.2.11}$$

This implies that the conditional distribution of θ_{t+1} given the information set $\mathcal{I}_t$ is

$$\theta_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\beta\theta_t, \sigma_\eta^2\right). \tag{3.2.12}$$

Since we assume that all agents are aware of the structure of the process, we model all prior beliefs as normal distributions. As anticipated we assign agents to three possible categories.

Informed agents. These agents receive the true information θ_{t+1} and base their forecast on it. One can think that this is due to agents having access to privileged or inside information. We prefer to associate this choice with empirical evidence provided by Huberman and Regev (2001) and Peng and Xiong (2006) supporting the idea of different classes of investors with different access to information. Some agents may possess knowledge to process domain-specific information that other, generalist agents, may lack. Their PLM coincides with the ALM. For them we have

$$\theta_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\beta\theta_t, \sigma_\eta^2\right). \tag{3.2.13}$$

Misinformed agents. These agents think they have perfect foresight like informed agents, but are actually basing their forecast on misinformation. Their PLM is given by

$$\theta_{t+1} = \beta\gamma_t + \sigma_\nu := \gamma_{t+1}, \tag{3.2.14}$$

that is they assume the process follows an AR(1), with same persistence parameter β but different noise term $\nu_{t+1} \sim \mathcal{N}(0, \sigma_\nu^2)$. This can be thought of as a misinformation shock. Their conditional distribution is then given by

$$\theta_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\beta\gamma_t, \sigma_\nu^2\right). \tag{3.2.15}$$

One can think of these traders as noise traders in De Long et al. (1990) style or akin to sentiment followers. One could naturally ask why these individuals are trading on noise. As Black (1986) puts it, *"One reason is that they like to do it. Another is that there is so much noise around that they don't know they are trading on noise. They think they are trading on information."* The latter explanation is what we argue can sustain this behavior in our model. Although these agents are not evaluating their source of information, they might be unable to detect their bias in the presence of noise.

Uninformed agents. These agents do not have any private signal, but their PLM coincides with the ALM. The only difference is that they condition their beliefs on their previous period expectation,

$$\theta_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\beta\tilde{\mathbb{E}}_{t-1}(\theta_t), \sigma_\eta^2\right). \tag{3.2.16}$$

The literature provides ample evidence that not all information is immediately processed by investors upon release. The simplest explanation is that of limited attention. Hirshleifer, Lim and Teoh (2009) finds that investor reactions to earnings announcement are weaker on days when there are multiple simultaneous news releases. On a similar note Dellavigna and Pollet (2009) show that investors take more time to process news on Friday. Tetlock (2011) shows that investors overreact to stale information and Gilbert, Kogan, Lochstoer and Ozyildirim (2012) demonstrate that this causes mispracing in the market by constructing a profitable trading strategy exploiting this finding. More recently Blankespoor, deHaan and Marinovic (2020) show that there seems to exist some "disclosure processing costs" that would make disclosures not public information, as usually defined, but a form of private information.

3.2.3 Information diffusion

Agents are socially connected and are organized in a network. Each agent represents a node, and nodes are entirely characterized by beliefs regarding the observable component of dividends θ_{t+1}. The network is static. All edges are exogenously determined and time-invariant. The edges represent the flow of information between nodes. At the beginning of time t agents have normal prior distribution according to their category. They then receive new data in two ways. Informed and misinformed agents observe θ_{t+1} and γ_{t+1} respectively. Moreover they observe them without noise, so that their posterior beliefs are immediately updated to the realizations of the variables[24]

$$\theta_{t+1}|\theta_{t+1} \sim \mathcal{N}(\theta_{t+1}, 0) \quad \gamma_{t+1}|\gamma_{t+1} \sim \mathcal{N}(\gamma_{t+1}, 0). \tag{3.2.17}$$

Uninformed agents do not receive any private information, but they can receive information from their peers. Formally, for agent i, we assume that they can observe node j prior mean if an edge exists between node i and node j. Based on this information they update their beliefs in a Bayesian way but incorporating a behavioral component. When an agent is connected to another node in the network, we assume that they construct an implicit variance evaluating the forecasting error of the node over time. To do so we compute an exponential moving average of the forecasting error, given as the squared difference between the last observable payoff and the payoff prediction ($\mu_{j,t-1}$) implied by source j, that is

$$EMA_{j,t} = w\left(y_{t-1} - \left(\frac{dR}{r} + \mu_{j,t-1}\right)\right)^2 + (1-w)EMA_{j,t-1}. \tag{3.2.18}$$

Then, to map this to a comparable variance of the given source, we multiply the ratio between source j exponential moving average and their own, with the prior variance

$$\sigma^2_{j,t} = \sigma^2_\eta \frac{EMA_{j,t}}{EMA_{i,t}}. \tag{3.2.19}$$

In simple terms, if they observe that node j has been more accurate then themselves they attach higher confidence to the information received by that node. They then update their beliefs by applying the Bayes rule to the case of receiving information from K different sources and given in the following proposition.

[24] This is in a sense an abuse of notation as a Normal distribution with 0 variance is a degenerate distribution with support at the single point θ_{t+1}, known as the Dirac's delta function. However this notation is convenient since it allow us to model the evolution of the beliefs of this category of agents in the general framework.

Proposition 4 (Bayesian Updating of Beliefs)**.** *Assume agents have a normal prior distribution with parameters* (μ_0, σ_0^2) *and receive* K *new information, each with a Normal likelihood* (μ_k, σ_k^2)*,* $k = 1, 2 \dots K$*. Then agents posterior distribution is Normal, with posterior mean given by:*

$$\mu_P = \frac{\sum_{k=0}^{K} \left(\mu_k \cdot [A]_k^{\bar{A}-1} \right)}{\sum [A]^K}, \tag{3.2.20}$$

and posterior variance:

$$\sigma_P^2 = \frac{\prod_{j=0}^{K} \sigma_j^2}{\sum [A]^{\bar{A}-1}}, \tag{3.2.21}$$

where, $A = \{\sigma_0^2, \sigma_1^2, \sigma_2^2, \dots \sigma_K^2\}$*,* $\bar{A}$ *is the cardinality of set* A*.* $[A]^J$ *is the set of all distinct combinations of products of size* J *from set* A*,* $\sum [A]^J$ *sums over all the elements in the set,* $[A]_k^J$ *indicates the combination that does not include* σ_k^2*.*

Proof. See appendix (3.D.1).

In every period t agent i uses this mechanism to update their beliefs and obtain their posterior distribution $\theta_{t+1} \sim \mathcal{N}(\mu_{P,t}, \sigma_{P,t}^2)$. To build intuition consider the simplest case in which an agent is connected only to one other node. Then the posterior parameters are given by

$$\mu_P = \frac{\mu_0 \sigma_1^2 + \mu_1 \sigma_0^2}{\sigma_0^2 + \sigma_1^2}, \quad \sigma_P^2 = \frac{\sigma_0^2 \sigma_1^2}{\sigma_0^2 + \sigma_1^2}. \tag{3.2.22}$$

Notice that this updating procedure is equivalent to using the Kalman filter to filter out the noise in the signal, as we show in appendix 3.B. If one then moves to $K = 2$, hence considering an agent with two connections, posterior parameters are given by

$$\mu_P = \frac{\mu_0 \sigma_1^2 \sigma_2^2 + \mu_1 \sigma_0^2 \sigma_2^2 + \mu_2 \sigma_0^2 \sigma_1^2}{\sigma_1^2 \sigma_2^2 + \sigma_0^2 \sigma_2^2 + \sigma_0^2 \sigma_1^2}, \quad \sigma_P^2 = \frac{\sigma_0^2 \sigma_1^2 \sigma_2^2}{\sigma_1^2 \sigma_2^2 + \sigma_0^2 \sigma_2^2 + \sigma_0^2 \sigma_1^2}. \tag{3.2.23}$$

The posterior mean can be seen as a weighted average of the prior means the agent has access to, with weight given by $[A]_j^{\bar{A}-1} / \sum [A]^{\bar{A}-1}$. The noisier the alternatives are, the more an agent will rely on a particular source. This mechanism can be seen as a particular case of the naive updating proposed in Golub and Jackson (2010) but the novelty of our approach is that we simultaneously derive the posterior variance. This object while of no relevance in most information diffusion works, has a crucial role in our work, given the risk averse behavior of our agents. Also different in our case is that the information exchange happens only one time per time step, therefore avoiding any potential bias given by repeated information as is the case in DeMarzo et al.

(2003). Morevoer the only source of heterogeneity is given by the different categories of agents, but within the same class, beliefs are ex-ante homogenous. This implies that there are only two innovations or shocks entering the network at each time step. The combination of this property with the autoregressive structure of the component agents are interested in, make so that aggregating over correlated information is unavoidable but not deleterious. Of particular interest are then the following situations.

1. An agent considers source k to be absolutely certain. Then in our model we have

$$\sigma_k^2 = 0 \implies \mu_p = \mu_k, \tag{3.2.24}$$

$$\sigma_k^2 = 0 \implies \sigma_p^2 = 0. \tag{3.2.25}$$

   That is, when an agent uses only one source of information and is totally confident in it, the posterior mean is equal to the signal, with variance 0.

2. An agent completely disregards source k, then

$$\lim_{\sigma_k^2 \to +\infty} \mu_p = \frac{\sum_{k=0}^{K} \left(\mu_k \cdot [B]_k^J \right)}{\sum [B]^{\bar{B}-1}}, \tag{3.2.26}$$

$$\lim_{\sigma_k^2 \to +\infty} \sigma_p^2 = \frac{\prod_{j=0}^{K} \sigma_j^2}{\sum [B]^{\bar{B}-1}}. \tag{3.2.27}$$

   where $B = A \setminus \sigma_k^2$. When agents believe that a source of information is totally unreliable, their posterior mean and variance is equal to the one they would get by omitting the source of information.

3.2.4 Payoffs and prices

With the posterior beliefs regarding the observable component of the dividend θ_{t+1}, we can compute beliefs regarding future payoffs. First, the subjective expectations are $\tilde{\mathbb{E}}_t(\theta_{t+1}) = \theta_{t+1}$ for informed agents, $\tilde{\mathbb{E}}_t(\theta_{t+1}) = \gamma_{t+1}$ for misinformed agents, and $\tilde{\mathbb{E}}_t(\theta_{t+1}) = \mu_{P,t}$ for uninformed agents. Using the expression for p_t allows us to compute, (see Appendix 3.C), the conditional expectation:

$$\tilde{\mathbb{E}}_t(y_{t+1}) = \frac{dR}{r} + \frac{R\tilde{\mathbb{E}}_t(\theta_{t+1})}{R - \beta}, \tag{3.2.28}$$

and the conditional variance:

$$\tilde{\mathbb{V}}_t(y_{t+1}) = \sigma_\varepsilon^2 + \tilde{\mathbb{V}}_t(\theta_{t+1}) + \frac{\tilde{\mathbb{V}}_t(\tilde{\mathbb{E}}_{t+1}(\theta_{t+2}))}{(R-\beta)^2} + 2\mathbb{C}\tilde{\mathbb{O}}\mathbb{V}_t\left(\theta_{t+1}, \frac{\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})}{R-\beta}\right), \quad (3.2.29)$$

which clarifies that heterogeneity in beliefs is completely described by θ_{t+1}. While the expected value of the stochastic payoff is given for every category by equation (3.2.28), the conditional variance is specific for each category. In particular, starting from equation (3.2.29), we have:

Informed agents

$$\tilde{\mathbb{V}}_t(y_{t+1}) = \sigma_\varepsilon^2 + \frac{\sigma_\eta^2}{(R-\beta)^2}, \quad (3.2.30)$$

since for these agents, θ_{t+1} is not a random variable at time t, and

$$\tilde{\mathbb{V}}_t(\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})) = \tilde{\mathbb{V}}_t(\theta_{t+2}) = \sigma_\eta^2. \quad (3.2.31)$$

Misinformed agents

Similarly, for misinformed agents, we have:

$$\tilde{\mathbb{V}}_t(y_{t+1}) = \sigma_\varepsilon^2 + \frac{\sigma_\nu^2}{(R-\beta)^2}. \quad (3.2.32)$$

Uninformed agents

Deriving the conditional variance for uninformed agents is more challenging. In particular, the term $\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})$ depends on the network structure and the updating mechanism we have described. This makes the future forecast a random variable, whose distribution depends on the entire network topology, making the derivation analytically intractable. However, since we are deriving the conditional variance of payoffs under the assumption that agents operate as representative agents, we have:

$$\tilde{\mathbb{E}}_{t+1}(\theta_{t+2}) = \beta\tilde{\mathbb{E}}_t(\theta_{t+1}), \quad (3.2.33)$$

which is known at time t and, therefore, not a random variable. Under this assumption, we have:

$$\tilde{\mathbb{V}}_t(y_{t+1}) = \sigma_\varepsilon^2 + \tilde{\mathbb{V}}_t(\theta_{t+1}). \quad (3.2.34)$$

3.2.5 The role of social learning

In this section, we highlight the importance of our social learning model by comparing the full model to two simpler cases that are naturally nested within our framework.

The first case features traders' heterogeneity but no social learning. In this framework, agents keep their expectations constant at their prior beliefs. Defining I as the total number of agents, and λ and ξ as the proportions of informed and misinformed agents respectively, and imposing a prior mean of 0 for uninformed agents, the resulting price evolves according to:

$$p_t = \frac{d}{r} + \frac{\lambda\theta_{t+1} + \xi\gamma_{t+1}}{R - \beta}. \tag{3.2.35}$$

The second case modifies this baseline by adding perhaps the most notable social learning mechanism, that of Degroot (1974), which occurs when uninformed agents assign equal weight to all sources of information available. For this case, we consider a fully connected network, allowing us to derive analytical results. In this framework, every uninformed agent receives λI signals from informed agents, ξI signals from misinformed agents, and $(1-\lambda-\xi)I$ signals from uninformed agents (including themselves).

Moreover, every uninformed agent has identical beliefs. In particular, assuming that the homogeneous variance attached to every piece of information is $\sigma^2 > 0$, we have for uninformed agents:

$$\tilde{\mathbb{V}}_t(\theta_{t+1}) = \frac{\sigma^2}{I}. \tag{3.2.36}$$

It is worth noting that the posterior variance approaches 0 as the network size approaches infinity. More precisely, the driving factor here is the number of signals each agent receives, which equals the total number of agents[25]. This result is general, as formalized in the following proposition:

Proposition 5 (Vanishing of Posterior Variance)**.** *Assuming that the variance associated with each source of information is positive and bounded, then:*

$$\lim_{K\to\infty} \tilde{\mathbb{V}}_t(\theta_{t+1}) = 0. \tag{3.2.37}$$

Proof. See Appendix 3.D.2.

This proposition implies that uninformed agents become increasingly confident in their posterior beliefs as they receive more signals, eventually acting as if their expectations are certain. This feature results in a higher quantity demanded in the market compared to the baseline case, leading to increased volume and more volatile returns.

[25]Since their own beliefs are included.

However, the symmetry in the updating process ensures that the return distribution remains symmetric.

The posterior mean for uninformed agents is given by:

$$\tilde{\mathbb{E}}_t(\theta_{t+1}) = \lambda\theta_{t+1} + \xi\gamma_{t+1} + (1-\lambda-\xi)\tilde{\mathbb{E}}_{t-1}(\theta_t). \tag{3.2.38}$$

The resulting pricing equation is then:

$$\begin{aligned} p_t = \frac{d}{r} + \left(\lambda \frac{\frac{\theta_{t+1}}{R-\beta}}{V_I} + \xi \frac{\frac{\gamma_{t+1}}{R-\beta}}{V_M} + (1-\lambda-\xi) \frac{\frac{\tilde{\mathbb{E}}_t(\theta_{t+1})}{R-\beta}}{V_U} \right) \\ \left(\frac{\lambda}{V_I} + \frac{\xi}{V_M} + \frac{(1-\lambda-\xi)}{V_U} \right)^{-1}, \end{aligned} \tag{3.2.39}$$

with:

$$V_I := \sigma_\varepsilon^2 + \frac{\sigma_\eta^2}{(R-\beta)^2}, \quad V_M := \sigma_\varepsilon^2 + \frac{\sigma_\nu^2}{(R-\beta)^2}, \quad V_U := \sigma_\varepsilon^2 + \frac{\sigma_\eta^2}{I}.$$

Rearranging, this reads as:

$$p_t = \frac{d}{r} + \frac{1}{R-\beta} \frac{\lambda\left(\theta_{t+1}\right) V_M V_U + \xi(\gamma_{t+1}) V_I V_U + (1-\lambda-\xi)\tilde{\mathbb{E}}_t(\theta_{t+1}) V_I V_M}{\lambda V_M V_U + \xi V_I V_U + (1-\lambda-\xi) V_I V_M}. \tag{3.2.40}$$

We now compare the two models to the full model, in which agents are connected in a small-world network. To ensure comparability, we simulate the three models for 30 different realizations of the stochastic processes using different seeds and show the distribution of the normalized model-implied returns in Figure 3.1. The parameters used are the same for all models and are fixed to the values in Table 3.1.

We compare the return distributions for all models and plot a quantile-quantile (QQ) plot for each return series. Both the baseline and the DeGroot model exhibit symmetric return distributions and no excess kurtosis. The introduction of our social learning mechanism is not only more realistic but also necessary to match empirical evidence.

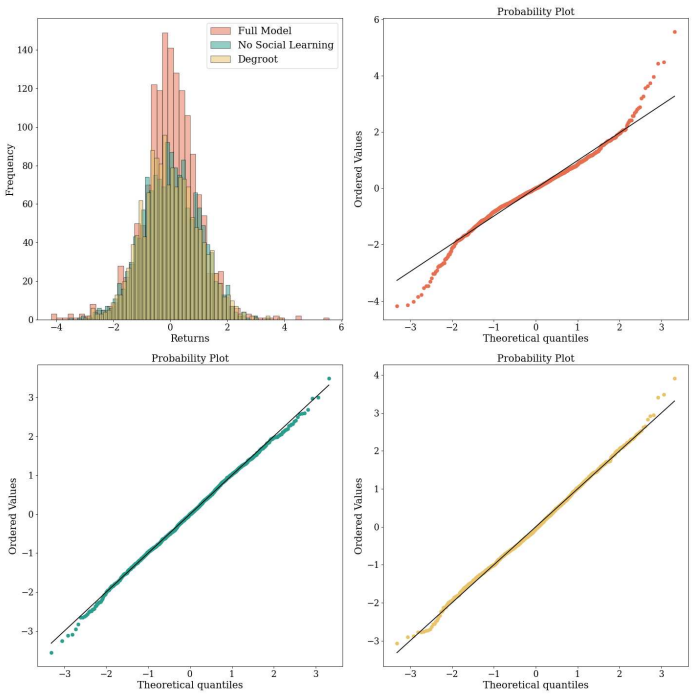

Figure 3.1: Return distributions for the three models.

In the next section, we turn to the calibration of the full model and the exploration of its properties.

3.3 Empirical Calibration

First, we summarize the model by offering a more accesible visualization of the sequence of events taking place in each time step t

1. Agents update their prior beliefs regarding the observable component of dividends θ_{t+1}, according to their own category;

2. Agents compute the average forecast error over payoffs of the nodes they are connected to. The last payoff considered by them in the computation is y_{t-1};

3. Agents use the updating mechanism to obtain their posterior beliefs regarding the observable component of dividends θ_{t+1};

4. The implied expected payoff y_{t+1} is computed by each agent;

5. The resulting price p_t and individual demands are computed.

Then we discipline the model and its many parameters by calibrating it with the following strategy. We target the daily returns of the TESLA stock from September 2023 to September 2024. In the literature a common choice to obtain a reliable calibration is to use some index, typically the S&P 500 index (Schmitt and Westerhoff, 2017). In this context we argue that it is more natural to work with a single stock and in particular with one which features a high percentage of retail investors participation. In TESLA case the ownerhisp attributed to the public is around 40%. Moreover individuals and insiders make up 12% of the total ownership, thus making the stock an ideal candiate to study the effects of information frictions. We then deal with the calibration of the model by following a three-step procedure.

First, some parameters are specific to the network topology. The one we use as our benchmark is the Watts and Strogatz (1998) Small World Network. The network is generated by starting with a regular lattice of a given degree, assigning each node to one of the three categories introduced in Section 3.2.3 according to specified proportions. Finally we rewire a fraction of edges randomly by a given probability of rewiring, introducing long-distance connections in the network. The main features of such a topology are local clustering, short-average path length and almost homogeneous degree of connection across nodes.

Second we keep some parameters constant at values which are arbitrarily chosen after exprimentation with the model, as they do not affect the model dynamics. Total time steps are sufficient to ensure that behaviors driven by initial conditions are absorbed in the long run. The number of agents in the model is of relative importance only when combined with network specific parameters. The gross risk free rate is the daily equivalent corresponding to the average risk free rate in the period September 2023-September 2024 as proxied by the Market Yeald on U.S. Treasury Securities at 1-year Constant Maturity. The constant component of dividends[26] is calibrated to match the daily closing price at the beginning of the period: $d = p_0 \cdot r$. The coefficient of constant risk aversion is at a value common in the literature, see for example Chetty (2006).

Third, for some parameters we explore their effect on the model for a wide range of values by Sobol sensitivity analysis Sobol (2001), Saltelli (2002), Saltelli (2010) implemented by using the the SALib python library (Herman and Usher, 2017). Sobol sensitivity analysis is a global variance-based method used to quantify the contribution of individual input variables, or in our case parameters, to the output variance of a

[26] Although TESLA does not currently pay dividends, one could interpret this as an expected average dividend based on expected earnings and an expected payout ratio. Mathematically this value has the only effect of shifting prices and has no impact on the dynamics.

model. Given a model of the form $Y = f(X_1, X_2, \ldots, X_k)$, where Y is a scalar output, the first-order effect of a factor X_i can be expressed as

$$V_{X_i}\left(\mathbb{E}_{X_{\sim i}}(Y|X_i)\right), \tag{3.3.1}$$

where X_i represents the i-th factor, and $X_{\sim i}$ denotes the matrix of all factors except X_i. The inner expectation operator $\mathbb{E}_{X_{\sim i}}(Y|X_i)$ represents the mean of Y over all possible values of $X_{\sim i}$ while keeping X_i fixed. The outer variance operator is taken over all possible values of X_i. The associated sensitivity measure, known as the first-order sensitivity index, is then given by:

$$S_i = \frac{V_{X_i}\left(\mathbb{E}_{X_{\sim i}}(Y|X_i)\right)}{V(Y)}. \tag{3.3.2}$$

S_i is normalized, as $V_{X_i}\left(\mathbb{E}_{X_{\sim i}}(Y|X_i)\right)$ ranges between zero and $V(Y)$. $V_{X_i}\left(\mathbb{E}_{X_{\sim i}}(Y|X_i)\right)$ measures the first-order (additive) effect of X_i on the model output, while $\mathbb{E}_{X_i}\left(V_{X_{\sim i}}(Y|X_i)\right)$ is customarily referred to as the residual.

Another commonly used variance-based sensitivity measure is the total effect index, defined as:

$$S_{T_i} = \frac{\mathbb{E}_{X_{\sim i}}\left(V_{X_i}(Y|X_{\sim i})\right)}{V(Y)} = 1 - \frac{V_{X_{\sim i}}\left(\mathbb{E}_{X_i}(Y|X_{\sim i})\right)}{V(Y)}. \tag{3.3.3}$$

The total effect index S_{T_i} measures the total contribution of X_i, which includes both the first-order and higher-order effects (interactions) of the factor. One way to interpret this is by considering that $V_{X_{\sim i}}\left(\mathbb{E}_{X_i}(Y|X_{\sim i})\right)$ represents the first-order effect of $X_{\sim i}$, so that $V(Y)$ minus $V_{X_{\sim i}}\left(\mathbb{E}_{X_i}(Y|X_{\sim i})\right)$ must give the contribution of all terms in the variance decomposition that include X_i.

In Figure 3.2 we report the total order Sobol index for the seven parameters in the y-axis, computed on three outputs of the model, namely: price variance in excess of a model with a representative fully informed investor which we use as benchmark, skewness and kurtosis of the model returns. All results are obtained from using 128 samples for each of the 9 parameters of interest, with 30 different seeds per combination, resulting in a total of 34560 simulations.

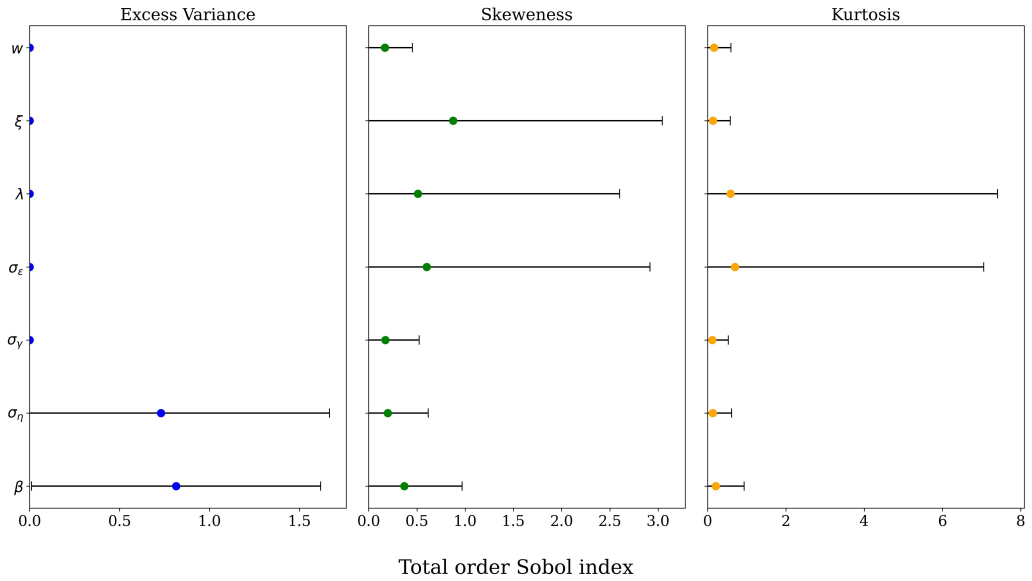

Figure 3.2: Sobol Sensititivty Analysis

Excess variance seems to be driven by two parameters, β and σ_η, which combine to determine the unconditional variance of the observable component of dividends. The skewness of the model returns is driven by the proportions of informed and misinformed agents, λ and ξ, and by the size of the unobservable noise component of dividends σ_ε. A similar result is obtained for the kurtosis of the model returns, although now the proportion of misinformed agents seems to play a smaller role. The number of informed and misinformed agents in the model however has both a direct and indirect effect. The former is due to the fact that these agents are also actively operating in the market and therefore their beliefs are directly reflected in the price. The latter is due to the effect that their communication has on the beliefs of uninformed agents, which in turn affect the price. Since we are mainly interested in this last effect we decide to keep the proportion of informed and misinformed agents at a constant value. The value we chose is 5% of the population for both classed, in order to have a reasonable number of agents in each category while still being able to attribute the majority of the price dynamics to the uninformed agents. The volatilities of the unobservable and observable component of dividends combine to determine the noise-to-signal ratio of the information. We therefore normalize the value of σ_ε to 1. Finally, as already mentioned, β and σ_η jointly determines the variance of the information. We therefore fix one of them, namely $\beta = 0.5$ and use the other to vary the variance of the information. This leaves us with two parameters to calibrate σ_η and σ_ν and we do so by using two moments of the TESLA stock returns: skewness and kurtosis. For this we use the Sequential Neural Posterior Estimation (SNPE) proposed by Papamakarios and Murray (2016)

and popularised to economics ABMs by the recent work of Dyer, Cannon, Farmer and Schmon (2024). SNPE is a likelihood-free inference method that uses neural networks to directly approximate the posterior distribution of model parameters.[27] The key idea behind this method is to simulate data from the model at different parameter values and use a neural network to learn the conditional density $q_\phi(\theta|x)$, where θ are the parameters and x is the observed data (e.g., TESLA stock skewness and kurtosis).

The process works as follows:

1. **Simulate Data**: For a given set of parameters θ_n, simulate data x_n from the ABM.

2. **Neural Network Training**: Use the simulated parameter-data pairs (θ_n, x_n) to train a Mixture Density Network (MDN). The MDN outputs a conditional probability distribution $q_\phi(\theta|x)$, which serves as an approximation to the true posterior $p(\theta|x)$.

3. **Posterior Approximation**: After training, the neural network provides an approximation of the posterior distribution for any observed data x_o, by maximizing the likelihood
$$\max_{\phi} \frac{1}{N} \sum_{n=1}^{N} \log q_\phi(\theta_n|x_n). \tag{3.3.4}$$

4. **Sequential Refinement**: The method iteratively refines the prior distribution $p(\theta)$ based on the learned posterior, improving the efficiency of the simulations by focusing on plausible regions of the parameter space.

We refer the reader to Dyer et al. (2024) for an in depth evaluation of the method. While other methods such the Simulated method of moments as in Franke and Westerhoff (2012) and Franke and Westerhoff (2016) could be used, this method is particularly efficient as it requires only enough simulations to train the neural network. Moreover we assess the ability of the method to recover the true posterior by using a synthetic dataset generated by the model in appendix 3.E, which serves to show that in our setting the neural network approximation is accurate.

For the real data estimation we take an agnostic position on the prior. Given that we are estimating variances we use a uniform prior distribution, with support $[0.1, 2]$ and report the results as well as the other parameter values in table 3.1. The posterior distribution is extremely well behaved and uni-modal as we can report in Figure 3.3,

[27] This is in contrast with classical Bayesian methods in which the posterior is computed as the product of the likelihood and the prior.

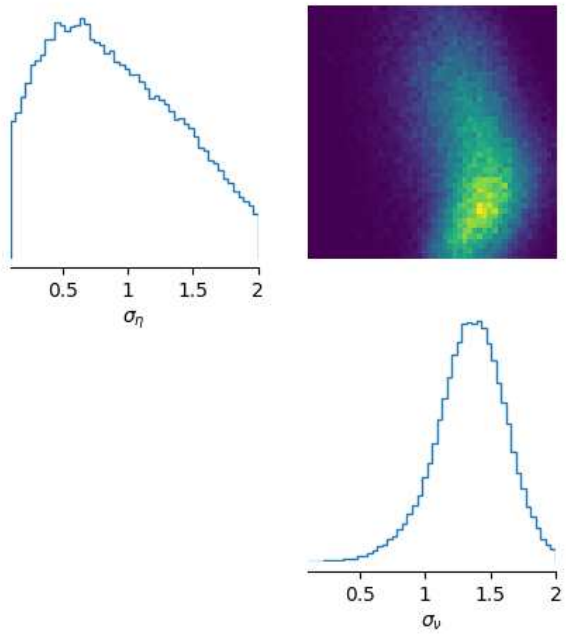

Figure 3.3: Posterior Distribution

with median values $\sigma_\eta = 0.84$ and $\sigma_\nu = 1.36$. It is interesting to notice that shocks to misinformation are slightly lager than shocks to information. This is consistent with misinformation being more sensational than actual information, in order to generate more attention.

Table 3.1: Parameter Calibration

Calibrated Parameters				
Parameter	Range Explored	Value	Source	Description
T	-	1500	Own Calibration	Total Time Steps
I	-	200	Own Calibration	Number of Agents
R	-	1,0001	Average daily risk free rate 09:2023 - 09:2024	Gross Risk Free Rate
d	-	0,021	Directly calibrated	Constant Component of Dividends
a	-	1	(Chetty, 2006)	Coefficient of constant risk aversion
σ_ε	$[0.0001, 2]$	1	Normalized	Std of the Unobservable Component of Dividends
w	$[0.1, 0.9]$	0.9	Own Calibration	Memory parameter of the Exponential Moving Average
λ	$[0.01, 0.3]$	0.05	Own Calibration	Proportion of Informed Agents
ξ	$[0.01, 0.3]$	0.05	Own Calibration	Proportion of Misinformed Agents

Estimated Parameters					
Parameter	Prior	Posterior			Description
		Mean	Median	Std	
σ_η	Uniform (0,2)	0.89	0.84	0.47	Std of Information shocks
σ_ν	Uniform (0,2)	1.34	1.36	0.26	Std of Misinformation shocks

3.4 Numerical Simulations

With our calibration we can then show key features of the model. We do so in Figure 3.4 by using the average over 30 Montecarlo simulations with different stochastic seeds. In panel (a) we plot the network structure used in the simulations. Agents position is randomly determined in the beginning on the experiment and then kept fixed, in order to allow us to track each agent evolution in the different simulations. Panel (b) shows a scatter plot of cumulative profits and average forecast error for each agent. The accuracy of uninformed agents lies between that of the other two categories. In this configuration, misinformation does not appear to spread significantly into the network. Although there is a generally linear positive relationship between accuracy and profits, there are some notable exceptions. While most uninformed investors incur small gains or even losses, a small subset earns large profits, sometimes almost as high as those of informed agents. The reason for this pattern is that these agents are directly linked to a source of information. They receive information relatively early and act on before others, allowing them to profit the most from the delayed price adjustment caused by the gradual spread of information in the network. Moreover as already pointed out, uninformed agents with many connections might have a lower posterior variance than informed agents, therefore taking larger positions and profiting more from the price movement in the direction of their beliefs. Lastly we look at the impact on returns. In panel (c) we use a QQ plot to identify the presence of fat tails. We compare the simulated quantiles with the theoretical quantiles from a Normal distribution with the same mean and standard deviation of model's returns. In table (3.4) we report the summary statistics of the model returns. Skewness and kurtosis are the target moments we have used for the calibration and are 0.04 and 5.50 in the simulated data and 0.25[28] and 5.54 in the empirical counterpart. We also report the profit (or loss) incurred by the misinformed agents in the model. In this case it is negative, showcasing that while misinformation is spreading in the network it does not diffuse outside of the local cluster of connections to misinformed agents.

[28] While normally one would expect negative skewness in financial markets, the period we considered was associated with a series of favorable earnings reports and positive growth expectations.

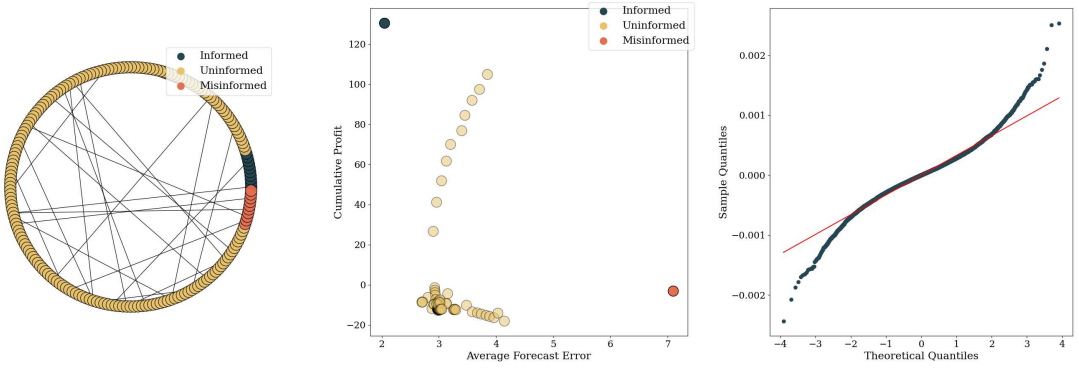

(a) Network Structure (b) Cumulative Profits (c) QQ Plot

Figure 3.4: Small World Network

Note: Data are obtained from 30 simulations with different stochastic seeds. **Network parameters** are: Network density = 0.1, Probability of rewiring = 0.01.

We can then analyze the implications of the model for the absorption of information in the market using an impulse response function (IRF) analysis. In our framework, shocks propagate through two distinct channels. First, there is a *direct channel* as the shock directly influences expectations, affecting investors' demands and ultimately feeding into market prices. Second, there is an *indirect channel* as the resulting prices and communication between agents feed back into how uninformed traders update their weighting matrices, which determine the relative importance they assign to different sources of information.

If we were to introduce a shock to the θ process at time τ, we could not attribute the resulting effects only to this specific shock because earlier shocks would still influence both the network's structure and the ongoing diffusion of information. To isolate the causal impact of a single shock, we proceed as follows. We begin by simulating the model for a sufficiently long period to ensure that the weighting matrix converges to a stable representative state. We then record this matrix and use it to initialize a new market scenario in which no further updates to the weighting matrix occur. By suppressing all noise until time $\tau = 10$, agents have no reason to alter their weightings, which therefore remain fixed at the initial distribution. Under these controlled conditions, we introduce a one-standard-deviation shock to either θ (an information shock) or γ (a misinformation shock). This setup allows us to compute the resulting impulse response and entirely capture the direct effect of the shock. Figure 3.5 shows the impulse response of the price to a one-standard-deviation shock.

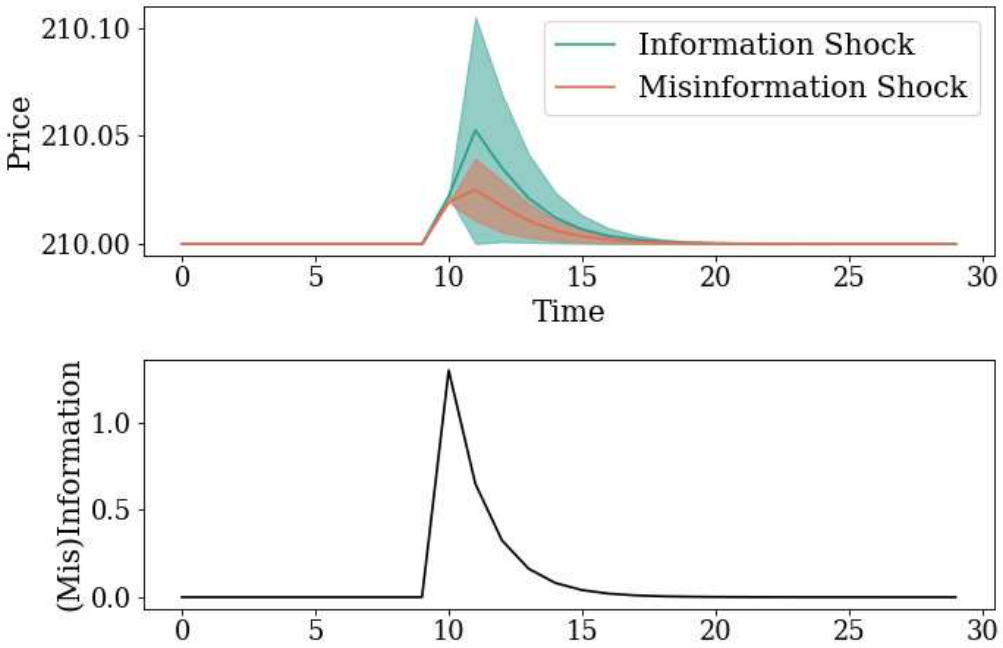

Figure 3.5: Impulse Response To a (Mis)Information Shock at time $\tau = 10$

The results align with patterns observed in empirical studies. While misinformation shocks do propagate into market prices, their overall impact is weaker than that of information shocks. In line with the findings of Clarke, Chen, Du and Hu (2021), the market appears able to filter out some of the noise generated by misinformation. By contrast, information shocks not only exert a stronger influence on prices but also continue to shape price dynamics over time. In both cases, the delayed peak of the IRF is consistent with the assumption that information diffuses gradually through the network and only reaches its full impact on prices once uninformed agents have processed it (Huberman and Regev, 2001).

3.4.1 The effect of different Network Topologies

The use of social networks has drastically reshaped communication in recent years. The network topology plays a crucial role in enabling rapid information sharing but also raises concerns about how beliefs are formed. In recent years, social media platforms have shifted from a friend-based structure, where connections are reciprocal, to a follower-based structure, where a few influencer nodes can have a disproportionately large number of connections. Moreover, the way many algorithms are designed to maximize user engagement can lead to the creation of echo chambers, where individuals are exposed only to information that confirms their existing beliefs. While the small-world

network structure we have used so far is a good approximation of the friend-based structure, we now explore the impact of polarization and follower-based structures on the model dynamics by adopting different network topologies. Crucially, to isolate the effect of the network structure, we repeat the analysis while keeping the variable parameters fixed.

Stochastic Block Network

The first scenario we analyze is a completely polarized society. We create it by using a Stochastic Block Model Holland, Laskey and Leinhardt (1983), in which we partition the nodes in order to create two clusters. Informed and misinformed agents are separated and assigned to either the information block or the misinformation one. We denote by density of intra-groups edges, the likelihood of having an edge between agents belonging to the same cluster. This is higher than the density of inter-groups edges, regulating connections between agents belonging to different blocks. Thus not only the network is partitioned, but communication between the two groups is scarce. We report the network specific parameters in table (3.2).

Table 3.2: Parameters of the Stochastic Block Network

Parameter	**Value**
Number of Partitions	2
Density of Intra-Groups Edges	0.1
Density of Inter-Groups Edges	0.001

As before we simulate the model 30 times with different stochastic seeds. Results are displayed in Figure (3.6).

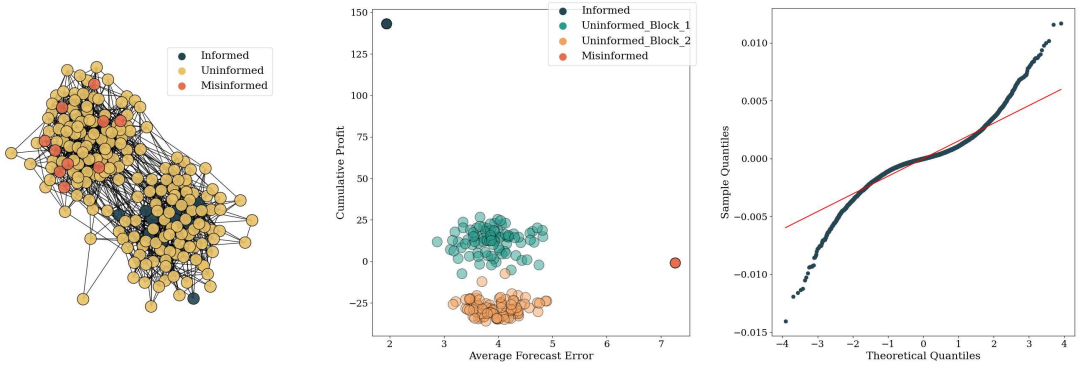

(a) Network Structure (b) Cumulative Profits (c) QQ Plot

Figure 3.6: Stochastic Block Network

Note: Data are obtained from 30 simulations with different stochastic seeds. **Network parameters** are: Number of Partitions = 2, Density of Intra-Groups Edges = 0.1, Density of Inter-Groups Edges = 0.001.

Panel (a) illustrates the network structure, which now consists of two clusters each exhibiting small-world architectures. Under these conditions, uninformed agents are less likely of receiving reliable information and misinformation simultaneously. The consequences of this segregation are clearly visible in Panel (b), where agents in the information block earn higher profits than those in the misinformation block. Panel (c) shows that, although the overall shape of the returns distribution remains qualitatively similar, the leptokurtosis is now more pronounced. Misinformed agents continue to incur losses, though these losses are smaller than in the baseline scenario. While these agents can influence their own cluster, their erroneous beliefs fail to fully penetrate the market's pricing mechanism. As a result, uninformed agents situated in the informed cluster capitalize on their disadvantaged counterparts in the misinformed segment of the network.

Scale Free Network Until now we have simulated societies in which agents have equal opportunities of sharing their beliefs, given that the average degree was rather homogeneous across the network. For our next scenarios we opt to work with societies in which certain individuals can have a disproportionate impact on the network[29]. This concept is similar to the one of "guru", discussed in Tedeschi, Iori and Gallegati (2012). Gurus are are agents that are most imitated by others and can emerge endogenously in the market. The main difference is that in their model, edges are created by a

[29]There are multiple example of investment platforms building on this feature. A famous one is the multi-asset investment plaform eToro, in whcih users can directly copy the porfolio allocation of *top-performing investors*.

mechanism of preferential attachment based on wealth. Instead we use an exogenous mechanism so by creating a directed[30] Scale-Free Network (Bollobás, Borgs, Chayes and Riordan, 2003) and forcefully allocating either informed or misinformed agents in the nodes with most outward connections. The network specific parameters are reported in table (3.3)

Table 3.3: Parameters of the Scale Free Network

Parameter	**Range**
Alpha	0.41
Beta	0.54
Gamma	0.05

Note: Alpha is the probability for adding a new node connected to an existing node chosen randomly according to the in-degree distribution. Beta is the probability for adding an edge between two existing nodes. Gamma is the robability for adding a new node connected to an existing node chosen randomly according to the out-degree distribution. For a detailed explanation of the parameters role we refer to page 2 of Benabou and Laroque (1992).

We begin by analyzing the case where informed agents are the most central. The network topology is displayed in panel (a) of Figure (3.7). Panel (b) highlights that informed agents benefit significantly from their prominent position in the network, earning much more than in previous scenarios. This is because they can directly reach most of the uninformed traders. Since communication is lagged, informed agents push the price in the direction in which they have already taken a position. This results in the majority of uninformed agents incurring small losses. As everyone receives the information almost simultaneously, there are no additional gains to be made; in fact, uninformed agents lose money on average to informed agents who acted earlier based on their private information. Even so, it is still better for uninformed agents to incorporate stale news into their forecasts, given the persistence of these shocks. This result offers a potential explanation for the empirical findings regarding stale information reported by Gilbert et al. (2012) and Tetlock (2011). The effect of this configuration on returns is a distribution with extremely low kurtosis and skewness, as shown in table (3.4) and panel (c). The market is as efficient as possible, given the inability of uninformed agents to access private information. As a result, there is almost no room for extremely large returns or extreme losses.

[30]It must be remarked that although until now we were using undirected networks, informed and misinformed agents were behaving in a dogmatic fashion. This is because a 0 prior variance in their beliefs implies non updating or posterior beliefs exactly equal to the prior.

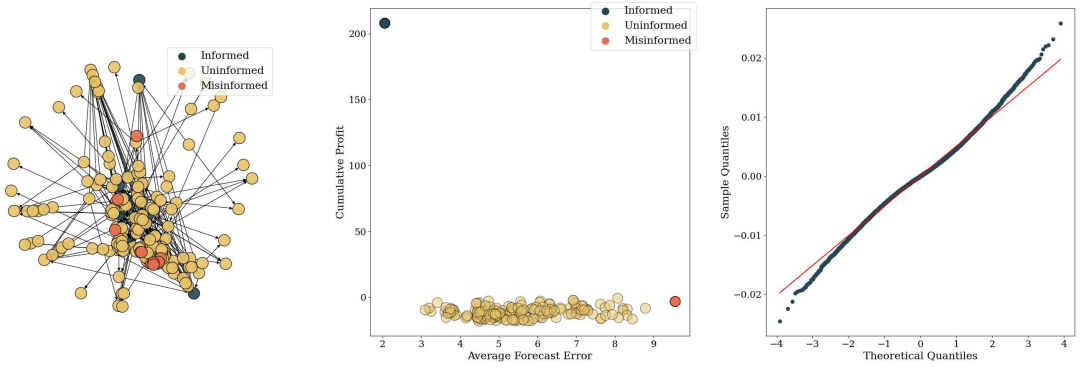

(a) Network Structure (b) Cumulative Profits (c) QQ Plot

Figure 3.7: Barabasi-Albert Graph, Informed

Note: Data are obtained from 30 simulations with different stochastic seeds. **Network parameters** are: Alpha = 0.41, Beta = 0.54, Gamma = 0.05.

We then analyze the second scale-free society, in which the most connected nodes are misinformed agents. The network in panel (a) of Figure (3.8) has the same configuration as the previous one, with the only difference being the position of agents. Turning to the profit and accuracy analysis in panel (b), we can see that misinformation has successfully spread throughout the network. Misinformed agents now achieve positive and significantly high profits. This comes at the expense of uninformed agents, who are both losing more on average and showing a high degree of inequality in their profits. Having, in most cases, connections only to misinformed individuals or other uninformed agents with stale misinformation makes their participation in the market extremely unfruitful. Panel (c) confirms that returns are again extremely lepto-kurtic, as reported in table (3.4), and the skewness is now negative. This market is the least efficient of all the scenarios we have analyzed. However, some uninformed agents are still able to make profits. This is because the network structure allows certain agents to profit from future agents incorporating their beliefs, even though these beliefs are based on misinformation. This mechanism is extremely realistic and typical of pyramid or Ponzi scheme structures, in which, even if one is aware of participating in such a scheme, profits can still be made as long as enough new participants enter in the future.

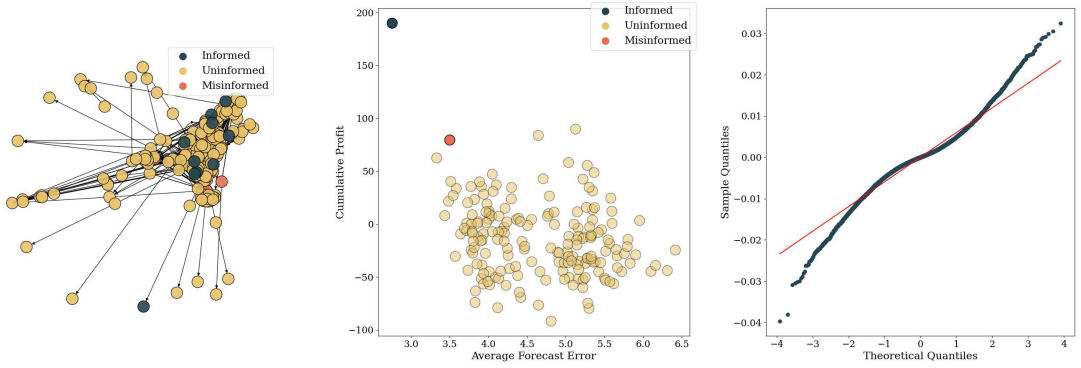

(a) Network Structure (b) Cumulative Profits (c) QQ Plot

Figure 3.8: Barabasi-Albert Graph, Misnformed

Note: Data are obtained from 30 simulations with different stochastic seeds. **Network parameters** are: Alpha = 0.41, Beta = 0.54, Gamma = 0.05.

Table 3.4: Summary of moments for different scenarios

	Profit of Misinformed	**Skewness**	**Kurtosis**
Small World	-4.30	0.04	5.50
Stochastic Block Network	-3.17	0.08	7.85
Scale Free Informed	0.29	0.04	3.77
Scale Free Misinformed	85.58	-0.01	5.44

3.5 Conclusion

We have presented an Agent-Based Model of a financial market to study the interplay between information diffusion and market prices and returns. In this setting agents are connected in a social network and can obtain information from their peers in order to form more accurate forecasts of the underlying dividend process. We proposed a novel mechanism of expectation formation when agents have to evaluate multiple sources of news simultaneously. This is based on Bayesian updating and provides an alternative to the full information rational expectations assumption, while imposing minimum departures from it. By means of numerical simulations we examined the efficiency implications of multiple social network structures. Leptokurtosis of returns and wealth inequality are exacerbated in societies in which misinformed agents occupy prominent positions.

Our approach in this chapter is purely positive: we do not address the reasons why certain network topologies emerge, yet this work yields valuable insights for both re-

tail investors and regulators. Our results indicate that in societies where misinformed agents hold prominent positions in terms of communication power, relying solely on past forecast accuracy to filter out misinformation may be insufficient. Although total welfare remains unchanged across different topologies, given the zero-sum nature of this simplified market, certain configurations may still be preferable, particularly those that mitigate wealth inequality. In this context, policymakers might consider promoting social media structures that prioritize friend-based connections rather than follower-based ones. Likewise, retail investors should try to assess the topology of the network in which they are embedded before deciding whether to incorporate a source of information in their investment strategy.

Appendix 3.A Forward Looking Price

$$\begin{aligned}
p_t &= R^{-1}\tilde{\mathbb{E}}_t\left(p_{t+1}+d_{t+1}\right)\\
&= R^{-1}\left[\tilde{\mathbb{E}}_t\left(p_{t+1}\right)+\tilde{\mathbb{E}}_t\left(d_{t+1}\right)\right]\\
&= R^{-1}\left[\tilde{\mathbb{E}}_t\left(R^{-1}\tilde{\mathbb{E}}_{t+1}\left(p_{t+2}+d_{t+2}\right)\right)+\tilde{\mathbb{E}}_t\left(d_{t+1}\right)\right]\\
&= R^{-1}\tilde{\mathbb{E}}_t(d_{t+1})+R^{-2}\tilde{\mathbb{E}}_t(d_{t+2})+R^{-2}\tilde{\mathbb{E}}_t(p_{t+2})\\
&= R^{-1}\tilde{\mathbb{E}}_t(d_{t+1})+R^{-2}\tilde{\mathbb{E}}_t(d_{t+2})+R^{-2}\tilde{\mathbb{E}}_t(R^{-1}\mathbb{E}_{t+2}\left(p_{t+3}+d_{t+3}\right))\\
&\vdots\\
&= \sum_{j=1}^{T}R^{-j}\tilde{\mathbb{E}}_t\left(d_{t+j}\right)+R^{-T}\tilde{\mathbb{E}}_t(p_{t+T}),
\end{aligned} \tag{3.A.1}$$

which implies, imposing $\lim_{T\to\infty}R^{-T}\tilde{\mathbb{E}}_t(p_{t+T})=0$ that in the limit $T\to\infty$

$$p_t=\sum_{j=1}^{\infty}R^{-j}\tilde{\mathbb{E}}_t\left(d_{t+j}\right). \tag{3.A.2}$$

Now we focus on the term $\tilde{\mathbb{E}}_t\left(d_{t+j}\right)$. We have that[31]

$$\tilde{\mathbb{E}}_t\left(d_{t+1}\right)=\tilde{\mathbb{E}}_t\left(d+\theta_{t+1}+\varepsilon_{t+1}\right)=d+\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)+\tilde{\mathbb{E}}_t\left(\varepsilon_{t+1}\right)=d+\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right), \tag{3.A.3}$$

since $\tilde{\mathbb{E}}_t(\varepsilon_{t+1})=0$.

Similarly

$$\tilde{\mathbb{E}}_t\left(d_{t+2}\right)=\tilde{\mathbb{E}}_t\left(d+\theta_{t+2}+\varepsilon_{t+2}\right)=d+\tilde{\mathbb{E}}_t\left(\theta_{t+2}\right)+\tilde{\mathbb{E}}_t\left(\varepsilon_{t+2}\right)=d+\beta\tilde{\mathbb{E}}_t(\theta_{t+1}), \tag{3.A.4}$$

since $\tilde{\mathbb{E}}_t(\varepsilon_{t+2})=0$ and $\tilde{\mathbb{E}}_t(\theta_{t+2})=\tilde{\mathbb{E}}_t(\beta\theta_{t+1}+\eta_{t+2})$, and in general

$$\tilde{\mathbb{E}}_t\left(d_{t+j}\right)=d+\beta^{j-1}\tilde{\mathbb{E}}_t(\theta_{t+1})\quad\text{for all}\quad j\geq 1. \tag{3.A.5}$$

Therefore we have

$$p_t=\sum_{j=1}^{\infty}R^{-j}\left(d+\beta^{j-1}\tilde{\mathbb{E}}_t(\theta_{t+1})\right)=d\sum_{j=1}^{\infty}R^{-j}+\beta^{-1}\tilde{\mathbb{E}}_t(\theta_{t+1})\sum_{j=1}^{\infty}\left(\frac{\beta}{R}\right)^{j}. \tag{3.A.6}$$

[31] Clearly if agents observe the realization of the stochastic component then $\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)=\theta_{t+1}$. We opt to keep the notation general because this allows us to simultaneously treat also agents that do not observe the realization of this noisy component.

Now we have

$$d\sum_{j=1}^{\infty} R^{-j} = d\sum_{j=0}^{\infty} R^{-j} - d = d\frac{R}{R-1} - d = \frac{d}{r}, \tag{3.A.7}$$

since $|R^{-1}| < 1$. Similarly

$$\begin{aligned}\beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)\sum_{j=1}^{\infty}\left(\frac{\beta}{R}\right)^j &= \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)\sum_{j=0}^{\infty}\left(\frac{\beta}{R}\right)^j - \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right) = \\ &= \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)\frac{R}{R-\beta} - \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right) = \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)\frac{\beta}{R-\beta}.\end{aligned} \tag{3.A.8}$$

So that finally

$$p_t = \frac{d}{r} + \frac{\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)}{R-\beta}. \tag{3.A.9}$$

Appendix 3.B Relationship of updating to Kalman Filter

Following the notation in chapter 13 of Hamilton (1994) we have the following state space representation for the observable component of dividends θ_t:

$$\begin{aligned}\text{(state equation)} \quad & \xi_t = F\xi_{t-1} + v_t, \\ \text{(measurement equation)} \quad & y_t = H\xi_t + w_t,\end{aligned} \tag{3.B.1}$$

in which all quantities are scalars, $\xi_t \equiv \theta_{t+1}$, $F \equiv \beta$, $H \equiv 1$. The variance-covariance matrix R associated with w_t is the scalar $\sigma^2_{j,t}$. In this contest the a priori variance covariance matrix is given by

$$P_{t|t-1} = \mathbb{E}(\xi_t - \hat{\xi}_{t|t-1})^2 = \sigma^2_\eta, \tag{3.B.2}$$

and the Kalman Gain

$$K_t = P_{t|t-1}H(H'P_{t|t-1}H + R)^{-1}, \tag{3.B.3}$$

collapses to

$$K_t = \frac{\sigma^2_\eta}{\sigma^2_\eta + \sigma^2_{j,t}}, \tag{3.B.4}$$

which is exactly the weight associated to the information received by source j in the case of being connected to source j only.

Appendix 3.C Derivation of conditional variance

$$\tilde{\mathbb{V}}_t(p_{t+1} + d_{t+1}) = \tilde{\mathbb{V}}_t(d_{t+1}) + \tilde{\mathbb{V}}_t(p_{t+1}) + 2\mathbb{C}\tilde{\mathbb{O}}\mathbb{V}_t(d_{t+1}, p_{t+1}). \tag{3.C.1}$$

The conditional variance of next period dividends is given by

$$\tilde{\mathbb{V}}_t(d_{t+1}) = \tilde{\mathbb{V}}_t(\theta_{t+1}) + \sigma_\varepsilon^2. \tag{3.C.2}$$

The conditional variance of next period price can be derived starting from the expression for p_{t+1} implied by equation (3.2.9)

We get

$$\tilde{\mathbb{V}}_t(p_{t+1}) = \tilde{\mathbb{V}}_t\left(\frac{d}{r} + \frac{\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})}{R-\beta}\right) = \frac{\tilde{\mathbb{V}}_t(\tilde{\mathbb{E}}_{t+1}(\theta_{t+2}))}{(R-\beta)^2}. \tag{3.C.3}$$

In a similar fashion we can derive the expression for the covariance of the future price and dividend as

$$2\mathbb{C}\tilde{\mathbb{O}}\mathbb{V}_t(d_{t+1}, p_{t+1}) = 2\mathbb{C}\tilde{\mathbb{O}}\mathbb{V}_t\left(\theta_{t+1}, \frac{\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})}{R-\beta}\right), \tag{3.C.4}$$

since d is a constant and ε_{t+1} is independent of any other variable. Rearranging we get equation (3.2.29).

Appendix 3.D Proofs

3.D.1 Proof of Proposition 4

The proof is by induction. First we prove the statement for the base case with only one source, that is $\bar{A} = 2$. This collapses to the normal case of Bayesian updating with conjugate normal prior, therefore we have:

$$\mu_P = \frac{\mu_1\sigma_0^2 + \mu_0\sigma_1^2}{\sigma_0^2 + \sigma_1^2}, \tag{3.D.1}$$

$$\sigma_P^2 = \frac{\sigma_0^2\sigma_1^2}{\sigma_0^2 + \sigma_1^2}. \tag{3.D.2}$$

Then we prove that if the statement holds for a generic set A with $\bar{A} = n$, then it holds also for B with $\bar{B} = n+1$. If the statement holds for $\bar{A} = n$, and receive an extra source of information, the new posterior distribution has parameters:

$$\mu_P = \frac{\mu_{K+1}\frac{\prod_{j=0}^{K}\sigma_j^2}{\sum[A]^{\bar{A}-1}} + \frac{\sum_{k=0}^{K}\left(\mu_k\cdot[A]_k^{\bar{A}-1}\right)}{\sum[A]^{\bar{A}-1}}\sigma_{K+1}^2}{\frac{\prod_{j=0}^{K}\sigma_j^2}{\sum[A]^{\bar{A}-1}} + \sigma_{K+1}^2} = \frac{\frac{\sum_{k=0}^{K+1}\left(\mu_k\cdot\left[A\cup\sigma_{K+1}^2\right]_k^{\bar{A}}\right)}{\sum[A]^{\bar{A}-1}}}{\frac{\prod_{j=0}^{K}\sigma_j^2}{\sum[A]^{\bar{A}-1}} + \frac{\sum\left[\left[A\cup\sigma_{K+1}^2\right]^{\bar{A}}\setminus\prod_{j=0}^{K}\sigma_j^2\right]}{\sum[A]^{\bar{A}-1}}} = \tag{3.D.3}$$

$$= \frac{\sum_{k=0}^{K+1}\left(\mu_k\cdot\left[A\cup\sigma_{K+1}^2\right]_k^{\bar{A}}\right)}{\sum\left[A\cup\sigma_{K+1}^2\right]^{\bar{A}}}. \tag{3.D.4}$$

$$\sigma_P^2 = \frac{\frac{\prod_{j=0}^{K}\sigma_j^2}{\sum[A]^{\bar{A}-1}}\sigma_{K+1}^2}{\frac{\prod_{j=0}^{K}\sigma_j^2}{\sum[A]^{\bar{A}-1}} + \sigma_{K+1}^2} = \frac{\frac{\prod_{j=0}^{K+1}\sigma_j^2}{\sum[A]^{\bar{A}-1}}}{\frac{\prod_{j=0}^{K}\sigma_j^2}{\sum[A]^{\bar{A}-1}} + \frac{\sum\left[\left[A\cup\sigma_{K+1}^2\right]^{\bar{A}}\setminus\prod_{j=0}^{K}\sigma_j^2\right]}{\sum[A]^{\bar{A}-1}}} = \frac{\prod_{j=0}^{K+1}\sigma_j^2}{\sum\left[A\cup\sigma_{K+1}^2\right]^{\bar{A}}}.$$

Therefore we have:

$$\mu_P = \frac{\sum_{k=0}^{K+1}\left(\mu_k\cdot[B]_k^{\bar{B}-1}\right)}{\sum[B]^{\bar{B}-1}}, \tag{3.D.5}$$

$$\sigma_P^2 = \frac{\prod_{j=0}^{K+1}\sigma_j^2}{\sum[B]^{\bar{B}-1}}, \tag{3.D.6}$$

which are the posterior mean and variance for the set $B = A \cup \sigma_{K+1}^2$ with $\bar{B} = \bar{A} + 1 = n + 1$, hence concluding the proof.

3.D.2 Proof of Proposition 5

Start from equation (3.2.21). Defining as e^k the posterior variance obtained when agent has K connections, we can express the posterior variance recursively for increasing values of K as

$$e^K = e^{K-1}\frac{\sum[A]^{K-1}\sigma_K^2}{\sum[A]^K}. \tag{3.D.7}$$

To prove this formulation we can again rely on induction. The base case is $K = 1$ for which

$$e^1 = \frac{\sigma_0^2\sigma_1^2}{\sigma_0^2 + \sigma_1^2}. \tag{3.D.8}$$

Then if the formula holds for K we have

$$e^{K+1} = e^K \frac{\sum [A]^K \sigma_{K+1}^2}{\sum [A]^{K+1}}, \tag{3.D.9}$$

and iterating

$$e^{K+1} = e^1 \frac{\sum [A]^1 \sigma_2^2}{\sum [A]^2} \frac{\sum [A]^2 \sigma_3^2}{\sum [A]^3} \cdots \frac{\sum [A]^K \sigma_{K+1}^2}{\sum [A]^{K+1}}, \tag{3.D.10}$$

ans substituting the formula for e^1 by noticing that $\sum [A]^1 = \sigma_0^2 + \sigma_1^2$ we have

$$e^{K+1} = \frac{\sigma_0^2 \sigma_1^2}{\sum [A]^1} \frac{\sum [A]^1 \sigma_2^2}{\sum [A]^2} \frac{\sum [A]^2 \sigma_3^2}{\sum [A]^3} \cdots \frac{\sum [A]^K \sigma_{K+1}^2}{\sum [A]^{K+1}}, \tag{3.D.11}$$

which simplifying gives exactly

$$e^{K+1} = \frac{\prod_{j=0}^{K+1} \sigma_j^2}{\sum [A]^{K+1}}. \tag{3.D.12}$$

The reason why the recursive formula is useful is that

$$\frac{\sum [A]^K \sigma_{K+1}^2}{\sum [A]^{K+1}} < 1, \tag{3.D.13}$$

that is

$$\sum [A]^K \sigma_{K+1}^2 < \sum [A]^{K+1}, \tag{3.D.14}$$

which holds noticing that $\sum [A]^{K+1} - \sum [A]^K \sigma_{K+1}^2 = \prod_{j=0}^{K} \sigma_j^2$ and $\sigma_j^2 > 0 \forall j$. This combined with the fact that the variances are bounded makes so that

$$\lim_{K \to \infty} e^K = 0. \tag{3.D.15}$$

Appendix 3.E Validation of the SNPE

To validate the model we generate synthetic data from the ABM and then we use the SNPE algorithm to recover the parameters. We fix $\sigma_\eta = \sigma_\nu = 0.7$ and keep the other parameters to their values in table 3.1. We choose a uniform prior for the parameters with support $[0.1, 2]$ and we provide the result of 10000 stochastic simulations to train the density estimator. The result is displayed in Figure 3.9 and we can see that the

posterior distribution is extremely well behaved with median values 0.95 and 0.71 respectively for σ_η and σ_ν. The value of the first parameter is slightly off, but well within the confidence interval, and most likely due the fact that the posterior looks flat in the region of the true value.

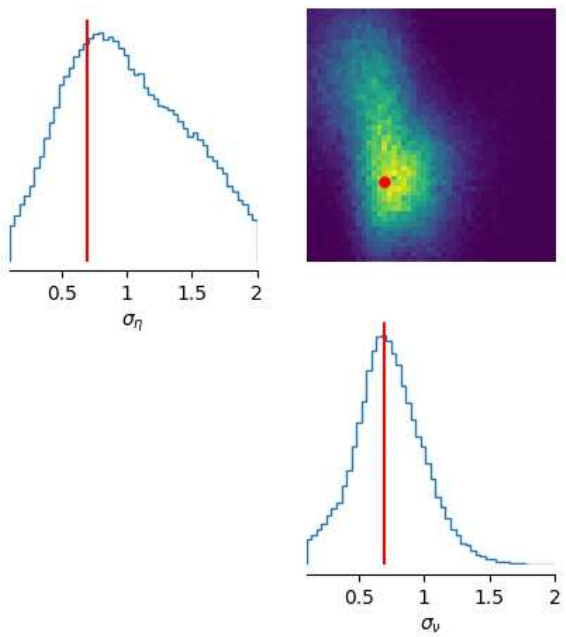

Figure 3.9: Validation of the SNPE.

Note: The Figure shows the posterior distribution of the parameters obtained by the SNPE algorithm. The red lines represent the true values of the parameters.

Chapter 4

Sticky information across the wealth distribution

4.1 Introduction

A recently expanding literature in macroeconomics focuses on building models that are consistent not only with macro aggregates but also with micro evidence. One of the main empirical findings on the aggregate level relates to the shape of output response to monetary policy shocks. An identified monetary policy shock, as in Romer and Romer (2004), induces a hump-shaped response of output. On the micro side, the literature has instead documented elevated and peaked on-impact marginal propensities to consume (MPCs) out of transitory income shocks (Fagereng, Holm and Natvik, 2021).

Models focusing on matching aggregate evidence have typically used New Keynesian models with a representative agent (RANK). To account for the delayed response of macro aggregates to shocks, these models are usually enriched with habits in the consumption behavior of households and frictions, such as sticky prices, on the firm side as in Christiano et al. (2005) or Smets and Wouters (2007). An alternative approach has been introducing deviations from the full-information rational expectations (FIRE) assumption. The main motif used by the literature has been rational inattention, as in Carroll (2003), which argues that agents may face an information processing constraint and, therefore, cannot process all available information immediately. Maćkowiak and Wiederholt (2009), Maćkowiak and Wiederholt (2010), and Zorn (2021) have shown that RANK models with some form of rational inattention can match the hump shape of the impulse response function.

To deal with micro-side evidence, the literature has instead focused on models with incomplete markets and idiosyncratic productivity risk, giving rise to wealth hetero-

geneity. This combination prevents some agents from achieving perfect consumption smoothing and helps deliver heterogeneous MPCs. The incorporation of these elements into New Keynesian models has created the so-called Heterogeneous Agent New Keynesian (HANK) models. The implications of these models, as in McKay, Nakamura and Steinsson (2016) and Kaplan, Moll and Violante (2018), for monetary policy, differ from the RANK benchmark in the General Equilibrium transmission which is mostly driven by indirect effects (via income) rather than direct effects (via interest rates). Concerning the Representative Agent setting, HANK models offer a solution to the forward guidance puzzle. However, they do not provide a stark difference in the effect of a current monetary policy shock on the economy.

Some recent papers have started to merge elements from the two traditions but have done so in isolation. They assume that the expectation formation process, even if not entirely rational, is uniform across the population and orthogonal to the state space. This chapter links the expectation formation process to wealth and studies the implications of this assumption for monetary policy transmission in a HANK model. The type of bounded rationality I use is the sticky information framework, as in Mankiw and Reis (2002) and Carroll et al. (2020). I use survey data, following the approach of Coibion and Gorodnichenko (2015), to show that households do not form their beliefs according to the FIRE assumption but display *stickiness* in incorporating new information. I then provide evidence that this stickiness is lower in households with higher wealth.

I use the HANK model with sticky expectations to revisit the mechanisms through which monetary policy is transmitted to the economy. First, by using a simple model, I show that inequality now matters for the transmission of monetary policy, even at the aggregate level. I also show that using a RANK model with sticky expectations misses the size and timing of the peak response. Second, I quantify the effect of this departure from rationality by estimating a quantitative model by matching simulated impulse responses to empirical ones for an identified Aruoba and Drechsel (2022) monetary policy shock. I show that stickiness is necessary to match the empirical estimates and that the FIRE version of the same model would overestimate the IRF of output by 0.4% on impact and 16% cumulatively. Lastly, I show that the model with sticky expectations commands a lower reaction to shocks that are further away in the future, thus lowering the forward guidance puzzle.

Relation to the literature. This chapter relates to three main strands of literature. The first one uses survey data to test deviations from the FIRE assumption.

Coibion and Gorodnichenko (2012), Coibion and Gorodnichenko (2015) and the review in Coibion, Gorodnichenko and Kamdar (2018) have shown that neither households nor firms are perfectly rational in their expectation formation process and have argued about the incorporation of this evidence in macroeconomic models. While generally information frictions have been studied on an aggregate level, some works exploit demographics data in survey data to study potential heterogeneity, with Madeira and Zafar (2015) being one of the first works in this area. Cloyne, Ferreira and Surico (2020) show that the consumption response to interest rate changes is higher in households with mortgages. While the mechanism highlighted in the paper is heterogeneity in MPCs, their findings would also be consistent with the hypothesis of heterogeneous expectations. Some recent works incorporate survey evidence into structural models. Gallegos-Dago (2023*b*) studies the effect of dispersed information in a RANK model, and Gallegos-Dago (2023*a*) extends it to a tractable HANK model in the flavor of Bilbiie (2024). Closer to this chapter is Nord (2022), who also finds heterogenous expectations driven by wealth but argues for a model of belief dispersion, with more disagreement being prevalent in low-wealth households. While the focus of that paper is on a partial equilibrium model of consumption and savings, I focus on the general equilibrium implications.

The second strand of literature considers the implications of bounded rationality for the transmission of monetary policy. Early contributions include Evans and Honkapohja (2003) and Bullard, Evans and Honkapohja (2008), who focus on adaptive learning and are interested in the implications for the determinacy of equilibrium. Perhaps a drawback of this theory of bounded rationality is that adaptive learning expectations are purely backward-looking. As in more recent years, central banks have relied on forward guidance as an alternative policy instrument, the stark result that adaptive learning would imply of no effect of forward guidance is not appealing. To obtain a more realistic effect of forward guidance, the literature has used information rigidities like information sparsity (Angeletos and Lian, 2018) or cognitive discounting (Gabaix, 2020) in a RANK setting and k-level thinking (Farhi and Werning, 2019) or incomplete information (Angeletos and Huo, 2021) in a HANK setting.

Lastly, the chapter belongs to the expanding literature on heterogeneous agents models being solved and estimated in sequence space rather than in state space, using the methodology introduced by Auclert, Bardóczy, Rognlie and Straub (2021). The same authors have shown how to adapt the methodology to deal with deviations from FIRE in Auclert, Rognlie and Straub (2020) and other works have build on this to deal with belief disagreement (Guerreiro, 2023) or noisy information and diagnostic

expectations (Bardoczy and Guerreiro, 2024). A common assumption of these works is that the expectation formation process does not depend on the state space and, in particular, on the wealth distribution. Methodologically, this chapter proposes a way to solve and estimate models in which this the assumption is relaxed.

Layout. The rest of the chapter is structured as follows. Section 4.2 presents the empirical evidence of wealth-dependent stickiness. Section 4.3 presents the simple model and highlights the main mechanisms at play. Section 4.4 presents the quantitative model and the estimation. Section 4.5 concludes.

4.2 From uniform to wealth dependent stickiness

I begin by laying out the classical sticky information framework. A continuum of agents on the unit interval populates the economy. Their forecast regarding a macroeconomic variable is given by

$$F_t(x_{t+h}) = \begin{cases} \mathbb{E}_t(x_{t+h}) & \text{with p} = (1-\theta) \\ F_{t-1}(x_{t+h}) & \text{with p} = \theta \end{cases} \tag{4.2.1}$$

Equation (4.2.1) says that in each period, an agent updates their information set with probability $(1-\theta)$, in which case they have rational expectations. With probability θ, they do not update their information set and stick with their last forecast. The economy as a whole has expectations that, on average, are given by

$$F_t(x_{t+h}) = \int_0^1 (1-\theta)\mathbb{E}_t(x_{t+h})di + \int_0^1 \theta F_{t-1}(x_{t+h})di = (1-\theta)\mathbb{E}_t(x_{t+h}) + \theta F_{t-1}(x_{t+h}). \tag{4.2.2}$$

This holds because θ and hence $F_t(x_{t+h})$ are independent of i and so constant across the population. To test whether FIRE is met in survey data, one can then use the fact that a rational forecast needs to satisfy

$$\mathbb{E}_t(x_{t+h}) = x_{t+h} - e_{t+h,t}, \tag{4.2.3}$$

with $e_{t+h,t}$ being the forecasting error and uncorrelated with information dated t or earlier. Combining equations (4.2.2) and (4.2.3) one gets

$$F_t(x_{t+h}) = (1-\theta)(x_{t+h} - e_{t+h,t}) + \theta F_{t-1}(x_{t+h}), \tag{4.2.4}$$

and rearranging

$$\frac{F_t(x_{t+h})}{1-\theta} = x_{t+h} - e_{t+h,t} + \frac{\theta}{1-\theta}F_{t-1}(x_{t+h}), \tag{4.2.5}$$

that is

$$F_t(x_{t+h}) + \frac{\theta}{1-\theta}F_t(x_{t+h}) = x_{t+h} - e_{t+h,t} + \frac{\theta}{1-\theta}F_{t-1}(x_{t+h}), \tag{4.2.6}$$

and finally

$$x_{t+h} - F_t(x_{t+h}) = \frac{\theta}{1-\theta}(F_t(x_{t+h}) - F_{t-1}(x_{t+h})) + e_{t+h,t}. \tag{4.2.7}$$

Equation (4.2.7) is testable on data since one can regress the forecasting error on the forecast revision $(F_t(x_{t+h}) - F_{t-1}(x_{t+h}))$. This can be done using Ordinary Least Squares (OLS) if the survey elicits forecasts for the same variable at different moments. The null hypothesis of FIRE is met when the coefficient associated with the forecast revision is 0, implying $\theta = 0$ and no degree of stickiness.

I proceed by first replicating the baseline estimation in Coibion and Gorodnichenko (2015) using data from the Survey of Professional Forecasters (SPF), the Livingston Survey (LS), and the Michigan Survey of Consumers (MSC). I report the specific questions asked in each survey in appendix 4.A. In the first, the target variable is the GDP/GNP inflation forecast for horizons starting at the current quarter to four quarters ahead. For the LS, agents are asked to forecast the CPI at six and twelve months ahead. The LS elicits forecasts at 6 and 12 months ahead. Therefore, t is the current semester. The SPF elicits forecasts for the next 4 quarters, therefore t is the current quarter. For both datasets, equation (4.2.7) can be estimated directly. The MSC has monthly waves, but forecast revisions cannot be directly obtained since individuals are only asked to forecast the percentage change in prices for the next 12 months. Specifically, the regression that one could run is

$$x_{t+h} - F_t(x_{t+h}) = c + \beta(F_t(x_{t+h}) - F_{t-1}(x_{t+h-1})) + \text{error}_t, \tag{4.2.8}$$

but now notice that error_t includes not only the rational forecast error but also the term $\beta(F_{t-1}(x_t) - F_{t-1}(x_{t+h}))$ and is thus not orthogonal to the regressor. As a result, OLS can not estimate this model. In this case, Coibion and Gorodnichenko (2015) suggests using oil price innovations as an instrument for the forecast revision. Oil price innovations are computed log-differences in the oil price. The estimates for the MSC are then obtained by two-stage least squares (2SLS) with quarterly frequency. To compute the forecast error for all surveys, I use the first release of the target variable

from the Federal Reserve Bank of Philadelphia Real-Time Data Set, as suggested by Croushore and Stark (2003). Table 4.1 reports the results.

Table 4.1: Information Rigidity for Households and Professionals

	LS	**SPF**	**MSC**
Constant	0.25	-0.06	-1.04
	(0.16)	(0.76)	(0.00)
Forecast Revision	0.92	0.95	1.21
	(0.00)	(0.01)	(0.00)
Sample	1969:6 - 2020:12	1969:1 - 2022:4	1982:2 - 2023:3
Observations	103	203	166
Instrument	-	-	Oil Price Innovations
F stat first stage	-	-	87.31

Note: *The table reports estimates of equation (4.2.7) using OLS for the Livingston Survey (LS), the Survey of Professional Forecasters (SPF), and the 2SLS for the Michigan Survey of Consumers (MSC). P-values obtained from Newye-West robust standard errors are in parentheses.*

From equation (4.2.7), one can directly retrieve the degree of stickiness from the estimated regression coefficient as $\theta = \frac{\hat{\beta}}{1+\hat{\beta}}$, resulting in $\theta = 0.48$ in the LS, $\theta = 0.49$ in the SPF, and $\theta = 0.50$ in the MSC. As discussed, this result is derived under the assumption of homoscedasticity in the updating probability across the population.

In what follows, I will provide evidence of heterogeneity in the probability of updating across the wealth distribution. The MSC provides demographic information on multiple dimensions, including respondents' education level, age, and stock holdings. Although the latter is not a direct measure of wealth, it is likely highly correlated with it[32]. My first step is to partition the dataset into different demographic groups and then estimate equation (4.2.7) for each group. For this analysis, I move to a monthly frequency, exploiting the MSC's monthly waves to obtain a larger sample and more powerful tests.

Table 4.2 reports the results for the whole sample, the top 10% and the bottom 90% of the wealth distribution. The first observation is that the estimated coefficient on the whole sample is higher than in the quarterly estimation. The theoretical model implies that, defining θ_m as the stickiness parameter at the monthly level, at the quarterly level, one has $\theta_q = \theta_m^3$. Indeed the 99% confidence interval at the monthly level is

[32]See for example Nord (2022)

$\theta_m \in [0.52, 0.83]$ which would give a quarterly stickiness parameter of $\theta_q \in [0.14, 0.58]$, that includes the quarterly point estimate obtained by the OLS coefficient in table 4.1 of $\theta_q = 0.55$. Second, the estimated coefficient for the top 10% is lower and statistically less significant than for the whole sample and the bottom 90%. This aligns with the hypothesis that wealthier individuals are more likely to update their information set. Third, I check whether the coefficients between the top 10% and the bottom 90% are statistically different by reporting the value of the z-statistic of the coefficient difference. Not only is the result statistically significat, but also economically meaningful. The point estimate implies a degree of stickiness of 52% for the top 10% and of 73% for the bottom 90%. In other words, since the probability of updating their information set in each period is independent, at the top of the wealth distribution, it takes an average of 2 months for a household to update its information set, while it takes almost 4 months at the bottom[33].

Table 4.2: Information Rigidity for different Wealth Levels

	All	**Top 10%**	**Bottom 90%**
Constant	-0.78	-0.33	-0.99
	(0.00)	(0.00)	(0.00)
Forecast Change	3.02	1.08	2.74
	(0.00)	(0.06)	(0.00)
Sample	1986:03 - 2023:06	1990:01 - 2023:06	1990:01 - 2023:06
Observations	306	306	306
Instrument	Oil Price Innovations	Oil Price Innovations	Oil Price Innovations
F stat first stage	39.02	24.68	29.98
p-val difference	-	-	0.09

Note: *The table reports estimates of equation (4.2.7) by 2SLS for the Michigan Survey of Consumers (MSC) at monthly frequency. P-values obtained from Newey-West robust standard errors are in parentheses.*

The Effect of Age and Education Wealth is well known to be highly correlated with age and education. Older individuals might update their information set more often as they have lived through more business cycles and could have a better understanding of the economy. Similarly, more educated households might be better able to process relevant information and update their information set more often. One might, therefore,

[33]This is because if p is the probability of updating, then the number of periods needed to update the information set follows a geometric distribution with expected value $1/p$. So, in this case, for the top 10%, one gets $1/0.48 = 2.08$, while for the bottom 90%, one gets $1/0.27 = 3.7$.

conjecture that the result obtained in table 4.2 is driven by these two variables. To test for this hypothesis, I repeated the analysis by constructing different subgroups. I need to consider large bins for both variables to have enough observations in each group. For age, I pool individuals into three age groups: under 35 years old, between 35 and 65, and above 65. For education, I consider four groups: individuals with a high school diploma, individuals who attended college but did not graduate, individuals with an undergraduate degree, and individuals with postgraduate education. Figure 4.1 reports the point estimate and the 90% confidence interval for each group. The results confirm the role of wealth as the primary driver of heterogeneity. First, all the estimates of wealth groups in the top 10% are not statistically different from zero, suggesting no degree of stickiness for these groups. Second, the differences among coefficients for the different age and education groups are not statistically significant. The only exception is between the group of individuals younger than 35 and older than 65 in the top 10% of the wealth distribution. However, this result is more the byproduct of a negative estimated point estimate for the younger group, which is in contrast with the underlying theory. There are only 295 observations in this subgroup, which might not be enough to obtain a reliable estimate.

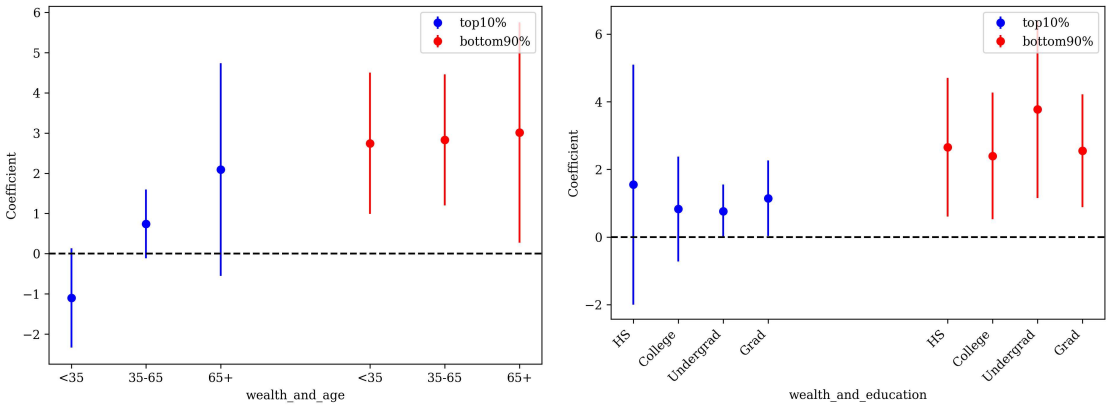

Figure 4.1: The role of age and education

4.3 Model

I begin by laying out the simple model I use to explore the implications of wealth-related stickiness. The model is purposely stylized to highlight the main mechanisms at play.

4.3.1 Households

The economy is populated by a continuum of households, which are heterogeneous in two dimensions: wealth and idiosyncratic productivity. Household $i \in [0, 1]$ solves

$$\begin{aligned} &\max_{c_{it}, a_{it}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(u(c_{it}) - v(n_{it}) \right) \\ &\text{s.t.} \quad c_{it} + a_{it} \leq (1 + r_t) a_{it-1} + z_{it} \\ &\qquad a_{it} \geq \underline{a} \end{aligned} \tag{4.3.1}$$

Utility from consumption and disutility from labor are given by $u(c) = \frac{c^{1-\sigma^{-1}}}{1-\sigma^{-1}}$ and $v(n) = \zeta \frac{n^{1+\phi^{-1}}}{1+\phi^{-1}}$. σ and ϕ are the elasticity of intertemporal substitution, and the Frisch elasticity, and ζ is the disutility of labor. $z_{it} = (1 - \tau_t) y_{it}$ is post tax real income, and $y_{it} = \frac{W_t}{P_t} e_{it} n_{it}$ is labor income. The idiosyncratic productivity e_{it} follows a Markov chain with transition matrix Π^e and is normalized such that $\mathbb{E}_i(e_{it}) = 1$. Households face a borrowing constraint $\underline{a}$, and the real interest rate is r_t.

4.3.2 Firms

A representative and perfectly competitive firm operates using a linear production technology:

$$Y_t = X_t N_t, \tag{4.3.2}$$

where Y_t represents aggregate output, N_t is the level of aggregate labor, and X_t captures productivity. Prices are flexible, and real wage is given by:

$$w_t = \frac{W_t}{P_t} = X_t. \tag{4.3.3}$$

Then inflation, $\pi_t = \log(P_t / P_{t-1})$ is given by:

$$\pi_t = \pi_t^w - \log(X_t / X_{t-1}), \tag{4.3.4}$$

where $\pi_t^w = \log(W_t / W_{t-1})$ is wage inflation.

4.3.3 Labor market

Following the approach of Auclert, Rognlie and Straub (2024) and early works by Schmitt-Grohé and Uribe (2005) and Erceg, Henderson and Levin (2000) I model the

labor market with sticky wages[34]. A continuum of unions indexed by $k \in [0,1]$ hires households and aggregates the efficient hours of work supplied to the union into a specific task $N_k t = \int e_{it} n_{ikt} di$. The different union supplies are then aggregated into the total labor supply by a competitive labor market-packer with the constant elasticity of substitution (CES) aggregator

$$N_t = \left(\int_0^1 N_{kt}^{\frac{\varepsilon-1}{\varepsilon}} dk \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{4.3.5}$$

and the total supply is sold to the representative firm at the wage W_t. I introduce nominal rigidities by assuming that unions face quadratic adjustment costs in changing the nominal wage

$$\frac{\psi}{2} \int_0^1 \left(\frac{W_{kt}}{W_{kt-1}} \right)^2 dk.$$

In each period, unions set a uniform wage and allocate hours uniformly across their workers so that the efficient amount of worked hours $\int e_{it} n_{it} di$ is equal to the aggregate N_t. All unions are symmetric, so they set the same wage W_t in equilibrium. Under these assumptions, the linearized wage Phillips curve can be derived as in Appendix 4.C and given by

$$\pi_t^w = \kappa_w \left(\sigma^{-1} \hat{c}_t + \phi^{-1} \hat{n}_t - (\hat{y}_t - \hat{\tau}_t - \hat{n}_t) \right) + \beta \mathbb{E}_t(\pi_{t+1}), \tag{4.3.6}$$

where $\hat{c}_t$, $\hat{n}_t$ and $\hat{y}_t$ are the log-deviations of consumption, labor, and output from their steady-state values, and $\hat{\tau}_t = d\tau_t/(1-\tau)$ and κ_w is the slope of the Phillips curve.

4.3.4 Government

The government chooses the path of government spending G_t, debt B_t and taxes τ_t to satisfy the government budget constraint

$$B_t = (1 + r_t) B_{t-1} + G_t - T_t, \tag{4.3.7}$$

with $T_t = \tau_t Y_t$ and where for now I work under the assumption that spending is set exogenously, bonds are fixed at the level B, and taxes adjust to satisfy the budget constraint.

[34]Using sicky prices and flexible wages, as common in the New Keynesian literature, leads to an implausible result when embedding heterogenous agents into the model as it would imply countercyclical profits. I refer to Broer, Hansen, Krusell and Öberg (2020) and Auclert et al. (2024) for a broader discussion.

4.3.5 Monetary authority

The monetary authority sets the nominal interest rate according to a Taylor rule

$$i_t = r + \phi_\pi \pi_t + \epsilon_t, \tag{4.3.8}$$

in which r is the natural real interest rate, ϕ_π is the response of the nominal interest rate to inflation, and ϵ_t is a monetary policy shock.

4.3.6 Equilibrium

Aggregate quantities are given by

$$C_t = \int_0^1 c_{it} di, \quad A_t = \int_0^1 a_{it} di. \tag{4.3.9}$$

Then, given initial conditions for nominal wage, price level, government debt, and the initial distribution of households over assets and productivity, the exogenous path of government spending, the equilibrium is a sequence of prices, quantities, and allocations such that households, firms, and unions optimize, the government budget constraint and the monetary policy rule are satisfied, and markets clear:

$$\begin{aligned} C_t + G_t &= Y_t, \\ A_t &= B. \end{aligned} \tag{4.3.10}$$

4.3.7 Beliefs

The model layout so far leverages the assumption that beliefs satisfy FIRE. Following the empirical evidence in section 4.2, I model households with sticky expectations. With respect to the general framework which I have laid down so far, I make the following additional assumptions. First, I assume that all households are aware of the steady state values of all variables in the economy[35]. Second, I assume that households have sticky information only with respect to aggregate variables and not idiosyncratic ones. Specifically, given that households know the steady states, they have sticky information concerning *shocks* to aggregate variables. This aligns with the literature and eases the comparison with existing work. Moreover, while I tested directly for stickiness to aggregate variables, it is not straightforward to test for stickiness to idiosyncratic shocks. Third, I follow Carroll et al. (2020) and Auclert et al. (2020) in assuming that once

[35]Note that in the general model given by equation (4.2.1), this is true if the economy starts at the steady state and $F_{-1}(x_{t+h}) = x_{-1}$.

a shock hits the economy, all agents become aware of it and update their information sets accordingly. The first two assumptions imply that the steady state of the model with sticky information (Sticky HANK model hereafter) coincides with the FIRE one. Finally, and the main novelty of this chapter, I assume that the wealth distribution endogenously determines the degree of stickiness. Since the survey data evidence is not enough to provide the full functional form for the relationship between wealth and stickiness, I assume[36] that a power function relates them

$$\theta(a) = \left(\frac{\underline{a} + c^*}{a + c^*}\right)^{\gamma}. \tag{4.3.11}$$

This functional form is convenient as it is determined by two parameters γ and c^* only, which can then be directly calibrated. Specifically, I will set them so that they match the average stickiness, on a quarterly basis, of the top 10% of the distribution and the bottom 90% by numerically solving the following system of equations

$$\begin{cases} \int_{\underline{a}}^{a_{90}} \left(\frac{\underline{a} + c^*}{a + c^*}\right)^{\gamma} dD(a) = 0.39 \\ \int_{a_{90}}^{\bar{a}} \left(\frac{\underline{a} + c^*}{a + c^*}\right)^{\gamma} dD(a) = 0.14, \end{cases} \tag{4.3.12}$$

where a_{90} is the 90^{th} percentile of the wealth distribution, and the integral is taken over the wealth distribution. The resulting parameters are $c^* = 0.06$ and $\gamma = 0.24$, which implies an average stickiness in the population of 0.42. The steady-state distribution of the wealth-dependent stickiness is displayed in Figure 4.2

4.3.8 Solution Method

Solving the Sticky HANK model is challenging for two reasons. First, one has to deal with household heterogeneity and incomplete markets. For this, I rely on the contribution by Auclert et al. (2020) by using the Sequence-Space representation of the model. Second, one has to consider the deviation from FIRE imposed by sticky expectations. For this, I rely on the methodology developed in Auclert et al. (2020), Guerreiro (2023), and Bardoczy and Guerreiro (2024). However, these works commonly assume that the deviations from FIRE are orthogonal to the model's state variables. This assumption is not met as the wealth distribution endogenously determines expectations in the

[36] A theory supporting this assumption is that of rational inattention in the style of Sims (2003). In this setting, beliefs' endogeneity makes it challenging to directly microfound the relationship between wealth and stickiness. In related settings, however, Guerreiro (2023) and Shabalina and Tzaawa-Krenzler (2025) find a similar relationship between attention and income.

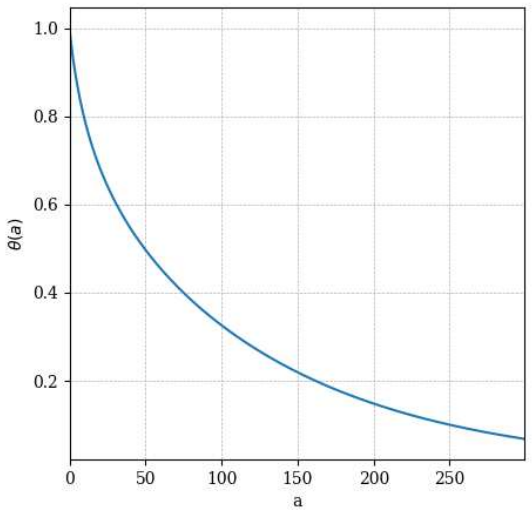

Figure 4.2: Steady state distribution of the wealth dependent stickiness

Sticky HANK model. A core contribution of this chapter is a method to solve the model in this case. To build sequentially towards this solution, I proceed as follows. First, I show how to compute the economy's steady state and then solve for first-order impulse responses to a shock in sequence-space thanks to the methodology in Auclert et al. (2020). As already remarked, the steady state of the Sticky HANK model and the FIRE one coincides. So, without having to recompute the steady state, I show how to compute impulse-responses in the Stikcy HANK by leveraging how the FIRE equivalents are calculated.

Steady State. In a steady state, aggregate variables are constant over time. Since the steady state is common knowledge, both under FIRE and sticky expectations, agents forecast perfectly aggregate variables. The challenging part with respect to a representative agent economy is to solve the household side of the model. The dynamic programming problem associated with the steady state is

$$\begin{aligned} V(a,e) = \max_{c,a'} & u(c) - v(n) + \beta \mathbb{E}[V(a',e')|e], \\ \text{s.t.} \quad & c + a' = (1+r)a + z, \\ & a' \geq \underline{a}. \end{aligned} \tag{4.3.13}$$

The solution method proceeds as follows. First, the state space is discretized. In this specific case, I use a grid with 500 points for the asset level and 7 possible values for the idiosyncratic productivity. The second step is to obtain the individual policies

for consumption and assets by backward iterating on the value function. The specific algorithm used is the endogenous grid method by Carroll (2006), which iterates on the derivative of the value function with respect to assets rather than to the value function itself. A key step, which I highlight because it will be important later on, is that, in general, the asset policy function will point households to an asset level that is not in the grid. To deal with this, the solution method uses a lottery that assigns the households to either the closest grid point below or above the optimal asset level. For an agent with desired asset a lying between two grid points a_j and a_{j+1}, the probability of being assigned to a_j is given by

$$\pi_j = \frac{a_{j+1} - a}{a_{j+1} - a_j}. \tag{4.3.14}$$

Finally, given the steady state asset policy function and the exogenous Markov process for idiosyncratic productivity, the steady state distribution can be computed by iterating the distribution forward. Aggregate steady-state quantities can then be computed as

$$C = \int c(a,e)dD(a,e), \quad A = \int a'(a,e)dD(a,e), \tag{4.3.15}$$

where $D(a,e)$ is the distribution of households over assets and productivity. In a steady state, the market clearing conditions are

$$C + G = Y, \quad A = B, \tag{4.3.16}$$

and a stationary equilibrium is then a sequence of prices, quantities, and allocations such that households, firms, and unions optimize, the government budget constraint and the monetary policy rule are satisfied, markets clear and the steady state constraint is satisfied.

FIRE dynamics away from the steady state. The study of dynamics in response to shocks hitting the economy leverages the Sequence-Space representation of the model. The object of interest is the *sequence* of states of an aggregate endogenous variable, for example, consumption in response to a change in an exogenous variable, such as the real interest rate. To denote infinite sequences I will use the notation $\{S_t\} \equiv (S_t, S_{t+1}, S_{t+2} \dots)$. The shocks considered will be MIT shocks, which are shocks that come as a surprise at time 0 in a steady-state economy. For sequences that start at time $t = 0$, I will also use exchangbly the notation $\mathbf{S} \equiv (S_0, S_0, S_0 \dots) \equiv \{S_0\}$. Then to first-order, the output $\mathbf{o}$ change in response to a change of an input $\mathbf{i}$ is given

by

$$d\mathbf{o} = \mathcal{J}^{o,i} d\mathbf{i}, \tag{4.3.17}$$

where $\mathcal{J}^{o,i}$ is the Jacobian summarizing the partial derivatives

$$\mathcal{J}^{o,i} = \begin{pmatrix} \frac{\partial o_0}{\partial i_0} & \frac{\partial o_0}{\partial i_1} & \dots \\ \frac{\partial o_1}{\partial i_0} & \frac{\partial o_1}{\partial i_1} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix}. \tag{4.3.18}$$

The most challenging Jacobians to compute pertains to the household side of the model. A brute force approach to computing them uses the same method described to obtain the steady state for each matrix entry. As an example, assume that one wishes to compute the partial derivative of consumption at date t to the interest rate at date s. Then, one would do a backward iteration starting from a period T at which the economy is assumed to be back in the steady state. The backward iteration can handle time-varying inputs, and one can carry it out by setting $\mathbf{r}$ at the steady state except at date s, in which it is shocked by an infinitesimal amount h After obtaining the respective policies, one would iterate forward the distribution and aggregate to compute C_t. Then the partial derivative would be given by

$$\frac{\partial C_t}{\partial r_s} = \frac{C_t - C}{h}. \tag{4.3.19}$$

This method is computationally expensive, requiring $T \times T$ backward and forward iterations for each Jacobian. One of the central contributions of Auclert et al. (2020) is a method to compute these Jacobians more efficiently. This method assumes that agents have perfect foresight, but Auclert et al. (2020) shows how to compute the Jacobians in the model with deviations from FIRE at almost no computational cost starting from the FIRE one. However, their method relies on the assumption that the deviations from FIRE are orthogonal to the model's state variables. Moving forward, I will show how to compute Jacobians for the households side in the Sticky HANK model in two ways. First, by brute force, then by leveraging the method in Auclert et al. (2020) and adapting it to the case where the deviations from FIRE are not orthogonal to the state variables.

Sticky dynamics away from the steady state. The brute force solution leverages the following idea. Consider an individual i with wealth a at time 0. Assume that this individual is going to keep their wealth constant over time. Then, when a shock is announced to hit the economy at time s, the probability that the individual will learn

about it at a time $\tau < s$ is $(1-\theta(a))\theta(a)^{\tau}$. Moreover, as long as the population of agents with the same asset level is large enough, then this probability coincides with the fraction of agents with the same asset level that updated their information set. In this case, across that asset level, at time τ, the economy reacts as if the shock hitting at time s has size $(1-\theta(a))\theta(a)^{\tau}$ of the original shock. The brute force algorithm then consists of computing the Jacobian as described in the previous part but rescaling the size of the shock to the correct level at each asset level: from h to $h(1-\theta(a))\theta(a)^{\tau}$. This method relies on the approximation that while agents are moving in the distribution, they move in such a way that

$$\theta(a_k) \approx \frac{\sum_{a_j} m_{a_j,a_k}\theta(a_j)}{\sum_{a_j} m_{a_j,a_k}}, \tag{4.3.20}$$

where m_{a_j,a_k} is the mass of agents that have wealth a_j at time $\tau-1$ and a_k at time τ. In other words, the probability of updating at time τ for an agent with wealth a_k is the weighted average of the updating probability of agents with wealth a_j at time $\tau-1$. I discuss the validity of this approximation in appendix 4.B. This method is computationally expensive, so I propose a different process for obtaining the Jacobians in the Sticky HANK model. The brute force method will then serve as a benchmark to test the validity of the new process.

The methodology builds on the one in Auclert et al. (2020) to compute Jacobians with deviations from FIRE. It needs, however, to be adapted to the case in which the degree of stickiness θ is endogenous and determined by the distribution. I begin by reassuming the method used in Auclert et al. (2020). This follows two assumptions. The first one is that it is possible to partition the population at any time t into two groups: those who have updated their information set at least once in a period $\tau \leq t$, thus learning about the shock and those who have not. This assumption is met also in this case. The second one is that the updating probability is orthogonal to idiosyncratic shocks and constant over the population and time. This assumption is not met in this case. If it were, then it would be possible to derive the sticky expectations Jacobian as

$$\mathcal{J}^{o,i} = (1-\theta)\sum_{\tau=0}^{\infty}\theta^{\tau}\mathcal{J}^{o,i,\tau}. \tag{4.3.21}$$

In the context of this chapter, a similar relationship can be derived, but this involved an additional step to deal with the endogenous evolving stickiness, as I explain below.

FIRE Jacobian manipulation

Consider a representative household with a time-varying stickiness parameter θ_t. I define the probability that this agent has not updated their information set in any period up to τ as

$$P_\tau = \begin{cases} 1 & \text{if } \tau = -1 \\ \prod_{k=0}^{\tau} \theta_k & \text{if } \tau \geq 0 \end{cases}. \tag{4.3.22}$$

Then, the probability of learning about a shock at time $\tau < t$ for the households can be derived as:

$$\ell_\tau = (1 - \theta_\tau) P_{\tau-1}.$$

Then, consider an economy populated by a continuum of representative households, this can be divided at any point in time into two groups: those who have updated their information set at least once in a period $\tau \leq t$, thus learning about the shock and those who have not.

Then the aggregate Jacobian describing the output sequence o response to input sequence i can be written as

$$\mathcal{J}^{o,i} = \sum_{\tau=0}^{\infty} \ell_\tau \mathcal{J}^{o,i,\tau}. \tag{4.3.23}$$

The insights provided in appendix D.3 of Auclert et al. (2020) still apply and can be used to derive the following relationship between the sticky information Jacobian and the Jacobian of the model with perfect foresight. Specifically, it is still true that a household learning at date τ about a shock at time s to input i will have the same response as a household learning at time 0 about the same shock at time $s - \tau$ to input i, shifted by τ periods. That is

$$\mathcal{J}^{o,i,\tau}_{t,s} = \mathcal{J}^{o,i,\tau-1}_{t-1,s-1} = \cdots = \mathcal{J}^{o,i,0}_{t-\tau,s-\tau}. \tag{4.3.24}$$

I also maintain the assumption that all households are aware of the shock when it hits at time s so that if $\tau > s$ then $\mathcal{J}^{o,i,\tau}_{t,s} = \mathcal{J}^{o,i,s}_{t,s}$. With this I can rewrite equation (4.3.23) as

$$\mathcal{J}^{o,i}_{t,s} = \sum_{\tau=0}^{s-1} \ell_\tau \mathcal{J}^{o,i,\tau}_{t,s} + \sum_{\tau=s}^{\infty} \ell_\tau \mathcal{J}^{o,i,s}_{t,s}. \tag{4.3.25}$$

The term $\sum_{\tau=s}^{\infty} \ell_\tau$ can be simplified as follows. First, note that

$$\ell_\tau = P_{\tau-1} - P_\tau, \tag{4.3.26}$$

since

$$P_{\tau-1} - P_\tau = P_{\tau-1} - (P_{\tau-1} \cdot \theta_\tau) = P_{\tau-1}(1 - \theta_\tau). \tag{4.3.27}$$

Then, the sum can be written as

$$S = \sum_{\tau=s}^{\infty} (P_{\tau-1} - P_\tau) = \left(\sum_{\tau=s}^{\infty} P_{\tau-1} \right) - \left(\sum_{\tau=s}^{\infty} P_\tau \right). \tag{4.3.28}$$

Then, noticing that the first sum is equal to the second sum shifted by one period, I have

$$S = P_{s-1} + \left(\sum_{n=s}^{\infty} P_n \right) - \left(\sum_{n=s}^{\infty} P_n \right) = P_{s-1}. \tag{4.3.29}$$

Therefore

$$S = P_{s-1} = \sum_{\tau=s}^{\infty} \ell_\tau. \tag{4.3.30}$$

Then (4.3.25) becomes

$$\mathcal{J}_{t,s}^{o,i} = \sum_{\tau=0}^{s-1} \ell_\tau \mathcal{J}_{t,s}^{o,i,\tau} + \mathcal{J}_{t,s}^{o,i,s} P_{s-1}. \tag{4.3.31}$$

Now, one can write

$$\mathcal{J}_{t,s}^{o,i} = \mathcal{J}_{t,s}^{o,i,s} P_{s-1} + \sum_{\tau=1}^{s-1} \ell_\tau \mathcal{J}_{t,s}^{o,i,\tau} + \ell_0 \mathcal{J}_{t,s}^{o,i,0}, \tag{4.3.32}$$

and applying (4.3.24) I have

$$\mathcal{J}_{t,s}^{o,i} = \mathcal{J}_{t-1,s-1}^{o,i,s-1} P_{s-1} + \sum_{\tau=0}^{s-2} \ell_{\tau+1} \mathcal{J}_{t-1,s-1}^{o,i,\tau} + \ell_0 \mathcal{J}_{t,s}^{o,i,0}. \tag{4.3.33}$$

Notice now that shifting (4.3.31) by 1 period in both t and s I have

$$\mathcal{J}_{t-1,s-1}^{o,i} = \sum_{\tau=0}^{s-2} \ell_\tau \mathcal{J}_{t-1,s-1}^{o,i,\tau} + \mathcal{J}_{t-1,s-1}^{o,i,s-1} P_{s-2}. \tag{4.3.34}$$

which is almost the same expression as the first two terms in (4.3.33), except for $\ell_{\tau+1}$ in place of ℓ_τ and P_{s-2} in place of P_{s-1}. It is possible to use the equivalence $P_{s-1} = P_{s-2}\theta_{s-1}$, but the relationship betwen $\ell_{\tau+1}$ and ℓ_τ can be derived as follows

$$\ell_{\tau+1} = (1-\theta_{\tau+1})P_\tau = (1-\theta_{\tau+1})\theta_\tau P_{\tau-1} = \ell_\tau \frac{(1-\theta_{\tau+1})\theta_\tau}{1-\theta_\tau}. \tag{4.3.35}$$

The last term is time dependent and prevents me from directly expressing $\mathcal{J}^{o,i}_{t,s}$ in terms of $\mathcal{J}^{o,i}_{t-1,s-1}$. To overcome this, I approximate[37] the term as

$$\frac{(1-\theta_{\tau+1})\theta_\tau}{1-\theta_\tau} \approx \theta_{s-1}. \tag{4.3.36}$$

With this I can write (4.3.33) as

$$\mathcal{J}^{o,i}_{t,s} = \mathcal{J}^{o,i}_{t-1,s-1}\theta_{s-1} + \ell_0 \mathcal{J}^{o,i,0}_{t,s}. \tag{4.3.37}$$

Now, notice that for $s = 0$ then (4.3.31) becomes $\mathcal{J}^{o,i}_{t,0} = \mathcal{J}^{o,i,0}_{t,0}$. For $t = 0$ and $s > 0$ then households react only if $\tau = 0$, so $\mathcal{J}^{o,i}_{0,s} = \ell_0 \mathcal{J}^{o,i,0}_{0,s}$. Combining these insights, and realizing that $\mathcal{J}^{o,i,0}_{t,s} = \mathcal{J}^{o,i,\text{FI}}_{t,s}$ since it just the full-information Jacobian, I can write

$$\mathcal{J}^{o,i}_{t,s} = \begin{cases} \mathcal{J}^{o,i}_{t-1,s-1}\theta_{s-1} + \ell_0 \mathcal{J}^{o,i,\text{FI}}_{t,s} & t > 0, s > 0 \\ \mathcal{J}^{o,i,\text{FI}}_{t,s} & s = 0 \\ \ell_0 \mathcal{J}^{o,i,\text{FI}}_{t,s} & t = 0, s > 0. \end{cases} \tag{4.3.38}$$

This is a convenient formulation since it relates the sticky information Jacobian to the full information Jacobian and the time-varying stickyness parameter. It can thus be computed efficiently as long as the path of θ_t is known. However, the path of θ_t is endogenously determined by the distribution of wealth in the economy, and in essence, this is a fixed point problem. To solve it, I can use the following insight. At the aggregate level, the stickiness of the economy in steady state is given by

$$\theta = \int \theta(a) dD(a). \tag{4.3.39}$$

I can then treat the path of θ_t as any aggregate output and compute the Jacobian mapping any shock to input i to the path of θ_t, $\mathcal{J}^{\theta,i}$ under FIRE assumption. I can then use the following algorithm to compute the effect of a shock to input i on the path of θ_t.

1. Guess an initial path for $\{\theta_t\}^0$.

2. Starting from $\mathcal{J}^{\theta,i,\text{FIRE}}_{t,s}$, compute $\mathcal{J}^{\theta,i}_{t,s}$ using (4.3.38).

[37] I discuss the validity of this approximation in Appendix 4.B.

3. Apply (4.3.17) to obtain the change in θ_t and add it to the steady state value to get a new path $\{\theta_t\}^1$.

4. Check if $||\{\theta_t\}^0 - \{\theta_t\}^1|| < \epsilon$, for an exogenously set tolerance level ϵ.

5. Run until convergence.

The resulting path of $\{\theta_t\}$ can then be used to obtain the sticky information Jacobian relating any output o to input i by applying (4.3.38) to the appropriate FIRE Jacobian.

4.3.9 Calibration

I calibrate the model to a quarterly frequency. The steady state of interest is the one with zero inflation, implying that nominal and real interest rates coincide. Table 4.3 reports the value of the parameters used in the calibration.

Table 4.3: Calibration

Parameter	**Description**	**Value**
Households		
σ	Elasticiy of intertemporal substition	0.5
ϕ	Firsch elasticity	0.5
ζ	Disitulity of labor	1.0
ρ_e	Persistence of idiosyncratic productivity shocks	0.92
sd_e	Standard deviation of idiosyncratic productivity shocks	0.92
n_e	Number of productivity states	11
$\underline{a}$	Minimum asset level	0
$\bar{a}$	Maximum asset level	4000.0
n_a	Number of asset grid points	250
γ	Sensitivity of stickiness to wealth	0.08
Firms		
X	Steady state TFP level	1.0
κ_w	Wage rigidity	0.16
Monetary		
ϕ_π	Taylor rule coefficient	1.5
r	Real interest rate	0.5%
Fiscal		
G/Y	Spending to GDP ratio	0.16
B	Bond supply	5.6
Targeted		
β	Discount factor	0.95

4.3.10 Monetary policy

I now consider the implications of the model for monetary policy. As a useful reference model, I also consider a representative agent (RANK) model, in which the only difference is the household side of the model. Specifically, the representative household

solves

$$\begin{aligned} &\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\sigma}}{1-\sigma} \\ &\text{s.t.} \quad C_t + A_t \leq (1+r_t)A_{t-1} + Z_t. \end{aligned} \tag{4.3.40}$$

The calibration for this model is also the same as in the HANK model, except for the value of the discount factor β, which is now set to $\beta = 1/(1+r)$. For consistency to the remainder of the chapter, I assume that the monetary authority shocks the beginning of the period interest rate r^{ante}. This will allow me to handle valuation effects when adding dividends. This can be accommodated into the model by introducing a simple rule mapping the ex-ante to the ex-post interest rate

$$\begin{cases} r_t = r_{t-1}^{\text{ante}} & \text{if } t \geq 1 \\ r_t = r & \text{if } t = 0 \end{cases} \tag{4.3.41}$$

I start in the FIRE framework and consider the response of aggregate output to a shock to the real interest rate which deviates 1% from its steady state on impact and returns to the steady state with an exponentially declining weight of 0.7. As remarked by Auclert et al. (2020), it is helpful to visualize the macro model as a Directed Acyclic Graph (DAG). Since B and G are constant, given the fiscal rule, T will adapt to match the fiscal authority budget constraint. Also, for the moment, I can focus on the real side of the economy, given that steady-state inflation is 0. Under these assumptions, the DAG, common for both HANK and RANK, is shown in Figure 4.3. Given that market clearing is 0, the DAG represents the following relationship in sequence space

$$\mathbf{Y} = \mathbf{G} + \mathcal{C}(\mathbf{Z}, \mathbf{r}), \tag{4.3.42}$$

where the notation $\mathcal{C}$ stresses that aggregate consumption and aggregate net of taxes income are functions of sequences. Indeed, also the ex-post interest rate and $\mathbf{Z}$ are functions. Then, in response to a shock to the real interest rate, differentiating (4.3.42) yields (since $\mathbf{G}$ is constant)

$$d\mathbf{Y} = \mathcal{J}^{C,r} d\mathbf{r} + \mathcal{J}^{C,Z} d\mathbf{Z}, \tag{4.3.43}$$

and by explicitly considering the dependence of the fiscal and ex-post interest rates

$$d\mathbf{Y} = \mathcal{J}^{C,r} d\mathbf{r} + \mathcal{J}^{C,Z} \left(\mathcal{J}^{Z,Y} d\mathbf{Y} + \mathcal{J}^{Z,r} d\mathbf{r} \right), \tag{4.3.44}$$

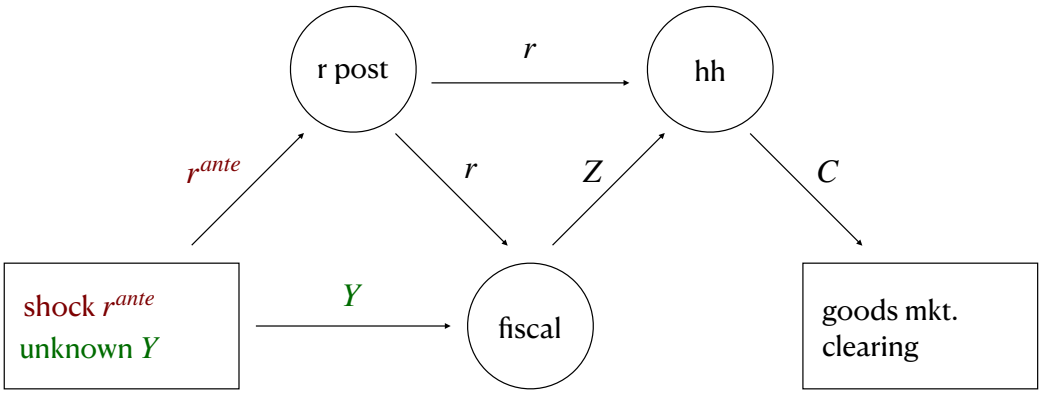

Figure 4.3: DAG of the model for monetary shock

whit $d\mathbf{r} = \left(\mathcal{J}^{r,r^{\text{ante}}} d\mathbf{r}^{\text{ante}}\right)$. This Jacobians in (4.3.43) are what Auclert et al. (2024) call intertemporal marginal propensities to consume (iMPCs) and are of particular interest since they are sufficient statistics for the response of aggregate output to a monetary policy shock. This formulation is also helpful in decomposing the effect of the shock into direct (via intertemporal substitution) and indirect (via income) effects. The left panel of Figure 4.4 shows the IRFs of the output gap in the HANK and RANK models to an interest rate shock in the right panel. Within the proposed calibration, the

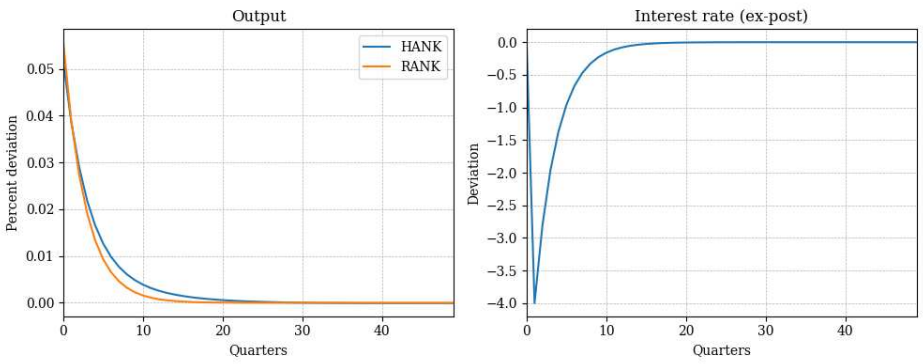

Figure 4.4: Monetary policy in the HANK and RANK

finding is that under FIRE, the difference in aggregate response from RANK to HANK is minimal. Indeed, this is a well-known result in the literature, (Ahn, Kaplan, Moll, Winberry and Wolf, 2018). While the aggregate response is similar, the decomposition

into direct and indirect effects in Figure 4.5 reveals that indirect effects via income are much more critical in HANK. Indeed, in the RANK model, marginal propensities to consume are zeros, and therefore, the only channel is the direct one. Again, however, the difference is minimal for a policymaker interested only in the aggregate response.

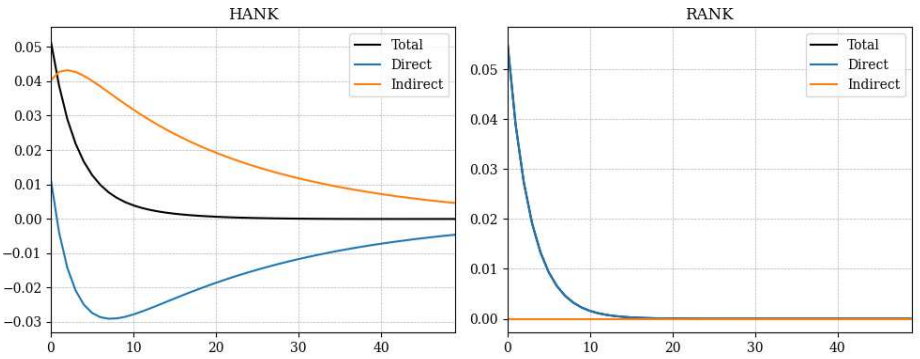

Figure 4.5: Decomposition

Now, I consider the same shock in the case in which agents have sticky information. I also consider a sticky version of the RANK model to maintain a fair comparison. Absent assets heterogeneity, however, I set the stickiness in RANK to the steady state value of the HANK model and constant for all agents. Then, I apply the algorithm described in section 4.3.8 to the iMPCs of the equation (4.3.43) to obtain the sticky information Jacobians.[38] The right panel of Figure 4.6 shows the IRFs of the HANK and RANK models under sticky information. In the second case, although the RANK model displays some dampening, it is only in the HANK version that the response does not peak on impact. When individuals have sticky information, inequality matters.

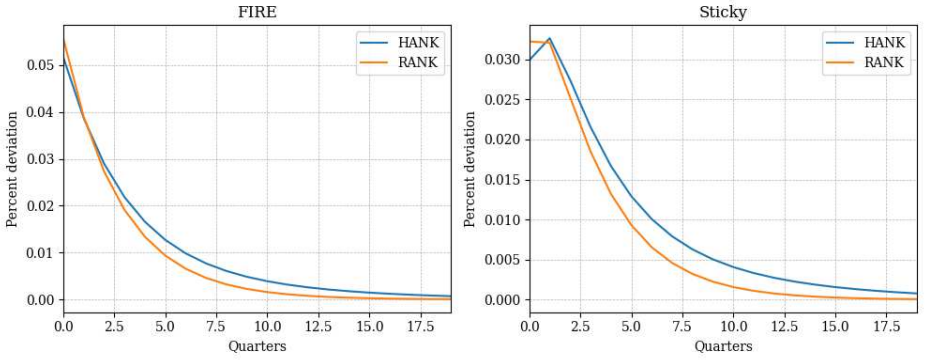

Figure 4.6: Monetary policy in the HANK and RANK under sticky information

[38]In the RANK case, since the stickiness is constant, there is no need to iterate, and the transformation in (4.3.38) can be directly applied.

4.4 Estimation

In this section, I move on to the estimation of the model on real data in order to quantify the effect of the information frictions on actual variables. To do so, I consider a slightly different model, including a more realistic supply side, that could accommodate the importance of investments in the transmission of monetary policy, as pointed out in Auclert et al. (2024).

4.4.1 Model

Households

The household side is almost the same, but now it is assumed that households also receive dividends from ownership of firms. Moreover taxes and dividends follow an incidence rule. Specifically, I assume that they are both proportional to household productivity. Their optimization problem is then

$$\begin{aligned} &\max_{c_{it},a_{it}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(u(c_{it}) - v(n_{it})\right) \\ &\text{s.t.} \quad c_{it} + a_{it} \leq (1+r_t)a_{it-1} + y_{it} - \tau_t \bar{\tau}(e_{it}) + d_t \bar{d}(e_{it}) \\ &\qquad a_{it} \geq \underline{a} \end{aligned} \tag{4.4.1}$$

Firms The main difference is the introduction of capital into the economy. The firm side is populated by a final good producer and a continuum of intermediate goods producers, indexed by j. Intermediate producers have Cobb-Douglas production function

$$y_{jt} = F(n_{jt}, k_{jt-1}) = n_{jt}^{1-\alpha} k_{jt-1}^{\alpha}, \tag{4.4.2}$$

and face adjustment costs when choosing their capital stock. I follow Auclert et al. (2020) in setting quadratic adjustment costs

$$\Psi_t^k \left(k_{jt}, k_{jt-1}\right) = \frac{k_{jt}}{k_{jt-1}} - (1-\delta) + \frac{1}{2\delta\epsilon_I} \left(\frac{k_{jt} - k_{jt-1}}{k_{jt-1}}\right)^2, \tag{4.4.3}$$

with δ the depreciation rate and ϵ_I the capital adjustment cost parameter, both positive. Aggregate investment then evolves by

$$I_t = K_t - (1-\delta)K_{t-1} + \Psi_t^k \left(K_t, K_{t-1}\right) K_{t-1}. \tag{4.4.4}$$

Price setting of intermediate firms is also subject to quadratic adjustment costs,

$$\Psi_t^p\left(p_{jt}, p_{jt-1}\right) = \frac{\mu}{\mu - 1}\frac{1}{2\kappa} - \left[\log(p_{jt}) - \log(p_{jt-1})\right]^2, \tag{4.4.5}$$

where μ is the constant elasticity of substitution of the final good producer and κ is the price adjustment cost parameter. The Phillips curve for aggregate inflation is then

$$\log(1+\pi_t) = \kappa\left(\frac{w_t}{F'_N(N_t, K_{t-1})} - \frac{1}{\mu}\right) + \frac{1}{1+r_{t+1}}\frac{Y_{t+1}}{Y_t}\log(1+\pi_{t+1}). \tag{4.4.6}$$

Dividends are given by output net of investments, the remuneration of labor , and the price adjustment costs, so that $d_t = Y_t - I_t - w_t N_t - \Psi_t^p(P_t, P_{t-1})Y_t$. The evolution of capital is determined jointly with the Tobin's Q

$$Q_t = 1 + \frac{1}{\delta\epsilon_I}\frac{K_t - K_{t-1}}{K_{t-1}}, \tag{4.4.7}$$

$$\begin{aligned}(1+r_{t+1})Q_t = \alpha_t \frac{Y_{t+1}}{K_t}\frac{w_{t+1}}{F'_N(N_{t+1}, K_t)} - \\ \left[\frac{K_{t+1}}{K_t} - (1-\delta) + \frac{1}{2\delta\epsilon_I}\left(\frac{K_{t+1}-K_t}{K_t}\right)^2\right] + \frac{K_{t+1}}{K_t}Q_{t+1}.\end{aligned} \tag{4.4.8}$$

Fiscal and Monetary authority

The fiscal authority follows a similar rule as before, but now taxes are proportional to labor, so $\tau_t w_t N_t = r_t B_t + G_t$. The Taylor rule now includes also a reaction to the output gap $i_t = r + \phi_\pi \pi_t + \phi_y\left(Y_t - Y_{\text{ss}}\right) + \epsilon_t$, and the Fisher equation reads $r_t = (1 + i_{t-1})/(1 + \pi_t)$.

Market Clearing

Asset market clearing is the same as before, but goods market clearing also includes investments and price adjustment costs so that $Y_t = C_t + I_t + G_t + \Psi_t^p(P_t, P_{t-1})Y_t$.

4.4.2 Empirical IRFs

I estimate the model by matching model-implied IRFs to empirical ones. The causal shock of interest is the estimated monetary policy shock from Aruoba and Drechsel

(2022), which builds on the Romer and Romer (2004) procedure but also uses Natural Language Processing techniques. It computes the shock as changes in the Federal Funds Rate orthogonal to all available FED forecasts and text-based time series.

To pin down the causal dynamic effect of this shock on macroeconomic aggregates, I use a Bayesian Vector Autoregression (BVAR) model, following Caravello, McKay and Wolf (2024). I target three outcomes: output, inflation, and the nominal interest rate, and collect their empirical impulse responses over 25 quarters by stacking them into the vector $\hat{\boldsymbol{\Upsilon}}$, with covariance matrix $\boldsymbol{\Sigma}$. Appendix 4.D provides details on the computation of these two objects.

To estimate the model, I proceed in two steps. First, I calibrate all parameters that are not directly estimated, as shown in Table 4.3. Then, I estimate the following set of parameters: the degree of household inattention γ, the coefficient of the Taylor rule reaction to inflation ϕ_π, the coefficient of the Taylor rule reaction to the output gap ϕ_y, the price adjustment cost parameter κ, the wage rigidity parameter κ_w, and the investment adjustment cost parameter ϵ_I. Collecting these parameters in the vector $\boldsymbol{\Psi} \equiv (\gamma, \phi_\pi, \phi_y, \kappa, \kappa_w, \epsilon_I)$, an approximate likelihood of the data $\hat{\boldsymbol{\Upsilon}}$ as a function of $\boldsymbol{\Psi}$ is given by:

$$p(\hat{\boldsymbol{\Upsilon}} \mid \boldsymbol{\Psi}) \propto \exp\left[-0.5\left(\hat{\boldsymbol{\Upsilon}} - \boldsymbol{\Upsilon}(\boldsymbol{\Psi})\right)' \boldsymbol{\Sigma}^{-1}\left(\hat{\boldsymbol{\Upsilon}} - \boldsymbol{\Upsilon}(\boldsymbol{\Psi})\right)\right]. \tag{17}$$

The posterior for $\boldsymbol{\Psi}$ given the policy shock causal effect data $\hat{\boldsymbol{\Upsilon}}$ is then:

$$p(\boldsymbol{\Psi} \mid \hat{\boldsymbol{\Upsilon}}) = \frac{p(\hat{\boldsymbol{\Upsilon}} \mid \boldsymbol{\Psi}) p(\boldsymbol{\Psi})}{p(\hat{\boldsymbol{\Upsilon}})},$$

where:

$$p(\hat{\boldsymbol{\Upsilon}}) = \int p(\hat{\boldsymbol{\Upsilon}} \mid \boldsymbol{\Psi}) p(\boldsymbol{\Psi}) d\boldsymbol{\Psi}.$$

I use a Metropolis-Hastings algorithm described in Appendix 4.E to trace out the posterior distribution. I estimate the model under FIRE and sticky information and compare the results in Figure 4.8, where the empirical IRFs are compared with the model-implied IRFs.

The latter are obtained by simulating the model with parameter values set to the posterior mode. The sticky model is highly accurate in matching the path of the nominal interest rate. It captures the hump shape in the output gap response but fails to match the peak of the reaction and overestimates the inflation response. There are two main reasons for this:

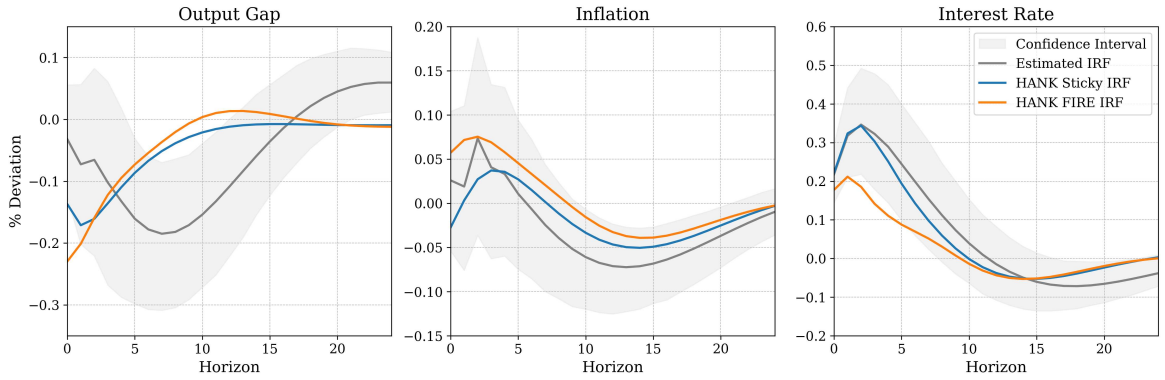

Figure 4.7: Empirical IRFs

First, the stickiness is calibrated exogenously and depends on the steady-state wealth distribution. The calibrated stickiness value, around 0.42, is far below the estimates of other studies, which typically find values close to 0.9 (Auclert et al., 2020) (Caravello et al., 2024). One could hypothesize that given the increasing inequality in recent decades, stickiness may also increase, leading to time-varying effects of monetary policy. However, since the model is stationary in this setting, shocks can only temporarily affect the wealth distribution itself, translating into temporary effects on stickiness. Exploring this further in future research could be worthwhile.

Second, the lack of frictions could play a role, especially on the supply side. This omission is deliberate to focus on the role of information frictions, but as argued, for example, in Auclert et al. (2020), the investment channel plays a crucial role in the transmission mechanism of boundedly rational HANK models. A more rigid investment side could help match the peak of the output gap response.

The FIRE model, on the other hand, while able to match inflation and interest rates to a reasonable extent, fails entirely in terms of the output gap response. I report the prior and posterior values of the estimated parameters in Tables 4.4 and 4.5.

Table 4.4: Prior and Posterior Distributions for the HANK FIRE model

Parameter	Prior Dist.	Mode	Mean	Median	5 percent	95 percent
ϕ_π	Gamma$(1.5, 0.5)$	1.400	1.350	1.377	1.192	1.482
ϕ_y	Gamma$(0.5, 0.5)$	0.490	0.477	0.476	0.427	0.539
κ_p	Gamma$(0.1, 0.1)$	0.150	0.144	0.144	0.091	0.208
κ_w	Gamma$(0.1, 0.1)$	0.120	0.105	0.104	0.056	0.166
ϵ_I	Gamma$(4, 2)$	1.700	1.686	1.685	1.634	1.750

Table 4.5: Prior and Posterior Distributions for the HANK sticky model

Parameter	Prior Dist.	Mode	Mean	Median	5 percent	95 percent
ϕ_π	Gamma$(1.5, 0.5)$	1.060	1.069	1.067	1.041	1.103
ϕ_y	Gamma$(0.5, 0.5)$	0.610	0.621	0.618	0.594	0.658
κ_p	Gamma$(0.1, 0.1)$	0.050	0.053	0.050	0.027	0.088
κ_w	Gamma$(0.1, 0.1)$	0.270	0.275	0.274	0.246	0.311
ϵ_I	Gamma$(4, 2)$	1.850	1.288	1.391	0.415	1.934

Monetary policy in the estimated model

As the only difference between the two models is the presence of sticky information, I explore its effect on the transmission of monetary policy shocks. I compute the on-impact and cumulative differences in the output gap response between the two models to quantify this. The cumulative difference is given by $\sum\{Y_t^{\text{sticky}}\}/\{Y_t^{\text{FIRE}}\} - 1$ and equals -2.18%. The on-impact difference is given by $Y_0^{\text{sticky}}/Y_0^{\text{FIRE}} - 1$ and equals 0.14%. The size of the shock is set to the average of the estimated monetary policy shock of Aruoba and Drechsel (2022), modeled as an AR(1) process with a persistence of 0.9. In Figure 4.8, I show the model-implied IRFs for output, inflation, and interest rate in response to the described shock.

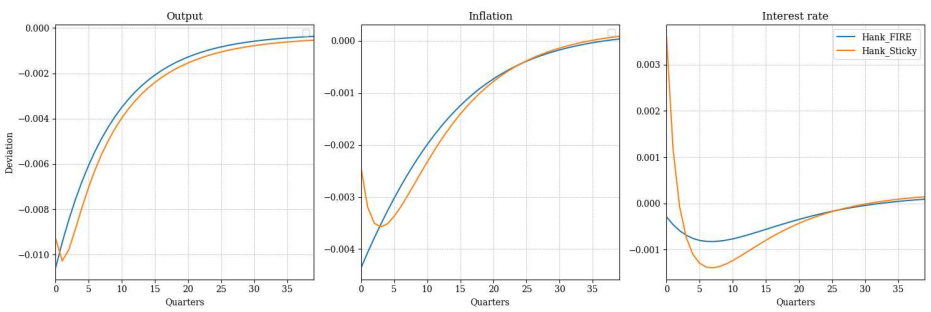

Figure 4.8: Monetary policy in the estimated model

The right panel is perhaps the most surprising, as it shows that a positive shock does not increase the nominal interest rate under FIRE. Given the Taylor rule specified above, the general equilibrium effects of a decrease in inflation and output more than compensate for the direct effect of the shock. This is not the case in the sticky version of the model, where the output and especially inflation responses are dampened. This leads to an increase in the nominal interest rate and to the documented hump shapes.

4.5 Conclusion

This chapter advances the understanding of monetary policy transmission by introducing wealth-dependent information stickiness into a HANK model. Through empirical analysis of survey data, I provide evidence of wealth-dependent stickiness, which is incorporated into the model using a novel computational methodology that extends the sequence-space methods introduced by Auclert et al. (2020). My approach modifies the computation of the General Equilibrium Jacobians to account for the endogenous evolution of the wealth distribution and state-dependent updating probabilities.

I use this methodology to efficiently solve the model, simulate the economy's response to monetary policy shocks, and estimate the model by matching impulse responses. The results reveal that wealth-dependent stickiness significantly alters the transmission mechanisms of monetary policy. Specifically, ignoring heterogeneity in information updating, as in representative agent models with uniform stickiness, leads to an underestimation of both the magnitude and the delay of the peak response to monetary shocks.

Appendix 4.A Survey Data

For the surveys used in the analysis, I report either a description of the variable or the question asked to the respondents.

Survey of Professional Forecasters: *Querterly expectations of inflation (measured by the GNP/GDP price index and, alternatively, the CPI).*

Michigan Survey of Consumers: *By about what percent do you expect prices to go (up/down) on the average, during the next 12 months?*

Livingston Survey: *Give the current-month (June and December) forecasts [of the CPI] and then base [your] six-month and 12-month forecasts on [your] current-month predictions.*

- Gini Index for the US. SIPOVGINIUSA from the FRED at yearly frequency. To obtain the index for each quarter or semester, I linearly interpolate between available observations.

When dealing with individual forecasts, I had to deal with missing values. In doing so, I followed the survey procedure as reported in the technical note. Specifically NaNs and "don't know if up or down" responses were dropped from the sample. However "don't know how much up" and "don't know how much down" were imputed by the mean of the other responses. In the orignial procedure the imputation is done by matching the distribution of the other responses. However, since the sample mean would be unaffected and higher-order moments are not used in this analysis, I opted for the more straightforward approach.

Appendix 4.B Method approximation

What is required for the approximation to be accurate is that agents move across the distribution in such a way that the probability of ending up in wealth level a is almost the same for agents starting at wealth levels $a-\Delta$ and $a+\Delta$, for any Δ. Two conditions must be met for this to happen. The first is that the mapping from current to future wealth levels must be almost linear. The second is that the wealth space is unbounded. To assess whether the first condition is met, in Figure 4.9, I plot net savings as a function of current wealth for different income levels in the steady state.

As can be seen, the relationship is almost linear, except for extremely low levels of wealth. Zooming in on the low-wealth region, one can see that the nonlinearity is accentuated only for the bottom 0.1% of the distribution. The second condition

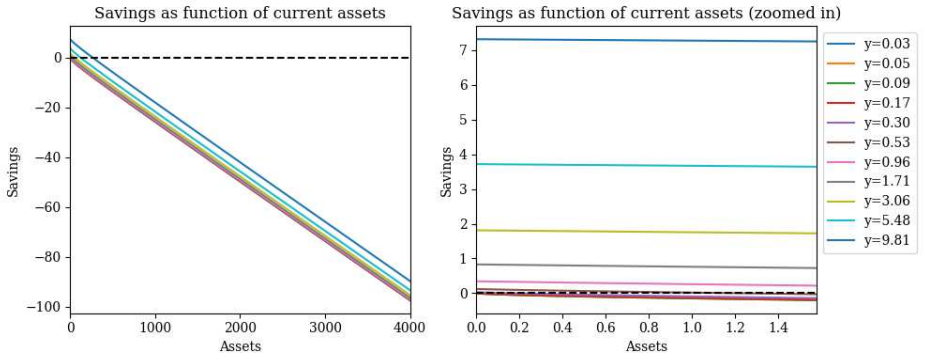

Figure 4.9: Net Savings as a Function of Assets for different Income Levels

is clearly not satisfied since the grid is discretized between a minimum and maximum asset level. Then, especially at the boundaries, the approximation will be less accurate. However, given the shape of the function mapping wealth to probability of stickiness, for high values of assets, the function is almost flat so that the error will be small. For example, at the 95% level, one has $\theta = 0.517$ while at the maximum asset level $\theta = 0.515$, a difference of 0.2 percentage points. At the bottom of the distribution, the error will be more significant. I can only argue that the proportion of agents in this region is small, as shown in Figure 4.10. The only exception is the lower bound in which there is a mass of 1.4% of agents in a steady state.

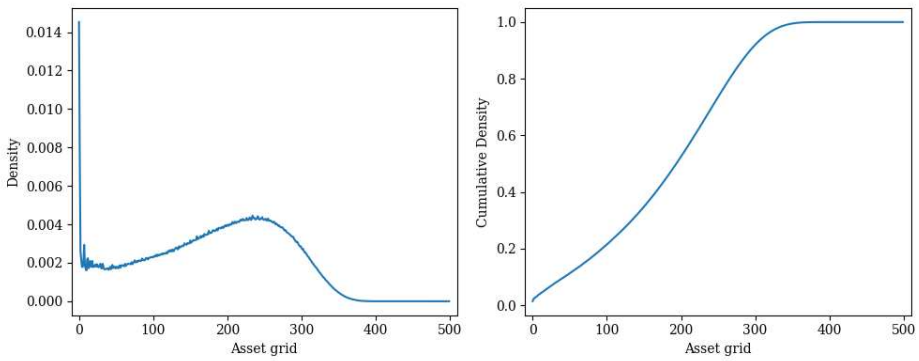

Figure 4.10: Wealth Distribution

I then use this benchmark to evaluate the accuracy of the approximation in equation (4.3.36). First, notice that altough the approximation deteriorates for large s, this effect is mediated by the fact that the term $\ell_{\tau+1} = (1 - \theta_{\tau+1})P_\tau$ gets smaller since $P_\tau = \prod_{k=0}^{\tau} \theta_k$ is rapidly decreasing in τ. To numerically show the accuracy of the approximation, I use the following exercise. I plot the difference between the path of the degree of stickiness implied by the brute force and by the approximation in Figure 4.11 for the interest rate shock depicted in the right panel of Figure 4.4.

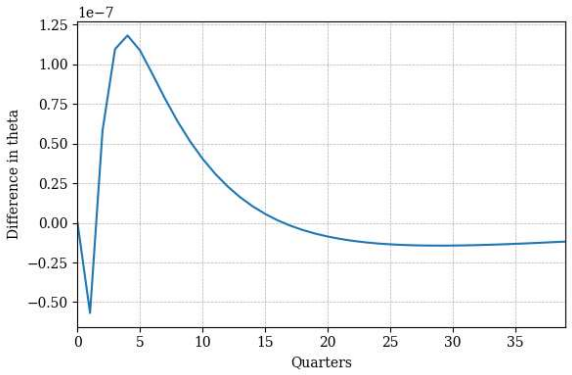

Figure 4.11: Difference between the brute force and the approximate method

The two methods are comparable up to six decimal places, and the difference gradually decreases over the time horizon.

Appendix 4.C Phillips Curve

The union k faces the following maximization problem at time t

$$\sum_{h=0}^{\infty} \beta^h \left[u'(C_{t+h})(1-\tau_{t+h}) \frac{W_{u,t+h} N_{u,t+h}}{P_{t+h}} - v'(N_{t+h}) N_{u,t+h} - \frac{1}{2\psi} \left(\frac{W_{u,t+h}}{W_{u,t+h-1}} - 1 \right)^2 \right], \tag{4.C.1}$$

which depends on the marginal utilities of aggregate quantities, which ignores the effect of the union decision on the distribution.[39] The union sets wages monopolistically, considering the demand curve of the labor packer

$$N_{kt} = \left(\frac{W_{kt}}{W_t} \right)^{-\varepsilon} N_t. \tag{4.C.2}$$

Under these assumptions, the first order condition reads

$$\left(e^{\pi_t^w} - 1\right) e^{\pi_t^w} = \psi(\varepsilon - 1) \left[-u'(C_t)(1-\tau_t) Y_t + \frac{\varepsilon}{\varepsilon - 1} v'(N_t) N_t \right] + \beta_t \left(e^{\pi_{t+1}^w} - 1\right) e^{\pi_{t+1}^w}, \tag{4.C.3}$$

and can be linearized to

$$\pi_t^w = \kappa_w \left(\sigma^{-1} \hat{c}_t + \phi^{-1} \hat{n}_t - (\hat{y}_t - \hat{\tau}_t - \hat{n}_t) \right) + \beta \pi_{t+1}, \tag{4.C.4}$$

[39] In this I follow Wolf (2021), Mckay and Wolf (2022) and Guerreiro (2023). Using the approach in Auclert et al. (2024) of considering an average utility , keeping track of these distributional consequences would provide little quantitative difference at a great computational cost.

where $\kappa_w = \psi\varepsilon v'(N)N$.

Appendix 4.D Empirical Impulse Responses

The estimation follows Caravello et al. (2024) and the empirical data is taken from the paper reprodcution files. The reduced-form VAR is given by

$$y_t = \sum_{\ell=1}^{p} A_\ell y_{t-\ell} + e_t, \tag{4.D.1}$$

where y_t is the vector of observables, A_ℓ is the and e_t is the vector of residuals. The Wold innovations η_t are then obtainied by $C^{-1}e_t$ where C is the Cholesky factor of the covariance matrix of the residuals $\Sigma_e = C\Sigma_\eta C'$. For the estimation, the values in the Aruoba and Drechsel (2022) series that are missing are imputed to be 0 and monthly data are aggregated by averaging. In addition to the monetary policy shock, the VAR includes the following variables:

- **Output Gap.** Row data is log output per capita (FRED series A939RX0Q048SBEA). Detrending is achieved by a regression of the variable al date $t+8$ quarters ahed on the four most recent values at date t, as recommended by Hamilton (2018).
- **Nominal Interest Rate.** Measured by the FRED series (FEDFUNDS).
- **Inflation.** Measured by the log-difference in GDP deflator (FRED series GDPDEF).

The shock is ordered first, the number of lags is set to 2, and a linear time trend is included. The estimation is done in a Bayesian setting, so to obtain a number of draws $N =$ for the impulse responses, the approach followed is that in Arias, Rubio-Ramírez and Waggoner (2018) of using a uniform-normal-inverse-Wishart posterior over the orthogonal reduced-form parametrization. The N draws are then used to compute the $\hat{\boldsymbol{\Upsilon}}$ by stacking the posterior mode of the impulse responses $\hat{\boldsymbol{\Upsilon}}_i \quad i = 1, \ldots, N$. The following steps compute the covariance matrix $\boldsymbol{\Sigma}$. First, the following object is constructed

$$\bar{\boldsymbol{\Sigma}} = \sum_{i=1}^{N} \left(\hat{\boldsymbol{\Upsilon}}_i - \hat{\boldsymbol{\Upsilon}}\right)\left(\hat{\boldsymbol{\Upsilon}}_i - \hat{\boldsymbol{\Upsilon}}\right)', \tag{4.D.2}$$

then, to accommodate this matrix for a small sample size, the matrix is transformed according to the following criteria:

- **Diagonal Elements.** The diagonal elements are preserved.

- **Off-Diagonal Elements less than H horizons apart.** The off-diagonal element corresponding to lags ℓ and j are scaled by the factor

$$\left(1 - \frac{|\ell - j|}{H}\right), \quad \ell, j = 1, \ldots, H$$

- **Off-Diagonal Elements more than H horizons apart.** Covariances for horizons more than H apart are set to 0.

Appendix 4.E Posterior Distribution

I use a standard Random Walk Metropolis Hastings algorithm with a multivariate normal for the proposal distribution. The variance-covariance matrix is initially assumed to be equal to the prior variance-covariance matrix, scaled by a constant. I use the first N_a draws to estimate the variance-covariance matrix of the proposal distribution, updating the proposal variance-covariance matrix to the observed variance-covariance matrix of parameters in the first N_a draws. Once updated, I sample another $N_b + N_c$ draws, burn the first N_b and keep the last N_c draws, which I use as our posterior distribution. I set $N_a = N_c = 100000, N_b = 50000$. The acceptance rates are 22.17% and 24.19% for the sticky and FIRE model, respectively.

Chapter 5

Non technical summary

5.1 Summary in English

This thesis explores how people and economic agents (like households, firms, and governments) form expectations about the future and how these expectations influence economic and financial outcomes. Traditionally, economists have assumed that agents form expectations rationally, meaning they use all available information perfectly to predict future events. This idea, known as Full Information Rational Expectations (FIRE), has been a cornerstone of economic modeling for decades. However, this assumption has been increasingly criticized for being unrealistic, especially after the 2008 financial crisis, which exposed the limitations of models relying on FIRE.

The thesis aims to address these limitations by exploring alternative ways of modeling expectations that are more realistic and better reflect how people actually think and behave. It does this in three main chapters, each focusing on a different aspect of expectation formation.

Chpter 2 studies a financial market in which different investors use different forecasting strategies to make decisions. In the presence of boundedly rational agents, with biased forecasts and trend following rules, we study the effect of two types of speculation: one based on fundamentalist and the other on rational expectations. We bring the model to data by estimating it on the Bitcoin Market with two contributions, relying on methods from Machine Learning. First, we construct the Bitcoin Twitter Sentiment Index (BiTSI) to measure sentiment regarding the asset. Second, we propose a new method for estimating the resulting heterogeneous agent model with rational speculators. We show that the switching finds support in the data and that while fundamentalist speculation amplifies volatility, rational speculation stabilizes the market.

Chapter 3 investigates the interplay between information diffusion in social networks and its impact on financial markets. We build a model in which agents receive and exchange information about an asset dividend process. A small proportion of the network has access to either private information or misinformation. Other agents can receive information only from their peers and will adjust their beliefs according to the confidence they have in the source of information. We examine, by means of simulations, how information diffuses in the network and provide a framework to account for the delayed absorption of new information into the market price, which is in contrast to what is predicted by classical financial models. We investigate the effect of the network topology on the resulting asset price and evaluate under which conditions misinformation diffusion can make the market more inefficient.

Chapter 4 focuses on the role of the inattention of households to economic information for the transmission of monetary policy in a Heterogeneous Agent New Keynesian (HANK) model. Using survey data, we provide empirical evidence that households do not form expectations according to the FIRE hypothesis but instead are slow in updating their information, with wealthier households updating twice as much on a monthly basis. We then incorporate this finding into a macroeconomic model and demonstrate that it significantly influences the aggregate responses to monetary shocks. Specifically, models that overlook heterogeneity in information updating tend to underestimate both the magnitude and the delay of the peak response to monetary policy shocks. Estimating the model by aligning simulated impulse response functions (IRFs) with empirical ones indicates that stickiness is essential for accurately capturing the dynamics observed in the data.

5.2 Samenvatting (Summary in Dutch)

Bibliography

Aalborg, H. A., Molnár, P. and de Vries, J. E. (2019), 'What can explain the price, volatility and trading volume of bitcoin?', *Finance Research Letters* **29**, 255–265.

Acemoglu, D. and Ozdaglar, A. (2011), 'Opinion dynamics and learning in social networks', *Dynamic Games and Applications* **1**(1), 3–49.

Acemoglu, D., Ozdaglar, A. and ParandehGheibi, A. (2010), 'Spread of (mis)information in social networks', *Games and Economic Behavior* **70**(2), 194–227.

Adam, K. and Nagel, S. (2023), Chapter 16 - expectations data in asset pricing, *in* R. Bachmann, G. Topa and W. van der Klaauw, eds, 'Handbook of Economic Expectations', Academic Press, pp. 477–506.

Ahn, S. H., Kaplan, G., Moll, B., Winberry, T. and Wolf, C. (2018), 'When inequality matters for macro and macro matters for inequality', *NBER Macroeconomics Annual* **32**, 1–75.

Angeletos, G. M. and Huo, Z. (2021), 'Myopia and anchoring', *American Economic Review* **11**, 1166–1200.

Angeletos, G. M. and Lian, C. (2018), 'Forward guidance without common knowledge', *American Economic Review* **108**, 2477–2512.

Anufriev, M., Gardini, L. and Radi, D. (2020), 'Chaos, border collisions and stylized empirical facts in an asset pricing model with heterogeneous agents', *Nonlinear Dynamics* **102**(2), 993–1017.

Arias, J. E., Rubio-Ramírez, J. F. and Waggoner, D. F. (2018), 'Inference based on structural vector autoregressions identified with sign and zero restrictions: theory and applications', *Econometrica* **86**(2), 685–720.

Armantier, O., Bruine de Bruin, W., Potter, S., Topa, G., van der Klaauw, W. and Zafar, B. (2013), 'Measuring inflation expectations', *Annual Review of Economics* **5**(Volume 5, 2013), 273–301.

Armstrong, J., Black, R., Laxton, D. and Rose, D. (1998), 'A robust method for simulating forward-looking models', *Journal of Economic Dynamics and Control* **22**, 489–501.

Aruoba, S. B. and Drechsel, T. (2022), 'Identifying monetary policy shocks: A natural language approach', *SSRN Electronic Journal* .

Auclert, A., Bardóczy, B., Rognlie, M. and Straub, L. (2021), 'Using the sequence-space jacobian to solve and estimate heterogeneous-agent models', *Econometrica* **89**, 2375–2408.

Auclert, A., Rognlie, M. and Straub, L. (2020), 'Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model', *Manuscript* .

Auclert, A., Rognlie, M. and Straub, L. (2024), 'The intertemporal keynesian cross', *Journal of Political Economy* pp. 4068–4121.

Axtell, R. L. and Farmer, J. D. (2023), 'Agent-based modeling in economics and finance: Past, present, and future', *Journal of Economic Literature* **forthcoming**.

Axtell, R. L. and Farmer, J. D. (Forthcoming), 'Agent-based modeling in economics and finance: Past, present, and future', *Journal of Economic Literature* .

Baig, A., Blau, B. M. and Sabah, N. (2019), 'Price clustering and sentiment in bitcoin', *Finance Research Letters* **29**, 111–116.

Banerjee, A. V. (1992), 'A simple model of herd behavior', *The Quarterly Journal of Economics* **107**(3), 797–817.

Barberis, N., Shleifer, A. and Vishny, R. (1998), 'A model of investor sentiment', *Journal of Financial Economics* **49**(3), 307–343.

Bardoczy, B. and Guerreiro, J. (2024), 'Unemployment insurance in macroeconomic stabilization with imperfect expectations', *Working Paper* .

Baur, D. G. and Dimpfl, T. (2018), 'Asymmetric volatility in cryptocurrencies', *Economics Letters* **173**, 148–151.

Baur, D. G. and Glover, K. J. (2014), 'Heterogeneous expectations in the gold market: Specification and estimation', *Journal of Economic Dynamics and Control* **40**, 116–133.

Benabou, R. and Laroque, G. (1992), 'Using privileged information to manipulate markets: Insiders, gurus, and credibility', *Quarterly Journal of Economics* **107**(3), 921–958.

Benettin, G., Galgani, L., Giorgilli, A. and Strelcyn, J.-M. (1980), 'Lyapunov characteristic exponents for smooth dynamical systems and for hamiltonian systems; a method for computing all of them. part 1: Theory', *Meccanica* **15**(1), 9–20.

Bertella, M. A., Silva, J. N., Correa, A. L. and Sornette, D. (2021), 'The influence of confidence and social networks on an agent-based model of stock exchange', *Complexity* .

Bilbiie, F. O. (2024), 'Monetary policy and heterogeneity: An analytical framework', *Review of Economic Studies* p. rdae066.

Biondo, A. E. (2020), 'Information versus imitation in a real-time agent-based model of financial markets', *Journal of Economic Interaction and Coordination* **15**(3), 613–631.

Black, F. (1986), 'Noise', *The Journal of Finance* **41**(3), 528–543.

Blankespoor, E., deHaan, E. and Marinovic, I. (2020), 'Disclosure processing costs, investors' information choice, and equity market outcomes: A review', *Journal of Accounting and Economics* **70**(2-3), 101344.

Boehl, G. and Hommes, C. (2025), 'Rational vs. irrational beliefs in a complex world', *Journal of Economic Behavior and Organization* p. forthcoming.

Bollobás, B., Borgs, C., Chayes, J. and Riordan, O. (2003), Directed scale-free graphs, *in* 'Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms', SODA '03, Society for Industrial and Applied Mathematics, USA, p. 132–139.

Bolt, W., Demertzis, M., Diks, C., Hommes, C. and van der Leij, M. (2019), 'Identifying booms and busts in house prices under heterogeneous expectations', *Journal of Economic Dynamics and Control* **103**, 234–259.

Boswijk, H. P., Hommes, C. H. and Manzan, S. (2007), 'Behavioral heterogeneity in stock prices', *Journal of Economic Dynamics and Control* **31**, 1938–1970.

Bourghelle, D., Jawadi, F. and Rozin, P. (2022), 'Do collective emotions drive bitcoin volatility? a triple regime-switching vector approach', *Journal of Economic Behavior and Organization* **196**, 294–306.

Branch, W. A. (2006), Restricted perceptions equilibria and learning in macroeconomics, *in* D. Colander, ed., 'Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model', Cambridge University Press, p. 135–160.

Bray, M. (1982), 'Learning, estimation, and the stability of rational expectations', *Journal of Economic Theory* **26**.

Bray, M. M. and Savin, N. E. (1986), 'Rational expectations equilibria, learning, and model specification', *Econometrica* **54**.

Brock, W. A. and Hommes, C. H. (1997), 'A rational route to randomness', *Econometrica* **65**(5), 1059–1095.

Brock, W. A. and Hommes, C. H. (1998), 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model', *Journal of Economic Dynamics and Control* **22**, 1235–1274.

Brock, W. A., Hommes, C. H. and Wagener, F. O. (2005), 'Evolutionary dynamics in markets with many trader types', *Journal of Mathematical Economics* **41**, 7–42.

Broer, T., Hansen, N. J. H., Krusell, P. and Öberg, E. (2020), 'The new keynesian transmission mechanism: A heterogeneous-agent perspective', *Review of Economic Studies* **87**, 77–101.

Buechel, B., Hellmann, T. and Klößner, S. (2015), 'Opinion dynamics and wisdom under conformity', *Journal of Economic Dynamics and Control* **52**, 240–257.

Bullard, J., Evans, G. and Honkapohja, S. (2008), 'Monetary policy, judgment, and near-rational exuberance', *American Economic Review* **98**, 1163–1177.

Cagan, P. D. (1956), The monetary dynamics of hyperinflation, *in* M. Friedman, ed., 'Studies in the Quantity Theory of Money', The University of Chicago Press, Chicago, pp. 25–117.

Caravello, T. E., McKay, A. and Wolf, C. K. (2024), 'Evaluating policy counterfactuals: A "var-plus" approach', *Working Paper* .

Carroll, C., Crawley, E., Slacalek, J., Tokuoka, K. and White, M. (2020), 'Sticky expectations and consumption dynamics', *American Economic Journal: Macroeconomics* **12**(3), 40–76.

Carroll, C. D. (2003), 'Macroeconomic expectations of households and professional forecasters', *Quarterly Journal of Economics* **118**.

Carroll, C. D. (2006), 'The method of endogenous gridpoints for solving dynamic stochastic optimization problems', *Economics Letters* **91**, 312–320.

Cheah, E. T. and Fry, J. (2015), 'Speculative bubbles in bitcoin markets? an empirical investigation into the fundamental value of bitcoin', *Economics Letters* **130**, 32–36.

Chen, C. Y. H. and Hafner, C. M. (2019), 'Sentiment-induced bubbles in the cryptocurrency market', *Journal of Risk and Financial Management* **12**, 53.

Chetty, R. (2006), 'A new method of estimating risk aversion', *American Economic Review* **96**(5), 1821–1834.

Chiarella, C. (1992), 'The dynamics of speculative behaviour', *Annals of Operations Research* **37**(1), 101–123.

Chiarella, C., He, X. Z. and Zwinkels, R. C. (2014), 'Heterogeneous expectations in asset pricing: Empirical evidence from the s&p500', *Journal of Economic Behavior and Organization* **105**, 1–16.

Chiarella, C., ter Ellen, S., He, X. Z. and Wu, E. (2014), 'Fear or fundamentals? heterogeneous beliefs in the european sovereign cds market', *Journal of Empirical Finance* **32**, 19–34.

Christiano, L. J., Eichenbaum, M. and Evans, C. L. (2005), 'Nominal rigidities and the dynamic effects of a shock to monetary policy', *Journal of Political Economy* **113**, 1–45.

Clarke, J., Chen, H., Du, D. and Hu, Y. J. (2021), 'Fake news, investor attention, and market reaction', *Information Systems Research* **32**(1), 35–52.

Cloyne, J., Ferreira, C. and Surico, P. (2020), 'Monetary policy when households have debt: New evidence on the transmission mechanism', *Review of Economic Studies* **87**, 102–129.

Coibion, O. and Gorodnichenko, Y. (2012), 'What can survey forecasts tell us about information rigidities?', *Journal of Political Economy* **120**, 116–159.

Coibion, O. and Gorodnichenko, Y. (2015), 'Information rigidity and the expectations formation process: A simple framework and new facts', *American Economic Review* **105**, 2644–2678.

Coibion, O., Gorodnichenko, Y. and Kamdar, R. (2018), 'The formation of expectations, inflation, and the phillips curve', *Journal of Economic Literature* **56**, 1447–1491.

Collin-Dufresne, P. and Fos, V. (2016), 'Insider trading, stochastic liquidity, and equilibrium prices', *Econometrica* **84**(4), 1441–1475.

Cont, R. and Bouchaud, J.-P. (2000), 'Herd behavior and aggregate fluctuations in financial markets', *Macroeconomic Dynamics* **4**(2), 170–196.

Cornea-Madeira, A., Hommes, C. and Massaro, D. (2019), 'Behavioral heterogeneity in u.s. inflation dynamics', *Journal of Business and Economic Statistics* **37**, 288–300.

Croushore, D. and Stark, T. (2003), 'A real-time data set for macroeconomists: Does the data vintage matter?', *Review of Economics and Statistics* **85**, 605–617.

D'Acunto, F. and Weber, M. (2024), 'Why survey-based subjective expectations are meaningful and important', *Annual Review of Economics* **16**(Volume 16, 2024), 329–357.

Dawid, H. and Delli Gatti, D. (2018), Agent-based macroeconomics, *in* C. Hommes and B. LeBaron, eds, 'Handbook of Computational Economics', Vol. 4 of *Handbook of Computational Economics*, Elsevier, pp. 63–156.

Day, R. H. and Huang, W. (1990), 'Bulls, bears and market sheep', *Journal of Economic Behavior and Organization* **14**(3), 299–329.

De Grauwe, P. and Rovira Kaltwasser, P. (2012), 'Animal spirits in the foreign exchange market', *Journal of Economic Dynamics and Control* **36**(8), 1176–1192.

De Long, J. B., Shleifer, A., Summers, L. H. and Waldmann, R. J. (1990), 'Noise trader risk in financial markets', *Journal of Political Economy* **98**, 703–738.

Degroot, M. H. (1974), 'Reaching a consensus', *Journal of the American Statistical Association* **69**(345), 118–121.

Dellavigna, S. and Pollet, J. M. (2009), 'Investor inattention and friday earnings announcements', *Journal of Finance* **64**(2), 709–749.

DeMarzo, P. M., Vayanos, D. and Zwiebel, J. (2003), 'Persuasion bias, social influence, and unidimensional opinions', *Quarterly Journal of Economics* **118**(3), 909–968.

Di Francesco, T. and Hommes, C. (2024), 'Sentiment-driven speculation in financial markets with heterogeneous beliefs: a machine learning approach', *Working Paper* .

Dieci, R. and He, X. Z. (2018), 'Heterogeneous agent models in finance', *Handbook of Computational Economics* **4**.

Diks, C. and Dindo, P. (2008), 'Informational differences and learning in an asset market with boundedly rational agents', *Journal of Economic Dynamics and Control* **32**(5), 1432–1465.

Dyer, J., Cannon, P., Farmer, J. D. and Schmon, S. M. (2024), 'Black-box bayesian inference for agent-based models', *Journal of Economic Dynamics and Control* **161**, 104827.

Erceg, C. J., Henderson, D. W. and Levin, A. T. (2000), 'Optimal monetary policy with staggered wage and price contracts', *Journal of Monetary Economics* **46**(2), 281–313.

Evans, G. W. and Honkapohja, S. (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press, Princeton.

Evans, G. W. and Honkapohja, S. (2003), 'Adaptive learning and monetary policy design', *Journal of Money, Credit, and Banking* **35**, 1045–1072.

Fagereng, A., Holm, M. B. and Natvik, G. J. (2021), 'Mpc heterogeneity and household balance sheets', *American Economic Journal: Macroeconomics* **13**, 1–54.

Fair, R. C. and Taylor, J. B. (1983), 'Solution and maximum likelihood estimation of dynamic nonlinear rational expectations models', *Econometrica* **51**, 1169–1185.

Fama, E. F. (1970), 'Efficient capital markets: A review of theory and empirical work', *The Journal of Finance* **25**(2), 383–417.

Farhi, E. and Werning, I. (2019), 'Monetary policy, bounded rationality, and incomplete markets', *American Economic Review* **109**, 3887–3928.

Fisher, P. G., Holly, S. and Hallett, A. J. H. (1986), 'Efficient solution techniques for dynamic non-linear rational expectations models', *Journal of Economic Dynamics and Control* **10**, 139–145.

Franke, R. and Westerhoff, F. (2012), 'Structural stochastic volatility in asset pricing dynamics: Estimation and model contest', *Journal of Economic Dynamics and Control* **36**(8), 1193–1211.

Franke, R. and Westerhoff, F. (2016), 'Why a simple herding model may generate the stylized facts of daily returns: explanation and estimation', *Journal of Economic Interaction and Coordination* **11**(1), 1–34.

Frijns, B., Lehnert, T. and Zwinkels, R. C. (2010), 'Behavioral heterogeneity in the option market', *Journal of Economic Dynamics and Control* **34**, 2273–2287.

Gabaix, X. (2020), 'A behavioral new keynesian model', *American Economic Review* **110**, 2271–2327.

Gale, D. and Kariv, S. (2003), 'Bayesian learning in social networks', *Games and Economic Behavior* **45**(2), 329–346.

Gallegos-Dago, J. E. (2023*a*), 'Hank beyond fire: Amplification, forward guidance and belief shocks', *SSRN Electronic Journal* .

Gallegos-Dago, J. E. (2023*b*), 'Inflation persistence, noisy information and the phillips curve', *SSRN Electronic Journal* .

Gardini, L., Radi, D., Schmitt, N., Sushko, I. and Westerhoff, F. (2022), 'Causes of fragile stock market stability', *Journal of Economic Behavior and Organization* **200**, 483–498.

Gardini, L., Radi, D., Schmitt, N., Sushko, I. and Westerhoff, F. (2025), 'On boom-bust stock market dynamics, animal spirits, and the destabilizing nature of temporarily attracting virtual fixed points', *Macroeconomic Dynamics* **29**, e35.

Gerotto, L., Pellizzari, P. and Tolotti, M. (2019), Asymmetric information and learning by imitation in agent-based financial markets, *in* F. De La Prieta, A. González-Briones, P. Pawleski, D. Calvaresi, E. Del Val, F. Lopes, V. Julian, E. Osaba and R. Sánchez-Iborra, eds, 'Highlights of Practical Applications of Survivable Agents and Multi-Agent Systems. The PAAMS Collection', Springer International Publishing, Cham, pp. 164–175.

Gilbert, T., Kogan, S., Lochstoer, L. and Ozyildirim, A. (2012), 'Investor inattention and the market impact of summary statistics', *Management Science* **58**(2), 336–350.

Golub, B. and Jackson, M. O. (2010), 'Naive learning in social networks and the wisdom of crowds', *American Economic Journal: Microeconomics* **2**(1), 112–149.

Goncalves, S. and Kilian, L. (2004), 'Bootstrapping autoregressions with conditional heteroskedasticity of unknown form', *Journal of Econometrics* **123**, 89–120.

Goodfellow, I. J., Bengio, Y. and Courville, A. (2016), *Deep Learning*, MIT Press, Cambridge, MA, USA.

Grossman, S. J. and Stiglitz, J. E. (1980), 'American economic association on the impossibility of informationally efficient markets', *The American Economic Review* **70**(3), 393–408.

Guerreiro, J. (2023), 'Belief disagreement and business cycles', *Working Paper* .

Gurdgiev, C. and O'Loughlin, D. (2020), 'Herding and anchoring in cryptocurrency markets: Investor reaction to fear and uncertainty', *Journal of Behavioral and Experimental Finance* **25**, 100271.

Guégan, D. and Renault, T. (2021), 'Does investor sentiment on social media provide robust information for bitcoin returns predictability?', *Finance Research Letters* **38**, 101494.

Hamilton, J. D. (1994), *Time Series Analysis*, Princeton University Press.

Hamilton, J. D. (2018), 'Why you should never use the hodrick-prescott filter', *The Review of Economics and Statistics* **100**(5), 831–843.

Harris, M. and Raviv, A. (1993), 'Differences of opinion make a horse race', *Review of Financial Studies* **6**, 473–506.

He, X. Z. and Li, K. (2012), 'Heterogeneous beliefs and adaptive behaviour in a continuous-time asset price model', *Journal of Economic Dynamics and Control* **36**, 973–987.

Heemeijer, P., Hommes, C., Sonnemans, J. and Tuinstra, J. (2009), 'Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation', *Journal of Economic Dynamics and Control* **33**(5), 1052–1072.

Herman, J. and Usher, W. (2017), 'SALib: An open-source python library for sensitivity analysis', *The Journal of Open Source Software* **2**(9), 97.

Hilliard, J. E. and Ngo, J. T. D. (2022), 'Bitcoin: jumps, convenience yields, and option prices', *Quantitative Finance* **22**(11), 2079–2091.

Hirshleifer, D., Lim, S. S. and Teoh, S. H. (2009), 'Driven to distraction: Extraneous events and underreaction to earnings news', *Journal of Finance* **64**(5), 2289–2325.

Holland, P. W., Laskey, K. B. and Leinhardt, S. (1983), 'Stochastic blockmodels: First steps', *Social Networks* **5**(2), 109–137.

Hommes, C. (2021), 'Behavioral and experimental macroeconomics and policy analysis: A complex systems approach', *Journal of Economic Literature* **59**, 149–219.

Hommes, C., Huang, H. and Wang, D. (2005), 'A robust rational route to randomness in a simple asset pricing model', *Journal of Economic Dynamics and Control* **29**, 1043–1072.

Hommes, C. and in 't Veld, D. (2017), 'Booms, busts and behavioural heterogeneity in stock prices', *Journal of Economic Dynamics and Control* **80**, 101–124.

Hommes, C., Sonnemans, J., Tuinstra, J. and Velden, H. V. D. (2005), 'Coordination of expectations in asset pricing experiments', *Review of Financial Studies* **18**, 955–980.

Hommes, C. and Sorger, G. (1998), 'Consistent expectations equilibria', *Macroeconomic Dynamics* **2**(3), 287–321.

Hommes, C., Sorger, G. and Wagener, F. (2013), Consistency of linear forecasts in a nonlinear stochastic economy, *in* G. I. Bischi, C. Chiarella and I. Sushko, eds, 'Global Analysis of Dynamic Models in Economics and Finance: Essays in Honour of Laura Gardini', Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 229–287.

Hommes, C. and Zhu, M. (2014), 'Behavioral learning equilibria', *Journal of Economic Theory* **150**(C), 778–814.

Hong, H. and Stein, J. C. (2003), 'Differences of opinion, short-sales constraints, and market crashes', *Review of Financial Studies* **16**, 487–525.

Hornik, K., Stinchcombe, M. and White, H. (1989), 'Multilayer feedforward networks are universal approximators', *Neural Networks* **2**, 359–366.

Huberman, G. and Regev, T. (2001), 'Contagious speculation and a cure for cancer: A nonevent that made stock prices soar', *Journal of Finance* **56**(1), 387–396.

Hutto, C. J. and Gilbert, E. (2014), 'Vader: A parsimonious rule-based model for sentiment analysis of social media text', *International Conference on Web and Social Media* .

Iori, G. (2002), 'A microsimulation of traders activity in the stock market: The role of heterogeneity, agents' interactions and trade frictions', *Journal of Economic Behavior and Organization* **49**(2), 269–285.

Jegadeesh, N. and Titman, S. (1993), 'Returns to buying winners and selling losers: Implications for stock market efficiency', *Journal of Finance* **48**, 65–91.

J.Parra-Moyano, D.Partida, Gessl, M. and S.Mazumdar (2024), 'Analyzing swings in bitcoin returns: a comparative study of the lppl and sentiment-informed random forest models', *Digital Finance* **6**, 427–439.

Kanoria, Y. and Tamuz, O. (2013), 'Tractable bayesian social learning on trees', *IEEE Journal on Selected Areas in Communications* **31**(4), 2721–2725.

Kaplan, G., Moll, B. and Violante, G. L. (2018), 'Monetary policy according to hank', *American Economic Review* **108**, 697–743.

Kenny, G. and Morgan, J. (2011), Some lessons from the financial crisis for the economic analysis, Occasional Paper Series 130, European Central Bank.

Khashanah, K. and Alsulaiman, T. (2016), 'Network theory and behavioral finance in a heterogeneous market environment', *Complexity* **21**.

Kukacka, J. and Barunik, J. (2017), 'Estimation of financial agent-based models with simulated maximum likelihood', *Journal of Economic Dynamics and Control* **85**, 21–45.

Kukacka, J. and Kristoufek, L. (2023), 'Fundamental and speculative components of the cryptocurrency pricing dynamics', *Financial Innovation* **61**(9).

Kyle, A. S. (1985), 'Continuous auctions and insider trading', *Econometrica* **53**(6), 1315–1335.

Liu, Y. and Tsyvinski, A. (2020), 'Risks and returns of cryptocurrency', *Review of Financial Studies* **34**, 2689–2727.

Liu, Y., Tsyvinski, A. and Wu, X. (2022), 'Common risk factors in cryptocurrency', *The Journal of Finance* **77**, 1133–1177.

Lof, M. (2015), 'Rational speculators, contrarians, and excess volatility', *Management Science* **61**, 1889–1901.

Lucas, R. E. (1972), 'Expectations and the neutrality of money', *Journal of Economic Theory* **4**(2), 103–124.

Lucas, R. E. and Prescott, E. C. (1971), 'Investment under uncertainty', *Econometrica* **39**(5), 659–681.

Lux, T. (1998), 'The socio-economic dynamics of speculative markets: Interacting agents, chaos, and the fat tails of return distributions', *Journal of Economic Behavior and Organization* **33**(2), 143–165.

Lux, T. (2009), 'Rational forecasts or social opinion dynamics? identification of interaction effects in a business climate survey', *Journal of Economic Behavior and Organization* **72**(2), 638–655.

Madeira, C. and Zafar, B. (2015), 'Heterogeneous inflation expectations and learning', *Journal of Money, Credit and Banking* **47**, 867–896.

Makarov, I. and Schoar, A. (2020), 'Trading and arbitrage in cryptocurrency markets', *Journal of Financial Economics* **135**, 293–319.

Mankiw, N. G. and Reis, R. (2002), 'Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve', *Quarterly Journal of Economics* **117**.

Mankiw, N. G., Reis, R. and Wolfers, J. (2004), Disagreement about Inflation Expectations, *in* 'NBER Macroeconomics Annual 2003, Volume 18', NBER Chapters, National Bureau of Economic Research, Inc, pp. 209–270.

Manski, C. F. (2004), 'Measuring expectations', *Econometrica* **72**(5), 1329–1376.

Manski, C. F. (2018), 'Survey measurement of probabilistic macroeconomic expectations: Progress and promise', *NBER Macroeconomics Annual* **32**, 411–471.

Marcet, A. and Sargent, T. J. (1989*a*), 'Convergence of least-squares learning in environments with hidden state variables and private information', *Journal of Political Economy* **97**.

Marcet, A. and Sargent, T. J. (1989*b*), 'Convergence of least squares learning mechanisms in self-referential linear stochastic models', *Journal of Economic Theory* **48**.

Maćkowiak, B. and Wiederholt, M. (2009), 'Optimal sticky prices under rational inattention', *American Economic Review* **99**, 769–803.

Maćkowiak, B. and Wiederholt, M. (2010), 'Business cycle dynamics under rational inattention', *Review of Economic Studies* **82**, 1502–1532.

McKay, A., Nakamura, E. and Steinsson, J. (2016), 'The power of forward guidance revisited', *American Economic Review* **106**, 3133–3158.

Mckay, A. and Wolf, C. K. (2022), 'Optimal policy rules in hank', *Working Paper* .

Muth, J. F. (1961), 'Rational expectations and the theory of price movements', *Econometrica* **29**(3), 315–335.

Nerlove, M. (1958), 'Adaptive expectations and cobweb phenomena', *The Quarterly Journal of Economics* **72**(2), 227–240.

Nord, L. (2022), 'Who cares about inflation? endogenous expectation formation of heterogeneous households', *SSRN Electronic Journal* .

Orléan, A. (1995), 'Bayesian interactions and collective dynamics of opinion: Herd behavior and mimetic contagion', *Journal of Economic Behavior and Organization* **28**(2), 257–274.

Panchenko, V., Gerasymchuk, S. and Pavlov, O. V. (2013), 'Asset price dynamics with heterogeneous beliefs and local network interactions', *Journal of Economic Dynamics and Control* **37**(12), 2623–2642.

Papamakarios, G. and Murray, I. (2016), Fast epsilon-free inference of simulation models with bayesian conditional density estimation, *in* D. Lee, M. Sugiyama, U. Luxburg, I. Guyon and R. Garnett, eds, 'Advances in Neural Information Processing Systems', Vol. 29, Curran Associates, Inc.

Peng, L. and Xiong, W. (2006), 'Investor attention, overconfidence and category learning', *Journal of Financial Economics* **80**(3), 563–602.

Romer, C. D. and Romer, D. H. (2004), 'A new measure of monetary shocks: Derivation and implications', *American Economic Review* **94**, 1055–1084.

Rusinowska, A. and Taalaibekova, A. (2019), 'Opinion formation and targeting when persuaders have extreme and centrist opinions', *Journal of Mathematical Economics* **84**, 9–27.

Saltelli, A. (2002), 'Making best use of model evaluations to compute sensitivity indices', *Computer Physics Communications* **145**(2), 280–297.

Saltelli, A. (2010), 'Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index', *Computer Physics Communications* **181**(2), 259–270.

Sargent, T. J. (1991), 'Equilibrium with signal extraction from endogenous variables', *Journal of Economic Dynamics and Control* **15**(2), 245–273.

Sargent, T. J. (2008), 'Evolution and intelligent design', *American Economic Review* **98**(1), 5–37.

Schmitt-Grohé, S. and Uribe, M. (2005), 'Optimal fiscal and monetary policy in a medium-scale macroeconomic model', *NBER Macroeconomics Annual* **20**.

Schmitt, N. (2021), 'Heterogeneous expectations and asset price dynamics', *Macroeconomic Dynamics* **25**, 1538–1568.

Schmitt, N. and Westerhoff, F. (2017), 'Heterogeneity, spontaneous coordination and extreme events within large-scale and small-scale agent-based financial market models', *Journal of Evolutionary Economics* **27**(5), 1041–1070.

Shabalina, E. and Tzaawa-Krenzler, M. (2025), 'Heterogeneous attention to inflation and monetary policy', *Manuscript* .

Shalen, C. T. (1993), 'Volume, volatility, and the dispersion of beliefs', *Review of Financial Studies* **6**, 405–434.

Shen, D., Urquhart, A. and Wang, P. (2019), 'A three-factor pricing model for cryptocurrencies', *Finance Research Letters* **34**, 101248.

Sims, C. A. (2003), 'Implications of rational inattention', *Journal of Monetary Economics* **50**, 665–690.

Smets, F. and Wouters, R. (2007), 'Shocks and frictions in us business cycles: A bayesian dsge approach', *American Economic Review* **97**, 586–606.

Sobol, I. M. (2001), 'Global sensitivity indices for nonlinear mathematical models and their monte carlo estimates', *Mathematics and Computers in Simulation* **55**(1-3), 271–280.

Tedeschi, G., Iori, G. and Gallegati, M. (2012), 'Herding effects in order driven markets: The rise and fall of gurus', *Journal of Economic Behavior and Organization* **81**(1), 82–96.

ter Ellen, S., Hommes, C. H. and Zwinkels, R. C. (2021), 'Comparing behavioural heterogeneity across asset classes', *Journal of Economic Behavior and Organization* **185**, 747–769.

Tetlock, P. C. (2011), 'All the news that's fit to reprint: Do investors react to stale information?', *Review of Financial Studies* **24**(5), 1481–1512.

Townsend, R. M. (1978), 'Market anticipations, rational expectations, and bayesian analysis', *International Economic Review* **19**.

Tramontana, F., Westerhoff, F. and Gardini, L. (2010), 'On the complicated price dynamics of a simple one-dimensional discontinuous financial market model with heterogeneous interacting traders', *Journal of Economic Behavior and Organization* **74**(3), 187–205.

Tramontana, F., Westerhoff, F. and Gardini, L. (2013), 'The bull and bear market model of huang and day: Some extensions and new results', *Journal of Economic Dynamics and Control* **37**(11), 2351–2370.

Urquhart, A. (2016), 'The inefficiency of bitcoin', *Economics Letters* **148**, 80–82.

Urquhart, A. (2018), 'What causes the attention of bitcoin?', *Economics Letters* **166**, 40–44.

Vozlyublennaia, N. (2014), 'Investor attention, index performance, and return predictability', *Journal of Banking and Finance* **41**(1), 17–35.

Waters, G. (2019), 'Bubbles and rationality in bitcoin', *Economic Notes* **48**(2), e12133.

Watts, D. J. and Strogatz, S. H. (1998), 'Collective dynamics of 'small-world' networks', *Nature* **393**(6684), 440–442.

Weber, M., D'Acunto, F., Gorodnichenko, Y. and Coibion, O. (2022), 'The subjective inflation expectations of households and firms: Measurement, determinants, and implications', *Journal of Economic Perspectives* **36**(3), 157–84.

Wolf, C. (2021), 'Interest rate cuts vs. stimulus payments: An equivalence result', *SSRN Electronic Journal* .

Wu, S., He, J. and Li, S. (2018), 'Effects of fundamentals acquisition and strategy switch on stock price dynamics', *Physica A: Statistical Mechanics and its Applications* **491**, 799–809.

Zorn, P. (2021), 'Investment under rational inattention: Evidence from us sectoral data', *SSRN Electronic Journal* .