Essays on Experimental and Behavioural Economics and Finance

Aleksei Chernulich

A Doctoral Thesis
Presented to the Faculty
of University of Technology Sydney
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Economics
Supervisor: Professor Mikhail Anufriev

December 2018

Abstract

Three essays of this thesis combine research projects devoted to markets, prices, and expectations. The first chapter provides results on an experimental test of a model with interacting boundedly rational agents. Adaptive switching between forecasting heuristics by heterogeneous agents brings instability to the price dynamics and generates bubbles and crashes. In the second chapter of this thesis, behavioural models of channelling attention to adaptive choice are empirically tested on data generated from laboratory experiments. According to the identified self-tuning model, subjects scale their attention to the task given the stakes. Computational analysis and simulations demonstrate the importance of this self-tuning model for generating price dynamics that balances on the edge of stability. The third chapter is an experimental investigation of the role of forecasting horizon length in generating excess price volatility. In markets with initially unstable prices with an increase in horizon length price dynamics stabilises. This finding can be partly explained by dis-coordination of subject on non-fundamental expectations in markets with longer horizons.

Certificate of Original Authorship

I, Aleksei Chernulich, declare that this thesis, is submitted in fulfilment of the requirements for the award of PhD in Economics, in the UTS Business School at the University of Technology Sydney.

This thesis is wholly my own work unless otherwise reference or acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis. This document has not been submitted for qualifications at any other academic institution. This research is supported by the Australian Government Research Training Program.

Signature: Production Note: Signature removed prior to publication.

Date: 2019.04.13

Acknowledgements

The presented thesis combines three papers with the results of three different, but interconnected, research projects. My work on such a diverse research agenda was only possible because of the help and guidance that I have received from my supervisors, colleagues, and friends, who oftentimes are the same people.

I would like to thank my principal supervisor, Professor Mikhail Anufriev, who invited me to join the University of Technology Sydney (UTS) and to work with him on several projects. I received support, guidance, as well as the freedom to work on my own projects. I would like to thank my informal co-supervisor and co-author Professor Jan Tuinstra, who invited me to work with him in Amsterdam on several occasions. I would like to thank my co-supervisor, Associate Professor David Goldbaum, who always has a minute to discuss my research for a couple of hours. I would like to thank Dr. Corrado Di Guilmi for his feedback on preliminary versions of my papers, which were more like preliminary versions of preliminary versions. I would like to thank my co-supervisor, Professor John Wooders, for both his support and very healthy scepticisms and criticisms, which are oftentimes very necessary. I would like to thank Professor Isa Hafalir, who helped me develop existing projects and start new ones.

I thank UTS for financial and academic support, the Australian Government for 'Research Training Program' provided and Australian Research Council for funding–I am grateful for financial support from the Australian Research Council's Discovery Projects funding scheme (Project Number DP140103566, Wooders). I thank professional editor Abigail Bergman for excellent editing services.

The Economics Discipline Group at UTS is a great team of researchers, who at different times contributed to my professional development and the development of my thesis. Without mentioning academic ranks and further details, I would like to thank Susumo Imai, Jacob Goeree, Philippos Louis, Youngki Shin, Antonio Rosato, Shiko Maruyama, Ming He, Toru Suzuki, Jun Zhang, Olena Stavrunova, Jinjgjing Zhang, Jonathan Livermore, Sergey Alexeev, Yanlin Chen, Evgeniya Goryacheva, Eamon McGinn, Mengheng Li, Esther Mirjam Girsberger, Joshua Chan, Mario Fiorini, Elif Incekara Hafalir, Mark Chan, Gordon Menzies, Massimo Scotti, Peter Siminski, Emil Temnyalov, Kentaro Tomoeda, Junji Xiao, Benjamin Balzer, Xin Zhao, Benjamin Young, Andrea Giovannetti, Lakmali Abeyesinghe, Carole Fawcett, Sally Browning, Nirada Manosorn, Inês Soares and Natascha McQuilty.

There are many people who helped me at different stages of the thesis development, but there is one person who has constantly, and understandably, helped me to finish the thesis at the earliest possible date. I am exceptionally grateful to my wife, Lera, who at this point knows about experiments, markets, and bounded rationality no less than I do.

To my mom, Elena—my main advisor throughout life—and to my dad, Victor—my main example to follow.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Overview

Three chapters of this thesis combine research on the intersection of the economics and finance disciplines and are focused on markets, prices, and expectations. Markets are important mechanisms of allocations, where efficiency of the distribution critically depends on the prices. Understanding the price dynamics and origins of possible deviations of the price from the fundamental values, known as bubbles and crashes, is especially important for public finance in terms of developing appropriate stabilisation policies and tailoring tax systems. To investigate the interplay between prices and price expectations, this thesis uses experimental and behavioural economics tools. Computational methods and simulations are utilised to study possible effects of behavioural expectation formation on the market prices.

1.2 Market Prices and Price Expectations

Economists study different kinds of markets including those for labour, goods, and services. However, the 2008 Financial Crisis demonstrated both the importance of financial markets to a global economy and our limited understanding of these markets. All major economies were harmed, growth rates dropped to negative values, and the global markets' capitalisation was reduced by half, yet only a few economists were able to identify that the financial markets were in disequilibrium before the fall. Our incomplete knowledge about financial markets is not surprising given the complexities of the system: millions of transactions are made daily within and between markets, reflecting changes in various observed and non-observed fundamental indicators, such as changes in the composition of traders and their beliefs, and changes in policy, structure, and technology. To address the issue of complexity, the research presented in this thesis breaks down the task of understanding the market forces into more tractable components and, in particular, focuses attention on the interplay between prices and price expectations in a stylised version of a market populated by agents who behave adaptively rather than rationally.

Expectations about future prices are essential drivers of markets. Prices reflect expectations of the participants of the market, and expectations are often formed based on observed prices. These mutual feedback effects between prices and price expectations play an important role in the presented analysis since agents do not always behave fully rationally and their expectations are often boundedly rational. Once the assumption about the full rationality is dropped, agents no longer have a "perfect foresight" and alternative models of forecasting become available.

1.3 Bounded Rationality

Rational Expectation Hypothesis (REH)–introduced and advocated by [78]–remains a corner stone of modern economics and, in particular, macroeconomics. Economic agents are modelled to rationally analyse the available information to form expectations, and act optimally based on these expectations, which in the end become self-fulfilling. This perfect foresight is a result of a combination of individual rationality and mutual consistency in the behaviour of agents. But the underlying assumptions of REH with regard to computational ability and informational availability to the agents appears to be very restrictive, as discussed in [98]. In fact, expectations that are observed in reality are shown to be non-rational in [77]. Models of bounded rationality, where perfect foresight is replaced by adaptive expectations and learning, are discussed in [91], and applied in a macroeconomic dynamics setup in [43]. In order to compromise the tendency of economic agents to optimise their behaviour with the inability to perfectly predict the future dynamic, agents are modelled to be sophisticated econometricians studying the data in order to form a prediction, following [99]. This approach of modelling agents to be "fundamental" analysers is extended by introducing a cost for information and computation. This informational cost, which may motivate agents to be less rational, can explain boundedly rational behaviour, and plays an important role in the model introduced in the second chapter of this thesis. For the complex investigation of the presented issues, instruments from the two relatively young fields—behavioural finance and experimental economics—are employed.

1.4 Experimental and Behavioural Approaches

Experimental and behavioural economics approaches are often used in conjunction. Lab-based experiments with human subjects are a natural and convenient way to test possible behavioural models of economic agent. Equally, any behaviour observed in the lab that is not consistent with the "standard" economic theory serves as empirical evidence for building behavioural models. Both fields have already attracted many followers and the importance of the developments was acknowledged in recent years by awarding Nobel Prizes to the most prominent researchers in these fields.

Behavioural economics and finance are typically said to date back to the works of Amos Tversky and Daniel Kahneman, where integration of psychological insights on decision-making into standard economic context was performed. The work of Tversky and Kahneman sheds light on the existence of "bounded rationality" and heuristics, which are simple rules used by people to make decisions in complex environments. In the presented thesis, heuristics are applied in the context of the price forecasting task: agents are modelled to be "boundedly rational" and use simple rules of thumb to predict future price movements. Robert Schiller attracted great attention to behavioural finance contributing with his empirical analysis of asset prices with a focus on the possible behavioural origins of market inefficiencies. The excess volatility phenomena, identified by Robert Schiller, serves as a motivation for the third chapter of the thesis, where the potential importance of forecasting horizons for market price stability is investigated. Richard H. Thaler contributed to the development of behavioural economics by advancing the study of behavioural biases and, in particular, developing policy recommendations that account for existing biases in decision-making. The policy implications are discussed in all chapters of the thesis.

Experimental economics offers a number of tools for detailed examination of both the individual and group behaviour of participants in the controlled laboratory environment. One of the earliest contributions was made by [100], who ran a canonical asset market experiment that generated bubbles and crashes. Lab experiments were designed and conducted as a part of the first and the third chapters of the thesis: individual binary choice task with the limited information experiment, and stylised financial market experiment, respectively. Analysis of the clean experimental data gives us a better understanding of both individual behaviour and market institutions.

1.5 Heterogeneous Agent Models

Three main chapters of the thesis all have close connections to the literature of Heterogeneous Agent Modelling (HAM), and, in particular, Heuristic Switching Models (HSM). There are several features in the setup of these models that are important for the analysis and generated predictions of HAM papers. First, the assumption about the fully rational representative "neo-classical" economic agent is dropped. Instead, heterogeneous "boundedly rational" agents, interacting in the market, are introduced. Second, the mutual feedback effects between prices and price expectations materialise through demand and supply forces and play a crucial role in the price discovery process. Prices reflect expectations of the participants of the market, and expectations are formed based on the observed prices using forecasting heuristics. Agents' adaptive switching between different heuristics can bring instability to the price dynamics and generate "stylised facts" of financial markets: bubbles and crashes, fat tails, and volatility clustering. Seminal papers by Brock and Hommes laid a theoretical foundation that sparked numerous HSM-based papers with computational analysis, empirical studies of the financial data, and experimental investigations. A simple setup that can generate complicated dynamics became an ideal instrument for both modelling and explaining the turbulence of financial markets.

Theoretical development of the HAMs dates back to two seminal papers by [17, 18]. These papers introduce a model with markets that are populated by boundedly rational agents. Instead of having "neoclassical" rational expectations, agents form expectations by using one of the two archetypical rules: sophisticated fundamental analysis and naïve trend following. This simplifying assumption rests on evidence from numerous studies that identify these two forces to be the main powers driving the "moods" of the markets (see [35], [1], and [48], among many others). Agents' adaptive choice between forecasting rules is a key ingredient of the model. In the second and third chapters of the thesis, I closely scrutinize different behavioural models of adaptive choice. Adaptive choice between forecasting rules is usually modelled with the logit response function: agents choose probabilistically based on the observed previous payoffs generated by the rules. The logit parameter of the model, which represents the Intensity of Choice (IoC), and the cost of fundamental prediction, which represents informational rent, are responsible for generating price instability. As the value of either of two parameters increases, the model predicts the price to diverge from the steady state and fluctuate chaotically. The second chapter provides experimental tests regarding the role of information cost in generating price instability, and the third chapter investigates how price instability may affect the IoC parameter values.

Ability of the HSMs to generate various stylised facts of financial markets–including excess volatility, fat tails, and volatility clustering—attracts attention from finance researchers. Several HSM modifications receives empirical support based on the data from different sources. The S&P500 data is used to estimate a model of switching between fundamentalist and chartist strategy in [31]. Annual US stock price data is used in [13], where adaptive switching between fundamentalists and trend-followers was shown to play a significant role. Survey data of expectations for foreign exchange markets is used in [53] to estimate a model of switching between heterogeneous expectations. Data from the Survey of Consumer Attitudes and Behaviour is used in [15], where a model of switching between three predictors—naive, adaptive, and sophisticated—was estimated. Overall, this branch of the HSM literature provides extensive evidence of the empirical validity of HSMs and the variety of estimates of the IoC parameter.

What remains a main challenge in the literature is a high level of heterogeneity in the IoC parameter estimates. The IoC plays a crucial role in the model mechanics because different values of the parameter model generate qualitatively different predictions over the price dynamics: from the steady state to deterministic chaos. Diversity of estimates, which were obtained from the financial markets data, imposes serious calibration issues, and raises questions concerning the validity of the assumption that the IoC parameter is constant. The second chapter provides additional experimental evidence on variability of the parameter, and the third chapter proposes a behavioural model that endogenises the changes in the IoC parameter values.

1.6 Learning-to-Forecast Experiments

In addition to empirical tests of HSMs based on the existing data sets, several economic experiments with human subjects were run to study the interplay between prices and price expectations in a controlled laboratory environment. In these experiments, most of the real market components, including trading, were isolated from the experimental setup to avoid mixing confounding factors. The only task for the participants was to predict the future price given the observed price patterns, while the price discovery process depended on the given forecasts. This experimental design is different from experiments in [92], [59], and [73], where prices are pre-generated, and is instead based on the setup introduced in [79]. A series of experiments, which are referred to as Learning to Forecast (LtF) experiments, provided an important insight on both individual forecasting behaviour and market price equilibration processes. The main contributions include those from [69] and [70]. In these LtF experiments, participants generated substantial price bubbles and subsequent crashes by forecasting price one or two periods ahead. Observed phenomena, to a large extent, can be explained by the tendency of subjects to extrapolate a price pattern and reinforce price deviations rather than calculate fundamental values and stabilise the market. An important side product of indirect comparisons of these experiments suggests that markets, where participants had to predict two periods ahead instead of one period ahead, exhibited higher levels of price volatility. This observation motivates the fourth chapter of the thesis, where a new experiment is explicitly designed to study the effects of increasing the length of the forecasting horizon on price volatility.

1.7 Research Questions

This thesis utilises approaches of behavioural and experimental economics and combines it with more standard theoretical, empirical, and computational analysis to answer a number of questions related to both economics and finance disciplines. Can some stylised facts of financial markets, such as the excess volatility or bubbles and crashes, be explained by behavioural models of expectation formation? In particular, can stylised facts of financial markets be explained by models with agents switching between simple forecasting strategies or heuristics, such as fundamentalists versus chartists? Can a parsimonious model of switching between profitable alternatives explain subjects' behaviour in both stable and unstable environments? Does the forecasting horizon of agents in the market affect price volatility, for example, through more aggressive trend-extrapolating behaviour? In a broad sense, all three chapters are devoted to the issues related to the mutual feedback loops that exist in the real markets between prices and price expectations, and between agents' adaptive behaviour and environmental changes in response to this adaptation.

1.8 Thesis Structure

This thesis is organised as follows. The second chapter is an experimental paper that tests the predictions of the stylised version of the seminal work by Brock and Hommes.[1] In that seminal paper, HSM agents adaptively switch between two forecasting heuristics: one is naïve but free and the other is

[1] This chapter is a joint work with Professor Mikhail Anufriev and Jan Tuinstra. The paper is published in the Journal of Economic Dynamics & Control in 2018.

rational but costly. If the cost of rational expectation is high and prices are stable, agents prefer to be naïve and use a previously observed price as their best predictor. The popularity of a naïve prediction destabilises the market and makes rational forecasting more attractive despite the cost to be paid. Experimental findings support the model predictions regarding the importance of the information cost that is paid by the rational market participants for acquiring information about market fundamentals–as the cost increases, the price dynamics become less stable. This intuitive result rests on the individual decisions of experiment participants, which the simple adaptive choice model explains well. But there is an important caveat: the parameter of the "attention" or "focus" of participants is estimated to be different in different experimental sessions. The third chapter of this thesis studies the attention adjustment processes and their importance for adaptive behaviour. The paper collects "stylised facts" of previously identified differences in attention and builds a number of behavioural models to capture these effects. The study utilises data from several laboratory experiments on adaptive choice and applies econometric techniques of running contests between different models using the data from the meta-experiment. Results show that participants adjust their attention to the choice task given the stakes: if payoffs are of a similar value, participants place less emphasis on the choice task. This finding explains why attention is different in different experimental sessions and offers new insight into HSM dynamics where adaptive switching plays a crucial role. The fourth chapter experimentally investigates one of the possible sources of the excess volatility of the price: the length of the forecasting horizon.[2] Theoretically, an increase in the forecasting horizon leads to stability due to weakening arbitrage motives. Results

[2]This chapter is a joint work with Professor Mikhail Anufriev and Jan Tuinstra.

show that predictions do not always hold because of behavioural effects that are absent in the standard model—participants tend to extrapolate the price trends in medium-long horizons, which destabilises the market price.

Chapter 2

An experiment on the Heuristic Switching Model

2.1 Introduction

In the last couple of decades *heterogeneous agents models* have become increasingly popular as a description of turbulence and volatility on financial markets. In these models different types of traders coexist on financial markets, as motivated by early empirical and theoretical studies. [48], for example, distinguish between *fundamentalists*, who use in-depth analysis of firms and their market environment to determine the fundamental value of an asset and believe that prices tend to revert back to this fundamental value, and *chartists*, who use technical analysis to identify patterns in prices and extrapolate those when predicting future prices (also see [1]). [35] and [25] show that the interaction between these different types of traders may lead to the emergence of endogenous fluctuations in asset prices in an otherwise stationary and deterministic environment.

An important class of heterogeneous agents models assumes that there is a large population of traders that adaptively switch between some archetypal types of behavior, or heuristics, on the basis of the relative performance of these heuristics, see, e.g., [17, 18], [28, 29] and [4, 5]. These models, often referred to as *heuristic switching models*, have been successful in describing stylized facts of financial markets, such as bubbles and crashes in asset prices, excess volatility, volatility clustering, and fat tails, and have become quite popular as a result. The most common approach of modeling switching between heuristics in this literature is through the so-called discrete choice model. A crucial role in that model is played by the *Intensity of Choice* (IoC) parameter, which measures how sensitive traders are with respect to differences in performance. [17, 18] showed, and follow-up research confirmed, that heuristic switching models generate excess volatility and many other stylized facts of financial markets when the IoC parameter is *sufficiently large*. The IoC parameter is thus pivotal for the dynamic properties of the heuristic switching models, and hence for their success in explaining financial market data and for the validity of policy implications.

In this chapter we present the results of a laboratory experiment that is designed to test the predictions of the heuristic switching model and to estimate the IoC parameter from aggregate decisions of the experimental subjects. In particular, we construct a decision environment which is a stylized version of the framework laid down in the seminal work by [17, 18]. This framework has served as a benchmark heuristic switching model in the literature[1] but thus far its basic assumptions and implications have never been tested in a controlled laboratory experiment.

In our experiment we let subjects choose between a costly stabilizing heuristic and a cheap destabilizing heuristic. Payoffs associated with the choices depend upon the distribution of subjects over the heuristics. An important prediction of the heuristic switching model is that if costs for the stabilizing heuristic increase (relative to the costs for the destabilizing heuristic), the dynamics of the state variable (e.g., the price) and the distribution of agents over heuristics become unstable and complicated endogenous fluctuations may emerge. In our experiment we indeed find that an increase in costs for the stabilizing heuristic initially leads to the type of bubbles and crashes that are typical for the standard heuristic switching model. However, for the case of high costs we also find that, over time, the subjects adapt their behavior such that the dynamics becomes more stable and is consistent with a steady state of the model (although the dynamics are still much more volatile than in the low cost case). In particular, the estimated values of the intensity of choice parameter are much smaller for high costs than for low costs, which suggests that the characteristics of the economic environment are an important determinant of the value of this parameter. This finding is important because in the standard heuristic switching model the intensity of choice parameter is assumed to be exogenously given.

this chapter contributes to a growing empirical literature that tries to estimate the heuristic switching model from market or survey data. [52],

[1] According to Google Scholar [17] and [18] together have been cited more than 3100 times as of December 5, 2017. Examples of theoretical contributions that build upon their framework can be found in [50], [26], [19], [16], [27] and [88], among many others. A number of papers use the framework to study the effect of various financial market policies, such as the imposition of Tobin transaction taxes [112], increasing the number of financial derivatives [20] and restricting short selling [9].

for example, estimate the IoC parameter on mutual fund allocations decisions, whereas [15] uses survey data to estimate a discrete choice model with switching between three heuristics. [13] and [31] use U.S. stock price data to estimate a heuristic switching model with a fundamentalist heuristic and a trend-following or a chartist heuristic, respectively. Survey data of expectations in foreign exchange markets are used in [53]. Finally, [34] estimate a discrete choice model with switching between fundamentalists and random walk believers, using U.S. macroeconomic data. Although these contributions provide compelling evidence that the heuristic switching model performs well as a description of market behavior for different types of data, the drawback is that the IoC parameter in these models typically needs to be jointly estimated with other behavioral parameters (e.g., the heuristics are often parametrized as well). Several studies, including [13], find that the estimated value of the IoC parameter is not significant. Studies that produce significant estimates of the IoC report very different values, depending on the specification of the heuristics.[2] Finally, there is a large variation in the estimated value of the IoC parameter between studies.[3] This uncertainty regarding the relevant range of values for this parameter, which plays such a crucial role in the dynamics of heuristic switching models, may deter policy-makers from using those models.

[2]For instance, [52] use the risk adjusted payoffs of investors and estimate the IoC parameter as 0.9, 1.9, and 6.53 for three different values of the risk aversion parameter. [34] report values of the IoC parameter ranging from 1.99 to 9.26, depending on the forecasting rules used by fundamentalists.

[3]Because the IoC parameter is not a scale-free parameter the variation in the estimated values of this parameter may be partly due to the differences in (average) performance levels. This conjecture is, however, difficult to verify in empirical studies as performance levels of the different heuristics are rarely reported. Recently, [103] proposed a model where switching is based on relative, instead of absolute, profit differences. This opens up the possibility to compare the estimated IoC for various asset classes, such as metals, real estate prices, and foreign exchange markets, directly. The estimates of the IoC in [103] still differ across markets, however.

Laboratory experiments have the advantage that the experimenters can control both the market environment in which the subjects operate as well as the information these subjects have about that market environment. Therefore laboratory experiments can be used to obtain relatively clean data concerning the relationships under study. So-called 'Learning to Forecast' experiments (see, e.g., [69], [70] and [57], and [64] for an overview) have been used to estimate different forecasting heuristics that subjects typically use. An important feature of these Learning to Forecast experiments is that they take into account the self-referential characteristics of dynamic market environments. That is, expectations about the future value of an economic variable feed back into the actual realization of that variable, and thereby in part determine the behavior of that variable. This is in contrast to earlier experiments on expectation formation where participants have to predict an exogenously generated time series (see e.g., [92], [59] and [73]).

Experiments can also be used to better understand how people exactly *switch* between different heuristics. A first step in that direction, which can be seen as complementary to the Learning to Forecast experiments where the heuristics themselves are estimated, is taken in the experiment conducted in [2]. In that experiment subjects have to choose between different alternatives, where the payoffs associated with the alternatives are *exogenously* generated. The aggregate choice data are used to estimate different versions of the discrete choice model. The experiment presented in this chapter differs from the one in [2] in that payoffs are *endogenously* generated in this new experiment, and determined by the aggregate choices of the subjects. This setting therefore represents the typical heuristic switching model better.

The remainder of the chapter is organized as follows. We start by introducing a stylized version of the heuristic switching model in Section 2.2. This model is a simplified version of the original Brock-Hommes model and we use it as a basis for our experimental environment. In Section 2.3 we discuss the experimental design and in Section 2.4 we formulate the hypotheses to be tested. In Section 2.5 we present the experimental results. We fit a simple discrete choice model to subjects' aggregate choices and argue that the IoC parameter changes over time in response to the decision environment. In addition, we show that our results are robust to increasing the number of subjects in the experimental groups and to increasing the number of decision periods. Concluding remarks are given in Section 2.6.

2.2 A stylized heuristic switching model

This section introduces the stylized heuristic switching model that will be used in the laboratory experiment presented in this chapter. Some more background on the heuristic switching model and a description of the main mechanism that leads to complicated dynamics in that model are given in Section 2.2.1. We subsequently describe the stylized heuristic switching model that we use in this chapter and which consists of two parts: the dynamics of the state variable, which we discuss in Section 2.2.2, and the discrete choice model that describes switching between heuristics, which we discuss in Section 2.2.3.

2.2.1 Complex dynamics from the interaction of heuristics

The heuristic switching model was introduced in [17]. They consider a cobweb market where a large number of producers have to decide how much to supply of a (non-storable) commodity that takes one period to produce. The realized market price for the commodity will be the price for which consumer demand equals the aggregate supply of the producers, where the latter is determined by the individual price expectations of the producers. In [17] it is assumed that there are two forecasting heuristics available to the suppliers: either *naive expectations* – where a producer uses the last observed price as his prediction for the next price, or *rational expectations* – where a producer knows the underlying market equilibrium condition, as well as the distribution of producers over the different heuristics, and uses this information to compute the market clearing price in the current period. Whereas the naive expectations heuristic uses very little (and publicly available) information, applying the rational expectations heuristic is more demanding. It requires substantially more information as well as the cognitive effort to process this information correctly. It therefore seems reasonable to assume that using the rational expectations heuristic comes at strictly higher (information) costs than using the naive expectations heuristic.

Every period producers decide which forecasting heuristic to use on the basis of relative past performance (after information costs have been deducted). It turns out that complicated dynamics may emerge if the sensitivity of producers with respect to the profit difference (i.e., producers' intensity of choice) of the heuristics is sufficiently high. The mechanism underlying these complicated dynamics can be described as follows. When the price is close to its steady state level, the forecasting heuristics give comparable predictions and, given their relative cost, producers have the incentive to choose the naive expectations heuristic. However, when almost all producers use naive expectations, the cobweb dynamics are unstable and market prices start to oscillate (see, e.g., [44]). As a consequence, forecasting errors under naive expectations increase and producers tend to switch to the more profitable rational expectations heuristic. When enough producers use the rational expectations heuristic, prices converge to their steady state value and forecasting errors of both heuristics are similar again. As the rational expectations heuristic still comes with a cost, producers switch back to the naive expectations heuristic and the whole cycle repeats.

This type of mechanism, inducing complicated dynamics, works in many other market environments as well. Competition between fundamentalists and chartists in financial markets provides the most celebrated application, see, e.g., [18], [29] and [30] for early examples and [66] for an overview. It is important to note that the mechanism is qualitatively robust to changes in the set of heuristics, to the type of market institution (e.g., Walrasian equilibrium clearing, a market maker that adjusts prices, or continuous double auctions) and to the direction of the expectations feedback.[4] Generally, if the IoC parameter and/or the costs for the stabilizing heuristic are high enough, the fraction of agents using the stabilizing heuristic at the steady state is going to be low and will not be sufficient to stabilize the dynamics. The steady state will then be unstable and prices as well as the distribution of traders over the heuristics will fluctuate endogenously. The precise

[4]See [51], [8], and [63], respectively. Positive expectations feedback, where actual prices respond positively to an increase in the average expected price, is typical for financial markets, whereas negative expectations feedback, with a negative price response to an increase in price expectations, is common for supply-side driven markets, such as the cobweb market discussed above.

characteristics of the dynamics as well as the threshold values of the IoC parameter and the costs for the stabilizing heuristic will depend on the features of the underlying market environment.

2.2.2 Dynamics of the state variable and performance of heuristics

Our aim is to test the key assumptions and implications of the heuristic switching model described above in a laboratory experiment with paid human subjects. Laboratory experiments have the advantage that they generate clean data on choice behavior, that information given to the subjects, as well as the underlying model, are under control of the experimenters, and that subjects have well-defined monetary incentives. In addition, laboratory experiments can be replicated. At the same time, it is important to use decision environments that are not too complicated, in order to ensure that subjects have a good understanding of the task they are asked to perform. For this reason we want to use a *stylized* heuristic switching model that still exhibits the main mechanism of the models described above, but which is straightforward to implement in a laboratory experiment.

To that end, consider an economy where agents repeatedly choose between two profitable alternatives ('heuristics'), A and B. The profits generated by these alternatives are determined by an underlying state variable x_t. More specifically, profits of choosing A and B in period t are given by

$$\pi_{A,t} = W_A + \gamma_A x_t^2, \qquad \text{and} \qquad \pi_{B,t} = W_B - \gamma_B x_t^2\,, \tag{2.1}$$

respectively, where W_A, W_B, γ_A and γ_B are nonnegative parameters, with $\gamma_A + \gamma_B = 1$. Note that profits for alternative A depend positively on the deviation of the state variable from zero, and the other way around for alternative B.[5]

The evolution of the state variable x_t depends upon the distribution of agents over the alternatives A and B. In particular, we consider

$$x_t = \lambda n_{B,t} x_{t-1} + \epsilon_t \,, \tag{2.2}$$

where $n_{B,t} \in [0,1]$ is the fraction of the population of agents that chooses alternative B in period t, λ is the feedback coefficient and ϵ_t is a small idiosyncratic random shock. The sign of λ determines whether the feedback in the market is negative or positive, with $\lambda < 0$ representing the cobweb model of [17] and $\lambda > 0$ representing the asset pricing model of [18]. In both cases, the state variable x can be interpreted as the *deviation* of the price from its fundamental value.

The mechanism leading to endogenous fluctuations, discussed in Section 2.2.1 for the case of the cobweb model, but also relevant for financial markets, is preserved in the system consisting of equations (2.1) and (2.2), provided that $W_B > W_A$ and $|\lambda| > 1$. To see this, note that when the state variable x_t equals its steady state value of $x^* = 0$, alternative B is more profitable than alternative A (at this point $\pi_B = W_B > W_A = \pi_A$) and

[5] In the standard model only the *difference* between profits matters for agents, cf. Eq. (2.4) in Section 2.2.3. In our setup this difference is $\pi_{A,t} - \pi_{B,t} = W_A - W_B + x_t^2$, which, due to $\gamma_A + \gamma_B = 1$, does not depend on the exact values of γ_A and γ_B. For the experiment we choose the values of W_B and γ_B such that profits of alternative B remain positive for a relatively large range of values of x_t. This is done in order to minimize the number of periods in the experiment where payoffs for alternative B become negative. If in some period t the state variable x_t is so large that $\pi_{B,t} < 0$, we assign payoff 0 to alternative B for that period. This approach, which ensures that participating subjects have positive earnings, is common for laboratory experiments.

therefore attracts agents. Thus for low absolute values of the state variable, fraction $n_{B,t}$ will increase over time. However, when $n_{B,t}$ becomes large enough (in particular, when $|\lambda n_{B,t}| > 1$), the dynamics of the state variable become unstable and x_t will diverge away from its steady state value. When the deviation of x_t from zero becomes sufficiently large (to be specific, when $|x_t| > \sqrt{W_B - W_A}$) alternative A becomes more profitable and agents will tend to switch to that alternative again, which stabilizes the dynamics, and so on.

The model presented here, therefore, corresponds to a stylized version of the heuristic switching model described above, with alternative A playing the role of the costly stabilizing heuristic and alternative B playing the role of the cheap destabilizing heuristic. In addition, we can think of $W_B - W_A$ as the costs for heuristic A, because it corresponds to the difference in profits between the cheap heuristic B and the costly heuristic A, when $x_t = 0$, that is, at the steady state. In the remainder of this chapter we will therefore denote costs for using heuristic A (relative to the costs for using heuristic B) by $C = W_B - W_A$, which we will assume to be strictly positive.[6]

2.2.3 The discrete choice model

The model given by equations (2.1)–(2.2) still lacks a description of how agents choose between alternatives A and B when they know the past performance of alternatives (i.e., profits $\pi_{A,t}$ and $\pi_{B,t}$) but not the underlying

[6] It is worthwhile to stress one difference with the models from [17, 18]. In our stylized version the value of the state variable in period t depends upon $n_{B,t}$. This allows us to write the model as a simpler one-dimensional dynamical system, see Eq. (2.5), as opposed to the two-dimensional dynamical systems in [17, 18]. The dynamics, as well as the mechanism driving them, is qualitatively equivalent between the two settings. See Appendix A.1 for a formal analysis.

profit generating mechanism. The experiment presented in this chapter is designed specifically to investigate that decision, using human subjects. The most common approach to this choice problem in the literature on heuristic switching models is the so-called *discrete choice model*, which – in one of its most basic forms – looks as follows. Let $P_{A,t}$ and $P_{B,t}$ be the probabilities that an agent chooses alternative A and B in period t, respectively. Probability $P_{B,t}$ is specified as

$$P_{B,t} = \frac{\exp[\beta\pi_{B,t-1}]}{\exp[\alpha + \beta\pi_{A,t-1}] + \exp[\beta\pi_{B,t-1}]} = \frac{1}{1 + \exp[\alpha + \beta(\pi_{A,t-1} - \pi_{B,t-1})]}\,, \tag{2.3}$$

with $P_{A,t} = 1 - P_{B,t}$. Here $\beta \geq 0$ is the Intensity of Choice (IoC) parameter and α is a parameter that measures the so-called *predisposition effect*. The intuition behind the discrete choice model is that an increase in the performance of one alternative, relative to the other alternative, increases the probability that an individual agent chooses the former. The IoC parameter β measures how sensitive this probability is with respect to the performance difference. The predisposition effect measures a possible bias that agents have towards one of the alternatives: a positive value of α implies that alternative A is chosen with a higher probability than alternative B when their performance is the same. Similarly, a negative value of α implies that agents are biased towards alternative B.

Assuming that there are many agents, with all of them choosing between alternatives A and B according to (2.3), the Law of Large Numbers implies that the *fraction* of agents choosing alternative B is given by

$$n_{B,t} = \frac{1}{1 + \exp[\alpha + \beta(\pi_{A,t-1} - \pi_{B,t-1})]}\,, \tag{2.4}$$

with $n_{A,t} = 1 - n_{B,t}$.

The canonical discrete choice model, as used in [17, 18] and in the vast majority of subsequent contributions, is given by (2.4) with $\alpha = 0$. However, estimated discrete choice models on survey data [15] and experimental data [2] suggest that the predisposition effect plays an important role in explaining human behavior and therefore we add it here as well.[7]

We can now complete the model presented in Section 2.2.2 by assuming that (2.4) describes how the population of agents chooses between alternatives. Substituting (2.1) and (2.4) in (2.2) and using $\gamma_A + \gamma_B = 1$ and $C = W_B - W_A$, we find that the state variable x_t evolves according to

$$x_t = f\left(x_{t-1}\right) + \epsilon_t = \frac{\lambda x_{t-1}}{1 + \exp[\alpha + \beta(x_{t-1}^2 - C)]} + \epsilon_t\,. \tag{2.5}$$

Appendix A.1 analyzes the dynamics of this non-linear model and shows that the key properties of the heuristic switching framework outlined in [17, 18] are preserved. It turns out that the dynamics of the state variable, given by the first order difference equation (2.5), may give rise to complicated dynamics, even in absence of random shocks (i.e., setting $\epsilon_t = 0$). This depends in particular upon the structural parameters λ and C and the behavioral parameters α and β.

[7] There exist alternative formulations of the discrete choice model, for example where performance is measured by a weighted average of past profits (instead of only the profits from the previous period) or where updating is asynchronous (that is, agents do not have the opportunity to change between alternatives every period), see [4]. As an alternative to the discrete choice model, the evolution of the fraction choosing one particular alternative can be modelled by the (exponential) replicator dynamics, see, e.g., [38], [108] and [94]. Derived from a process of imitation, this evolutionary model leads to a more sluggish adaptation of fractions, which is similar to the effect of asynchronous updating.

Note that when the dynamics of the state variable are given by (2.5), the fraction of choices for alternative B evolves as

$$n_{B,t} = \frac{1}{1 + \exp[\alpha + \beta(x_{t-1}^2 - C)]}. \tag{2.6}$$

When the state variable equals $x^* = 0$, the fraction of B-choices is given by $n_B^* = [1 + \exp(\alpha - \beta C)]^{-1}$. The following finding (which focuses on the case $\lambda > 0$) follows from the more general result proven in Appendix A.1.

Proposition 2.2.1. *Consider the first order difference equation* (2.5) *with* $\epsilon_t = 0$. *For* $0 < \lambda < 1$ *there is a unique globally stable zero steady state,* $x^* = 0$. *For* $\lambda > 1$ *the zero steady state is unique and locally stable as long as* $\lambda n_B^* < 1$ *and unstable when* $\lambda n_B^* > 1$. *Moreover, when* $\lambda n_B^* > 1$, *two non-zero steady states* x^+ *and* x^- *exist, with* $x^+ = \sqrt{C + (\ln(\lambda - 1) - \alpha)/\beta}$ *and* $x^- = -x^+$. *The associated steady state fractions are* $n_B^+ = n_B^- = \lambda^{-1}$. *The non-zero steady states are locally stable if* $\beta\left(1 - n_B^+\right)\left(x^+\right)^2 < 1$.

Clearly, an increase in the fraction of agents using alternative B inhibits stability of the zero steady state (recall that alternative B corresponds to the cheap but destabilizing heuristic). Also note that we assume that $C = W_B - W_A$, the costs for using the more sophisticated and stabilizing rule A, is positive. Clearly, an increase in C increases n_B^* and destabilizes the zero steady state, because – at the steady state – more agents use the destabilizing rule. Likewise, an increase in the IoC parameter β will destabilize the dynamics, since this also increases the fraction of agents using alternative B (since they will do better at the steady state). An increase in α, on the other hand, will promote stability. If the zero steady state loses stability, for example because β increases, this occurs through a so-called *pitchfork bifurcation*. In this bifurcation two new, non-zero, steady

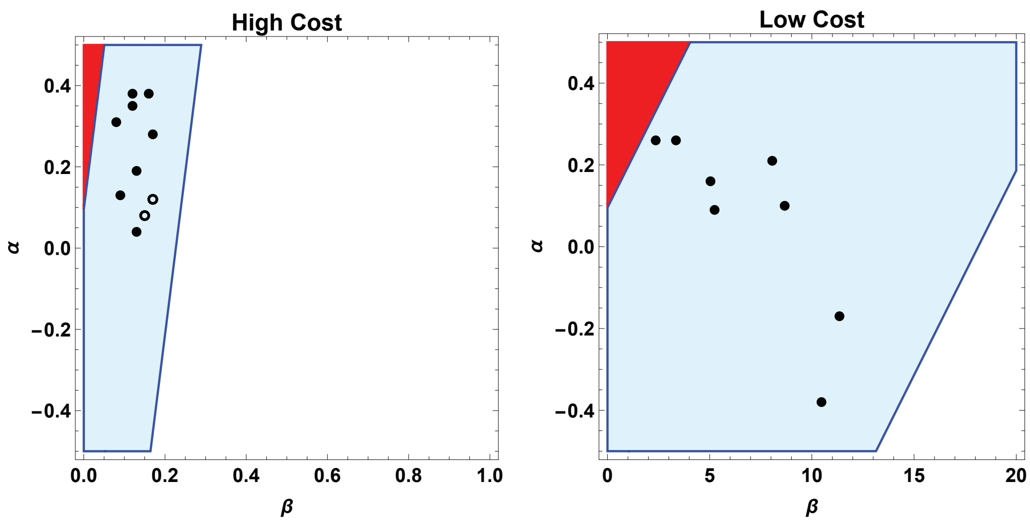

Figure 2.1: Bifurcation diagrams in (β, α)-coordinates for the stylized heuristic switching model, showing the regions of stability for the zero steady state (dark filled region), and non-zero steady states (light filled region). *Left:* Parameters corresponding to the High blocks in the experiment. *Right*: Parameters corresponding to the Low blocks in the experiment. Black points show the values of (β, α) estimated on subjects' choice data (see Section 2.5.2). The two black disks on the left panel show the estimated values from the *High Long* treatment (see Section 2.5.3).

states are created that are locally stable. The absolute deviation of these two steady states from zero is the same and depends upon the behavioral parameters of the model. With a further increase of β or C, these non-zero steady states become unstable as well.[8]

Fig. 2.1 shows the stability regions of the different steady states for the case of high costs, $C = 8$ (left panel) and low costs, $C = 0.1$ (right panel). The value of the parameter λ is 2.1, which is the value that will be used in the experiment. Note that for high costs the stability regions (both of the zero and of the non-zero steady states) are much smaller than for low costs (also note that the scale on the horizontal axis is very different).

[8]Note that for the case of $\lambda < 0$ we get a similar stability condition. However, in this case, instead of two non-zero steady states, a period two cycle $\{x^-, x^+\}$, with $x^- = f(x^+)$ and $x^+ = f(x^-)$, is created when the zero steady state x^* loses stability, see Appendix A.1.

The aim of the experiment in this chapter, outlined below, is to verify that the aggregate dynamics in our experiment will reproduce the patterns that are predicted by the heuristic switching model. In addition, we want to fit the discrete choice model on the aggregate choice data, and in particular estimate the relevant values of the intensity of choice parameter β and the predisposition parameter α.

2.3 Experimental Design

The experiment took place in June 2016 and October 2017 at the University of Technology Sydney Behavioral Laboratory. In total 80 subjects were recruited for four sessions, with 20 subjects participating in each session.[9] Most subjects are undergraduate students in economics and finance, and no subject participated in more than one session. Each session lasted approximately 90 minutes and subjects earned 25 Australian dollars (AUD) on average.

The task of the subjects in the experiment is to make a choice between two investment alternatives, labeled A and B. In the beginning of each session the 20 subjects that participate in the session are randomly divided in two groups of 10. Then, for 40 consecutive periods subjects make their binary choices, with their payoffs dependent on the distribution of choices of the 10 subjects in their group. After the first block of 40 decision periods is finished, subjects in the session are randomly rematched to form two new

[9] The Online Recruitment System for Economic Experiments (ORSEE), see [54], was used for recruiting participants. The experiment was conducted using the z-Tree software, see [46]. Two sessions of an additional treatment, with more periods and larger groups, took place in September 2017. We will discuss the results for this additional treatment in Section 2.5.3.

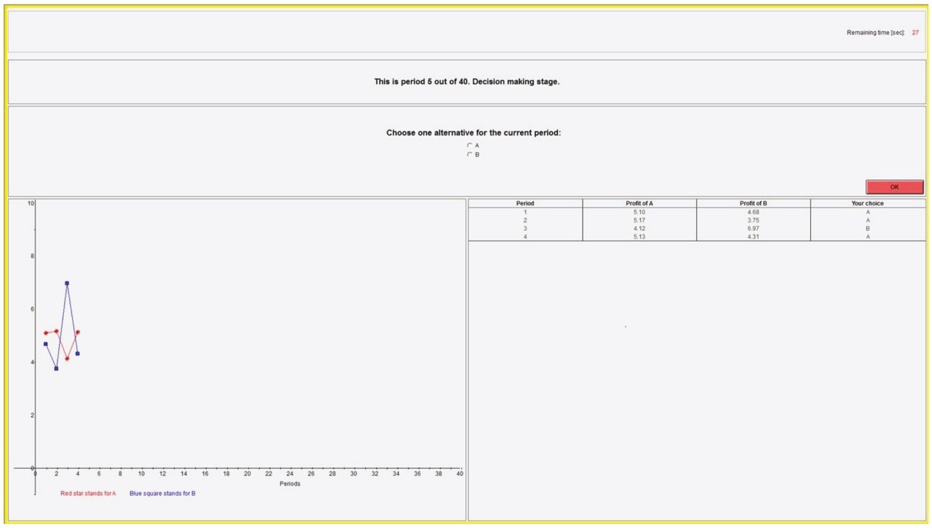

Figure 2.2: An example of the computer screen with interface elements. The upper part of the screen is used to submit a decision by clicking on one of two radio buttons, A or B. The lower part represents the history of profits for alternatives A and B in two formats: a graph on the left and a table on the right. The last column in the table shows the past choices of the subject.

fixed groups of 10 subjects. During the second block, the subjects have to make their binary decisions for another 40 periods, with their payofs depending on the choices of the subjects in their new group of 10 subjects.

At the beginning of every period subjects are provided with information about the past profits of the two alternatives, both in the form of a table and a graph, see Fig. 4.2. In addition they are informed about their own previous choices. As soon as all subjects in the group have made their decision for period t, the actual profits for the two alternatives are generated by (2.1), where the value x_t of the state variable[10] is determined by (2.2). The profits of both alternatives are then shown to the subjects together with their choice in period t. Subsequently they are asked to make their choice for period $t+1$.

[10] As explained in Section 2.2.2, this state variable can be thought of as the deviation of the price from the fundamental value. We do not ask subjects to predict the values of this variable (as they would do in the Learning to Forecast experiments). In fact, we do not even show the evolution of this variable to the subjects. This design has been chosen to focus exclusively on testing the discrete choice model, which assumes subjects only respond to the performance of the heuristics.

For both blocks in each session we choose $\lambda = 2.1$ (implying that the dynamics are unstable if, at equilibrium, at least half of the participants chooses alternative B, see Proposition 2.2.1), $\gamma_A = 0.6$ and $\gamma_B = 0.4$. The random shocks are IID distributed according to $\epsilon_t \sim N(0, 0.02)$. We use the same realization of random shocks for each block, and for each session. The only difference between blocks are the values of W_A and W_B, that are chosen in order to generate blocks with different *costs*, $C = W_B - W_A$, for the stabilizing heuristic. For the *High* blocks we choose $W_A = 1$ and $W_B = 9$ (i.e., high cost is given by $C = 8$) and for the *Low* blocks we choose $W_A = 4.95$ and $W_B = 5.05$ (i.e., low cost is given by $C = 0.1$).

The two treatments in our experiment, with two sessions each, only differ in the order of the blocks: treatment *High-Low* starts with a block with $C = 8$, followed by a block with $C = 0.1$, whereas the order is reversed for treatment *Low-High*.[11] We therefore have eight groups that make decisions in a High block and eight groups that make decisions in a Low block. Each of these sixteen groups consists of 10 subjects, with each subject participating in exactly one High and one Low block. In the remainder we will identify groups by block, session and group number (e.g., "High: session 1, group 2" refers to the second group in the first block of the first session).

By the end of the experiment, subjects are paid for their decisions according to the following procedure. For every subject we randomly choose two periods from the first block and two periods from the second block. The sum of the payoffs corresponding to the decision of the subject in these four periods is divided by two. This constitutes the subject's earnings in Australian dollars. In addition, each subject receives a 10 AUD participation

[11] Treatment *High-Low* was implemented in sessions 1 and 3, treatment *Low-High* was implemented in sessions 2 and 4.

fee. The procedure ensures that subjects have the financial incentive to make the best possible decision in every period of the experiment. Experimental instructions are provided in Appendix A.2.[12]

2.4 Hypotheses

If the discrete choice model (2.4) gives a good description of the aggregate behavior of subjects, then (2.5) and (2.6) should provide a good approximation[13] of the dynamics of the fraction $n_{B,t}$ and the state variable x_t. The dynamics will depend upon the specific characteristics of the economic decision problem (λ, γ_A, γ_B, W_A and W_B), which are chosen by the experimenters, and upon the behavioral parameters α and β, which describe the subjects' decision making. Since $C = W_B - W_A > 0$ in both blocks, in equilibrium we will have $n_B^* > 1/2$ (assuming α is not very high and $\beta > 0$). Given that $\lambda = 2.1$, we therefore expect the zero steady state, $(x^*, n_B^*) = (0, [1 + \exp(\alpha - \beta C)]^{-1})$, to be locally unstable in both the High and the Low blocks, see Proposition 2.2.1. Moreover, the values of W_A and W_B in the two blocks are such that for a large range of values of α and β the steady states (x^+, n^+) and (x^-, n^-) will be stable for the Low blocks, but unstable for the High blocks, see Fig. 2.1.

[12]Before the experiment starts subjects are required to solve a short quiz which is designed to check their understanding of the feedback effects and remuneration procedure. After the experiment, subjects are given a questionnaire that focuses on background information, such as demographic characteristics and field of study. The quiz and the questionnaire are available from the authors upon request.

[13]In the model fraction $n_{B,t}$ can take on any real value in the interval $[0, 1]$. In the experiment, however, there are 10 subjects in a group, and, therefore, this fraction can only take on 11 values $(0, 0.1, 0.2, \ldots, 1)$. Hence, (2.5) and (2.6) can only provide an approximation to the model dynamics.

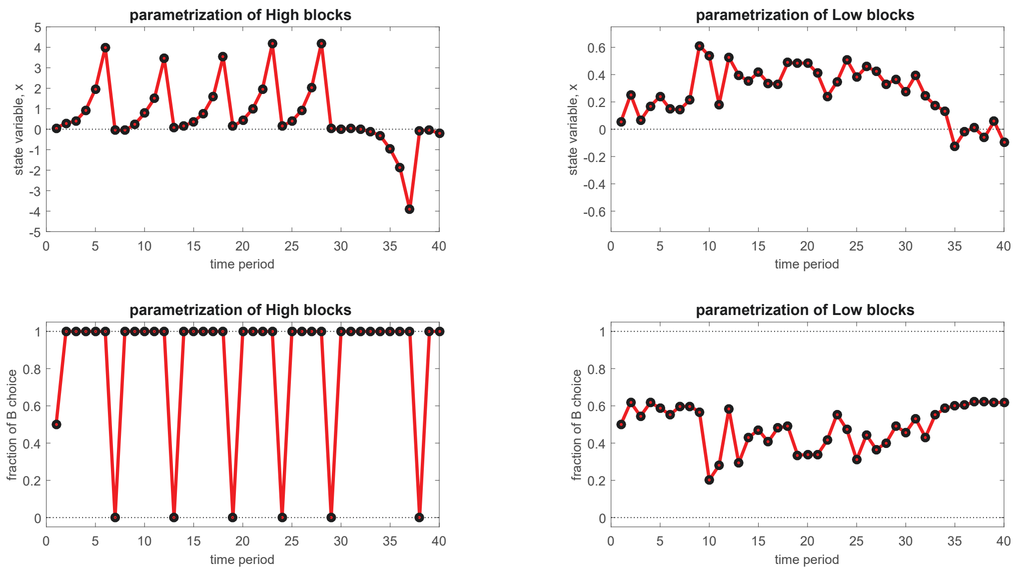

Figure 2.3: Simulations of the stylized heuristic switching model with $\lambda = 2.1$. *Left:* High blocks parametrization, $C = 8$. *Right:* Low blocks parametrization, $C = 0.1$. *Upper panels*: Dynamics of state variable, x. *Lower panels:* Dynamics of fractions of B-choices, n_B.

Fig. 2.3 illustrates the simulated dynamics of the stylized heuristic switching model (2.5)–(2.6) for the values of the structural parameters that we use in the experiment, and with the same realization of the random shocks ϵ_t that is used in the experiment. For these simulations we set the behavioral parameters to $\alpha = 0$ and $\beta = 5$. The left panels show the dynamics of x and n_B in the High blocks. We observe a pattern familiar from the original model of [18] with endogenous bubbles and crashes of the state variable (corresponding to the deviation of the price from the fundamental asset value in their framework). Note that due to the small random shocks, the realized state variable may become negative when it approaches zero. The right panels show the same dynamics for the Low blocks. Due to positive values of ϵ in the first couple of periods, the state variable becomes positive and stays close to the steady state value $x^+ = \sqrt{\ln(1.1)/5 + 0.1} \approx 0.345$ thereafter, although it may occasionally 'escape' from that steady state when a sufficiently low value of the noise term ϵ is realized (which happens in period $t = 35$).

The specific time series shown in Fig. 2.3 depend on the specific values of α and β, but the main properties of the model, i.e., instability of all steady states with endogenous bubbles and crashes for the High blocks and stability of the non-zero steady state for the Low blocks, hold for a large range of values of these behavioral parameters. Thus, if choice behavior is governed by the same heuristic switching model in both environments we expect less stability and more volatile dynamics in the High blocks than in the Low blocks, both in terms of the fraction of subjects choosing B, and in the state variable x_t. This leads to our first hypothesis on the dynamics of $n_{B,t}$ and x_t.

Hypothesis 1. *There is a substantial difference in the volatility of both $n_{B,t}$ and x_t between the High blocks and the Low blocks.*

If Hypothesis 1 is not rejected, it implies that the qualitative predictions of the heuristic switching model are confirmed. The next step is to investigate whether the discrete choice model (2.4) also gives a good quantitative description of the data. To that end, we fit the discrete choice model on the aggregate data. This gives our next hypothesis.

Hypothesis 2. *The endogenous variable $n_{B,t}$ can be described by a discrete choice model with one lag and a predisposition effect.*

If this hypothesis is confirmed, it provides experimental evidence for the relevance of the discrete choice model (2.4) as a description of aggregate decision making, and would thereby lend support to the use of heuristic switching models. Note that in [2] it was established that a discrete choice model with one lag and a predisposition effect is relatively successful in describing experimental data in a setting where payoffs for the different alternatives are exogenously generated. Hypothesis 2 checks whether this is also the case when there is feedback from subjects' decisions to the payoffs.

Ideally, the estimated discrete choice model – if it provides a good description of the data – is similar for the High and the Low blocks. Comparable values of the model's behavioral parameters would suggest that human decision making is independent of the specifics of the economic environment and governed by the same underlying laws. Indeed, Hypothesis 1 implicitly assumes that the decision process is the same for the High and Low blocks and so the variation in the dynamics occurs because of the difference in the other characteristics of the High and Low blocks (i.e., the difference in costs $C = W_B - W_A$). Our third and final hypothesis deals with investigating this issue.

Hypothesis 3. *There is no significant difference between the discrete choice models estimated on data from the High blocks, and the discrete choice models estimated on data from the Low blocks.*

In the next section we present the experimental data and test these three hypotheses.

2.5 Experimental results on switching

In this section we will discuss the experimental results. We start out with presenting the experimental data in Section 2.5.1. We will also provide some descriptive statistics and test Hypothesis 1 in that section. In Section 2.5.1

we estimate the discrete choice model on the experimental data and discuss Hypotheses 2 and 3. Finally, Section 2.5.3 is devoted to the analysis of an additional treatment with high costs, which features larger groups and more periods.

2.5.1 The experimental data

First, we depict the evolution over time of the fraction $n_{B,t}$ of subjects choosing alternative B. Figs. 2.4 and 2.5 show this fraction in the High and Low blocks, respectively. A quick visual inspection of these figures suggests that the time series of $n_{B,t}$ may indeed be slightly more volatile in the High blocks. For example, out of 8×40 observations the boundary values of 0 and 1 are reached 23 times (7.2%) in the High blocks, whereas they are only reached 10 times (3.1%) in the Low blocks. However, the standard deviations of $n_{B,t}$ in the different groups, reported in Table 2.1, appear to be roughly similar for the High and the Low blocks. Indeed, the difference in standard deviations between High and Low blocks is not statistically significant at the 5% level.[14]

To further investigate possible differences in the fractions, the top panels of Fig. 2.6 show the histograms of the fraction $n_{B,t}$ of subjects choosing alternative B in the High and Low blocks, pooled over all eight groups. These histograms also show a small, but distinct, difference between the two types of blocks. The distribution of choices in the Low blocks has a

[14]The Ansari-Bradley test, which is suitable for distributions with similar means and shapes, gives a p-value of 0.143, implying that we cannot reject the hypothesis of equal standard deviations of the two distributions. Note that, in order to preserve independence of the observations, we only used the first block in each session for this test (i.e., the four High blocks from Sessions 1 and 3 and the four Low blocks from Sessions 2 and 4). For all other test statistics (unless noted otherwise) we pool the data from the first and second High (respectively Low) blocks.

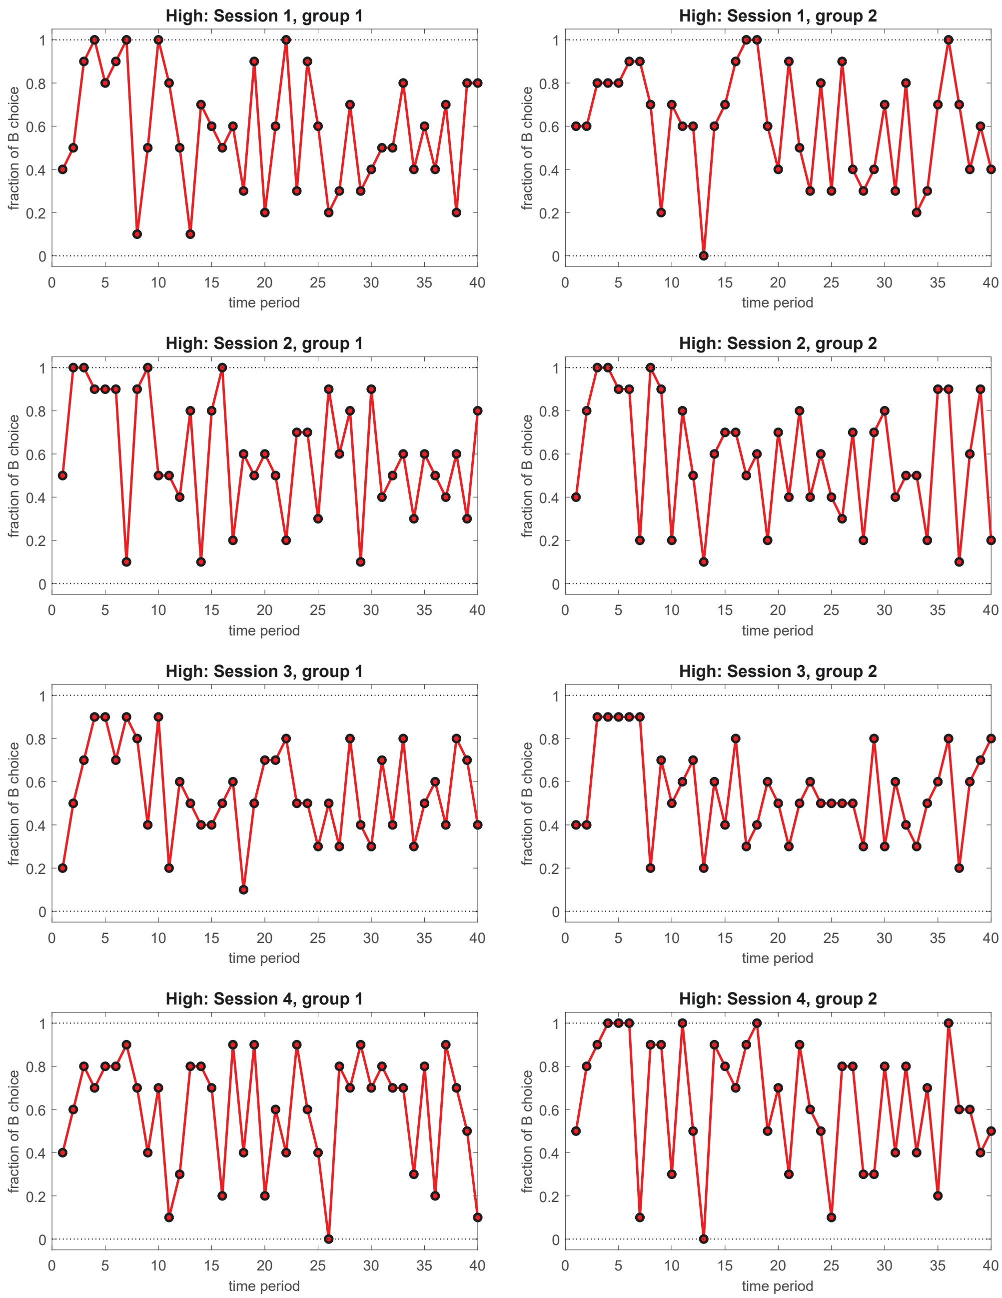

Figure 2.4: Fraction of choices of alternative B in High blocks.

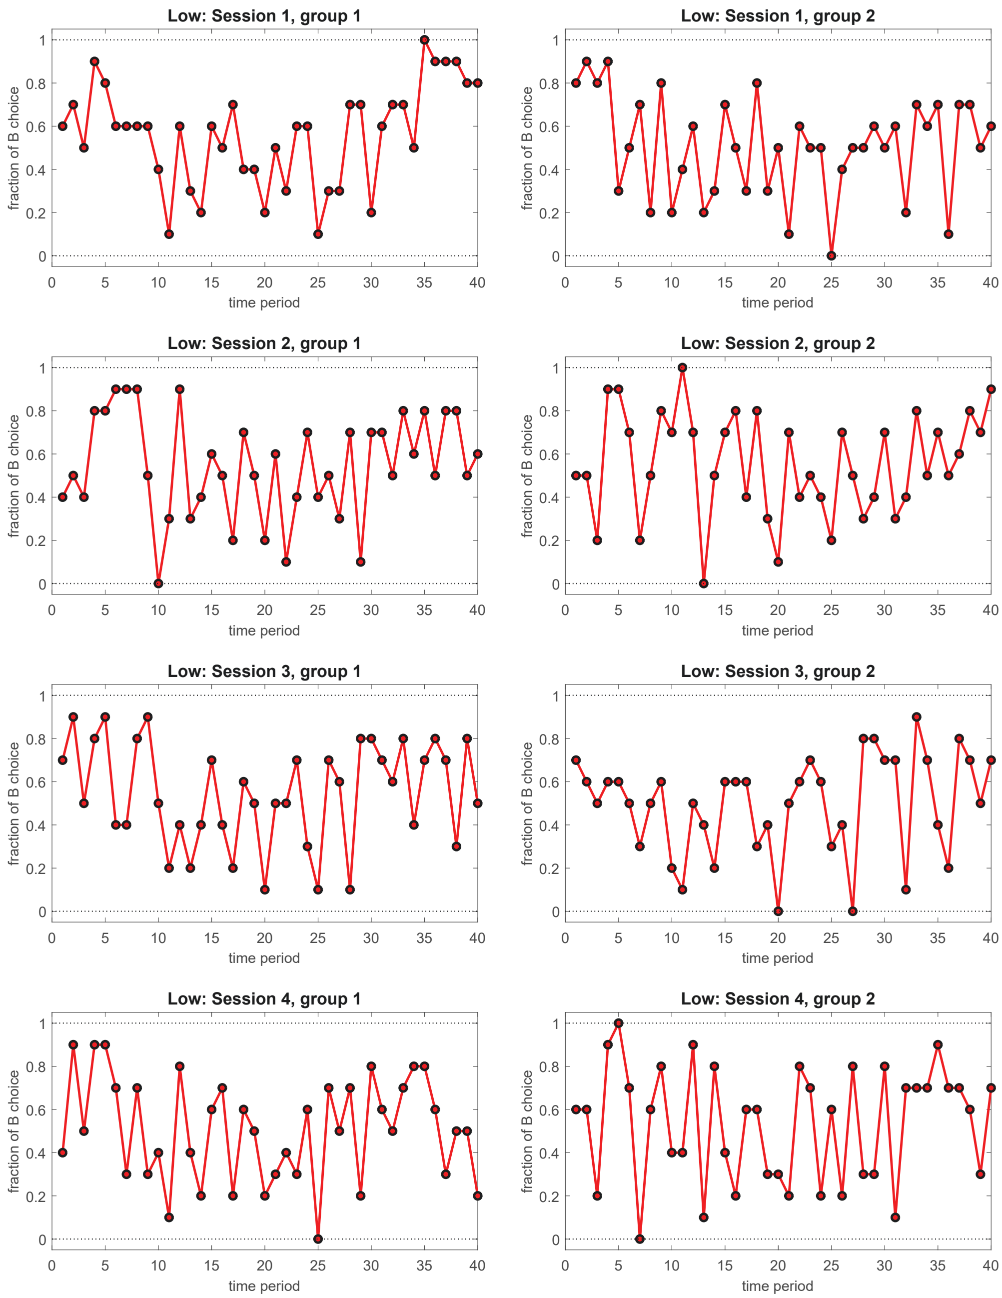

Figure 2.5: Fraction of choices of alternative B in Low blocks.

	Data	Fraction of B-choices, n_B Mean	Std. Dev.	State variable, x Mean	Std. Dev.
High	Session 1. Group 1	0.58	0.27	2.43	1.52
	Session 1. Group 2	0.61	0.25	2.26	1.52
	Session 2. Group 1	0.60	0.27	2.51	1.47
	Session 2. Group 2	0.58	0.28	2.39	1.27
	Session 3. Group 1	0.55	0.22	2.33	1.27
	Session 3. Group 2	0.55	0.21	2.35	1.11
	Session 4. Group 1	0.60	0.26	2.02	1.45
	Session 4. Group 2	0.64	0.29	2.15	1.42
	All High groups	0.59	0.26	2.43	1.52
High Long	Session 5. Group 1	0.53	0.18	2.62	0.92
	Session 6. Group 1	0.55	0.21	2.48	1.06
Low	Session 1. Group 1	0.56	0.23	0.26	0.22
	Session 1. Group 2	0.51	0.23	0.26	0.24
	Session 2. Group 1	0.55	0.24	0.17	0.21
	Session 2. Group 2	0.56	0.24	0.13	0.28
	Session 3. Group 1	0.55	0.24	0.32	0.22
	Session 3. Group 2	0.50	0.23	0.12	0.25
	Session 4. Group 1	0.51	0.24	0.31	0.18
	Session 4. Group 2	0.54	0.27	0.16	0.32
	All Low groups	0.53	0.24	0.26	0.22

Table 2.1: Descriptive statistics of the fraction n_B and the state variable x.

clear peak around $0.5 - 0.7$ (containing around half of the observations), whereas the distribution of fractions in the High blocks is much more evenly spread with substantially more observations of fractions $n_{B,t}$ close to 1. In fact, the hypothesis of equality of the two distributions is rejected at the 5% level.[15]

It is clear from the histograms of the fraction of B-choices that, although there is a difference between the High and Low blocks in our experiment, this difference is smaller than we would expect in case of stable and unstable dynamics in the heuristic switching model. Consider, for example, the time series of fractions given in the lower panels of Fig. 2.3 that are generated with $\alpha = 0$ and $\beta = 5$. The histogram of fractions corresponding to the lower left panel of Fig. 2.3 (High cost parametrization) will be bimodal, with many observations close to 0 or close to 1, whereas the histogram of fractions

[15]The Kolmogorov-Smirnov test for equality of distributions gives a p-value of 0.012.

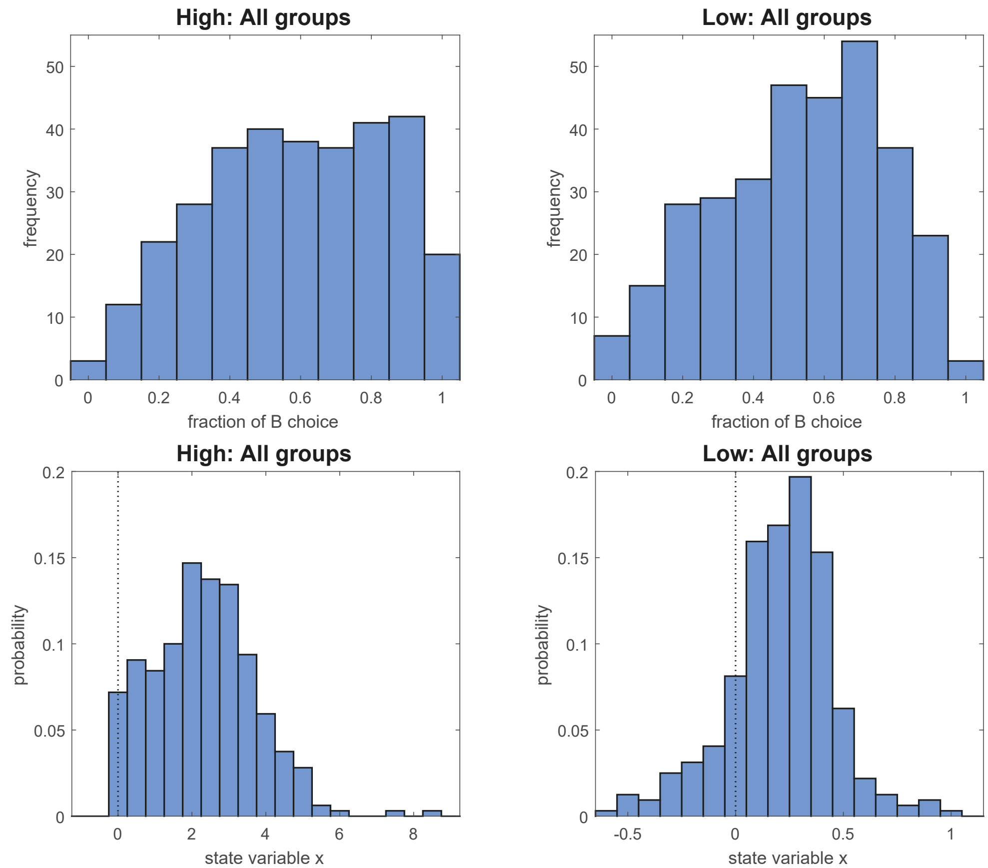

Figure 2.6: Histogram of fraction of B-choices (*Upper panels*). Density histogram of state variable, x (*Lower panels*). *Left:* High blocks. *Right:* Low blocks. The vertical dashed lines in the lower panels indicate the zero steady state of x.

corresponding to the lower right panel (Low cost parametrization) will be single-peaked. Whereas the latter is indeed consistent with the histogram for our Low block groups in Fig. 2.6, the histogram for the High block groups can hardly be described as bimodal.

Because the fraction of subjects choosing alternative B is high in many periods, deviations of x_t from zero should be quite persistent, at least in those periods, see Eq. (2.2). This is confirmed by inspection of Fig. 2.7 which shows the time series of x_t in all blocks of each session. Clearly, the mean and variance of x_t are much higher for the High blocks than for the Low

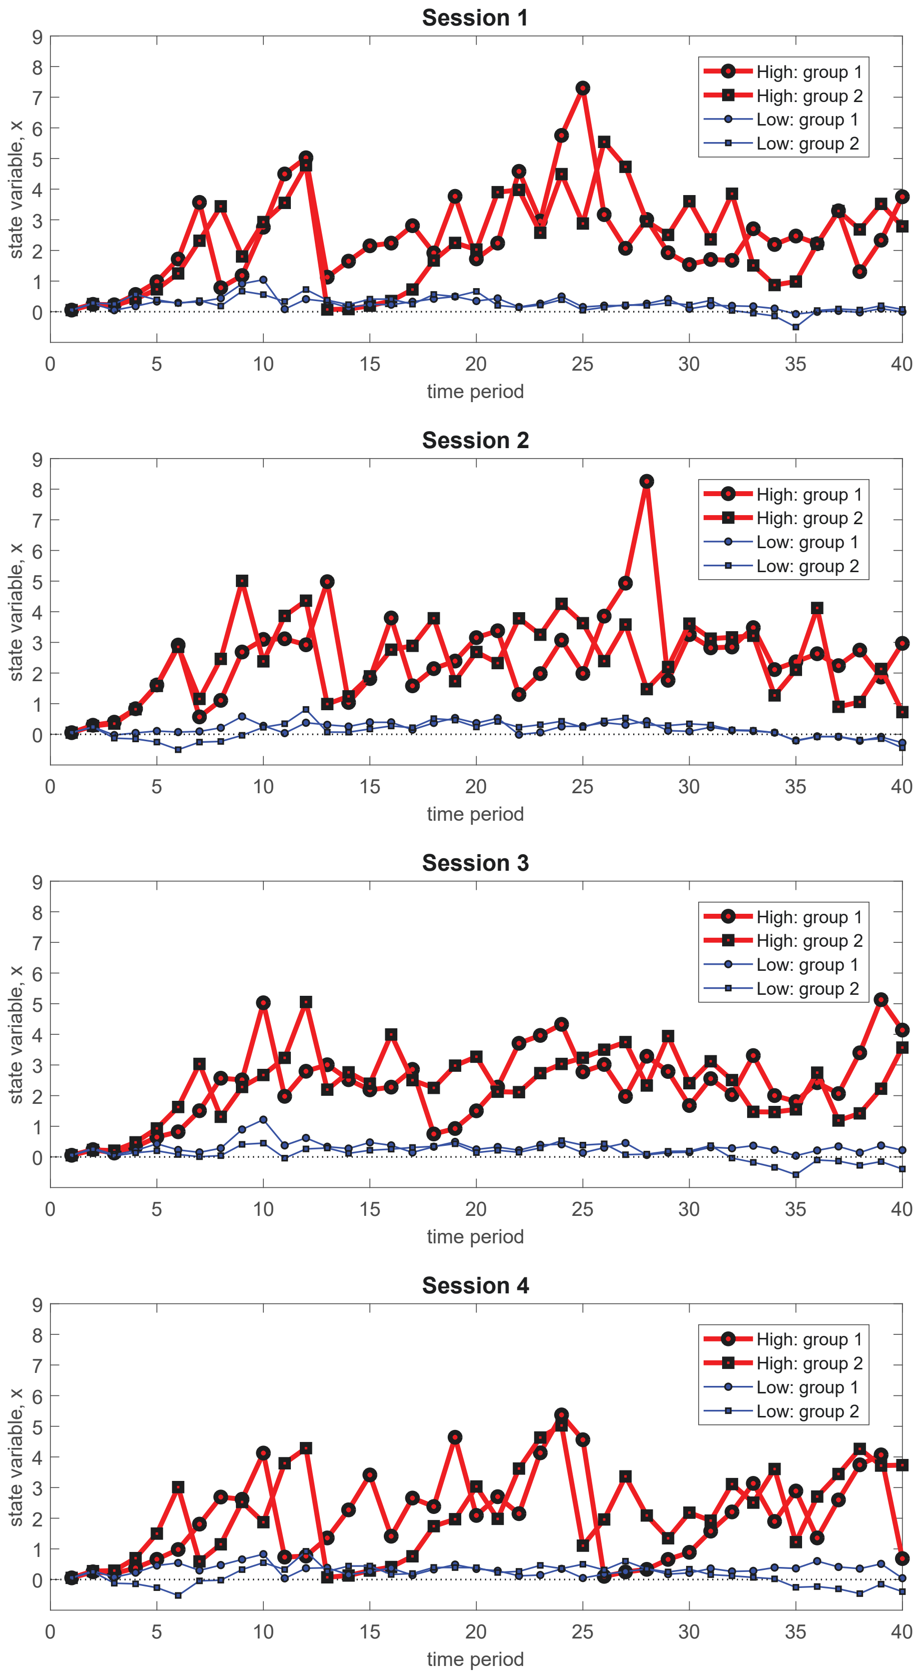

Figure 2.7: Time series of the state variable x_t in all 16 groups.

blocks, as can also be seen from the last two columns of Table 2.1. Both the difference in means and the difference in standard deviations is statistically significant at the 5% level.[16] Moreover, the mean values of x_t, both in the High and in the Low blocks, are significantly different from zero as well, and the standard deviation of these time series is significantly higher than they would be in a steady state of the model (where the standard deviation of x_t is equal to that of ϵ_t which, by construction, is approximately 0.14).[17]

Summarizing our results thus far, we conclude that, when looking at the dynamics of both $n_{B,t}$ and x_t, Hypothesis 1 cannot be rejected, although, particularly for $n_{B,t}$, the difference between blocks is smaller than expected.

A more detailed look at the time series of x_t leads to some interesting observations. First, the dynamics in the first half of the block are qualitatively similar for each of the eight High block groups. In each group the variable x_t increases in several consecutive periods after which it 'crashes' in one period, sometimes by a considerable extent.[18] After this crash x_t starts to increase again. This cycle is repeated two or three times in each of the eight groups. Interestingly, this type of dynamics is characteristic for the type of heuristic switching model studied in the literature on heterogeneous

[16]The Ansari-Bradley test of equal standard deviations (applied to the first blocks of each session again) gives a p-value of 0.003. Since the data looks closer to a normal distribution now, see the lower panels of Fig. 2.6, we also use an F-test for equal variances, which returns a p-value of 0.000. Similarly, the p-value of the t-test of equal means is 0.000. Finally, the Kolmogorov-Smirnov test for equality of distributions gives a p-value of 0.000.

[17]We can reject the hypothesis of zero mean for x_t, both for the High and for the Low blocks at the 1% significance level, using the t-test (p-values of 0.000 for both High and Low blocks). We can also reject the hypothesis that the variance of x_t is equal to the variance of ϵ_t, i.e., 0.02, in favor of the alternative hypothesis that the variance is higher, both for the High and for the Low blocks at a 1% significance level, using the Chi-Square test (p-values of 0.000 for both the High and Low blocks).

[18]Note that the initial value of x was chosen to be $x_0 = 0$. It follows that in the first period of the experiment $x_1 = \epsilon_1 = 0.0538$ (independent of the subjects' choices). If the first realization of the random variable ϵ_t would have been negative, it is likely that $x_t < 0$ for all t.

agents, see the upper left panel of Fig. 2.3 (also see Fig. 2 in [18]). However, in the second halves of the High blocks this structure in the dynamics disappears in each of the eight groups and the behavior of x_t becomes more irregular with no apparent structure. Second, the dynamics of x_t in the Low blocks are quite different and seem to be consistent with small irregular fluctuations around a fixed positive value, with a decrease in x_t in the last couple of periods.[19] Note that, since the non-zero steady state value x^+ from Proposition 2.2.1 depends upon the behavioral parameters α and β, it is difficult to test directly whether the mean of x_t equals x^+.[20]

2.5.2 Estimated discrete choice models

Our next step is to fit the discrete choice model (2.4) to the experimental data. That is, we estimate the discrete choice model separately for the aggregate choices in each group. Table 2.2 shows the results. Columns 3 and 4 give the estimated values and the standard deviations for the intensity of choice parameter β, and columns 5 and 6 give the estimated values and the standard deviations for the predisposition parameter α. We also estimated the model on the pooled data for all High block groups and all Low block groups.

The data for each of the sixteen groups can be described quite well by the discrete choice model, which confirms Hypothesis 2. In particular, the

[19] This decrease in the last couple of periods seems to be due, at least partially, to a large negative shock in period $t = 35$ of $\epsilon_{35} = -0.2944$.

[20] Since subjects do not observe the realized values of the state variable x_t, and because x_t only enters the profit functions (2.1) quadratically, its sign does not matter for the dynamics of fractions. Choosing $\lambda = -2.1$, instead of $\lambda = 2.1$ would therefore lead to the same experimental results (abstracting from the effect of random shocks), but with the state variable alternating between positive and negative values. For the Low blocks the dynamics of the state variable will then resemble a noisy period-two cycle, instead of a noisy non-zero steady state.

	Data	IoC		Predisposition		Zero SS		Non-Zero SS	
		Beta	**S.E.**	**Alpha**	**S.E.**	(x^{*}, n^{*})	$f'(x*)$	(x^{+}, n^{+})	$f'(x^{+})$
High	Session 1. Group 1	0.08	0.01	0.31	0.11	(0,0.58)	1.22	(2.30,0.48)	0.57
	Session 1. Group 2	0.12	0.02	0.38	0.11	(0,0.64)	1.34	(2.36,0.48)	0.31
	Session 2. Group 1	0.12	0.02	0.35	0.11	(0,0.64)	1.35	(2.42,0.48)	0.28
	Session 2. Group 2	0.17	0.02	0.28	0.11	(0,0.74)	1.56	(2.63,0.48)	-0.22
	Session 3. Group 1	0.09	0.02	0.13	0.11	(0,0.64)	1.35	(2.76,0.48)	0.29
	Session 3. Group 2	0.13	0.02	0.04	0.11	(0,0.72)	1.52	(2.90,0.48)	-0.11
	Session 4. Group 1	0.13	0.02	0.19	0.12	(0,0.69)	1.46	(2.69,0.48)	0.05
	Session 4. Group 2	0.16	0.02	0.38	0.12	(0,0.71)	1.49	(2.49,0.48)	-0.03
	All High groups	0.12	0.01	0.25	0.04				
High Long	Session 5. Group 1	0.15	0.02	0.08	0.09	(0,0.75)	1.58	(2.85,0.48)	-0.27
	Session 6. Group 2	0.16	0.02	0.10	0.09	(0,0.76)	1.60	(2.82,0.48)	-0.32
Low	Session 1. Group 1	3.35	0.84	0.26	0.11	(0,0.52)	1.09	(2.23,0.48)	0.82
	Session 1. Group 2	5.24	0.93	0.09	0.11	(0,0.61)	1.27	(2.32,0.48)	0.45
	Session 2. Group 1	11.35	1.66	-0.17	0.13	(0,0.79)	1.65	(2.35,0.48)	-0.46
	Session 2. Group 2	8.67	1.47	0.10	0.11	(0,0.68)	1.43	(2.32,0.48)	0.10
	Session 3. Group 1	2.36	0.67	0.26	0.10	(0,0.49)	1.04	(2.18,0.48)	0.92
	Session 3. Group 2	10.47	1.72	-0.38	0.13	(0,0.81)	1.69	(2.38,0.48)	-0.60
	Session 4. Group 1	5.04	1.04	0.16	0.11	(0,0.58)	1.23	(2.29,0.48)	0.54
	Session 4. Group 2	8.06	1.31	0.21	0.11	(0,0.65)	1.36	(2.29,0.48)	0.27
	All Low groups	5.71	0.42	0.11	0.04				

Table 2.2: Estimation of discrete choice model with two parameters, the intensity of choice parameter, β, and the predisposition parameter, α.

intensity of choice parameter β is positive and significantly different from 0 for all groups. The predisposition parameter α is positive (implying a predisposition towards alternative A) and significantly different from 0 in five of the eight High block groups, but only significantly different from 0 in three of the eight Low blocks (and positive for only two of those).[21]

One remarkable feature of the parameter estimates stands out from Table 2.2. Although the estimated intensity of choice coefficients for the same type of blocks have roughly the same order of magnitude, these coefficients are radically different between different types of blocks, with the estimates for the Low block groups about 50 times as high as those for the High block groups. The subjects therefore seem to be much less sensitive to profit differences in the 'unstable' High blocks than in the 'stable' Low blocks. Clearly, this means that we have to reject Hypothesis 3. This is broadly consistent

[21] Estimating the discrete choice model with the restriction $\alpha = 0$ leads to new estimates for β, which are quite close to the estimates in Table 2.2.

with the results of [2] who show that the estimated values of the intensity of choice increase when there is more structure in the time series of payoffs, and indeed the time series of payoffs in the Low blocks are less volatile and more predictable than those in the High blocks.[22]

For each group we superimposed the estimated values of α and β in the stability graphs of Fig. 2.1. First consider the eight data points corresponding to the Low blocks (the right panel). As we expected, all of them are in the region of the parameter space where the zero steady state (x^*, n^*) is unstable, but where the non-zero steady state (x^+, n^+) is locally stable. This is confirmed by the last four columns in Table 2.2, which show the zero and non-zero steady state and the slope of the dynamical system, at those steady states, respectively, that are implied by the estimated values of α and β. Comparing x^+ with the mean value in the experiment, given in column 4 of Table 2.1, suggests that the dynamics indeed converges to the non-zero (positive) steady state in each of those groups, although there is some excess volatility around that steady state.

Now consider the estimated discrete choice models for the High block groups, which are depicted in the left panel of Fig. 2.1. When designing the experiment we expected the dynamics in these High blocks to be unstable, and indeed the variations in x_t and, to a weaker extent, $n_{B,t}$, are higher in those blocks – see the discussion in Section 2.5.1. However, it turns out that the estimated parameters for these blocks are located in the region of local stability of the non-zero steady state, even if it requires very low levels of

[22] Another reason may be that the discrete choice model is not invariant with respect to changes in profit levels. That is, an increase in β is equivalent with exactly the same increase in profits, see Eq. (2.4). However, in our experiment average profits, as well as average profit differences, are roughly the same for both types of blocks (although the standard deviation of profits is higher in the High blocks).

the intensity of choice parameter β. Moreover, the non-zero steady state values of x^+ shown in Table 2.2, implied by the estimated values of α and β, are very close to the mean value of x_t in each of these eight High block groups. Although there is substantial volatility around the steady state, it seems that also in this case there is convergence to the non-zero steady state. Apparently, after experiencing the dynamics that is typical for the heuristic switching model in the first half of the block, subjects become more cautious and are able – at least to a certain extent – to stabilize the dynamics in this highly volatile and unstable environment, by behaving according to a discrete choice model with a low value of the intensity of choice parameter β.

2.5.3 Further evidence on the endogeneity of the IoC parameter

The analysis in Sections 2.5.1 and 2.5.2 suggests that subjects in our experiment on the heuristic switching model have a tendency to adapt their choice behavior to their environment. Partly due to the time and space constraints imposed by using laboratory experiments, the number of subjects and periods in our experiment are limited to 10 subjects per group, and 40 periods per block, respectively. A relevant question is whether our results are robust to increases in the group size and time horizon. To investigate this, we run an additional treatment that differs from the two treatments described and analyzed above in three ways: (i) the parametrization of the underlying model is fixed over the course of the treatment, and equal to that of the High blocks in the other two treatments; (ii) the experiment runs for 60 periods; and (iii) we increase the group size to 33 participants. Note

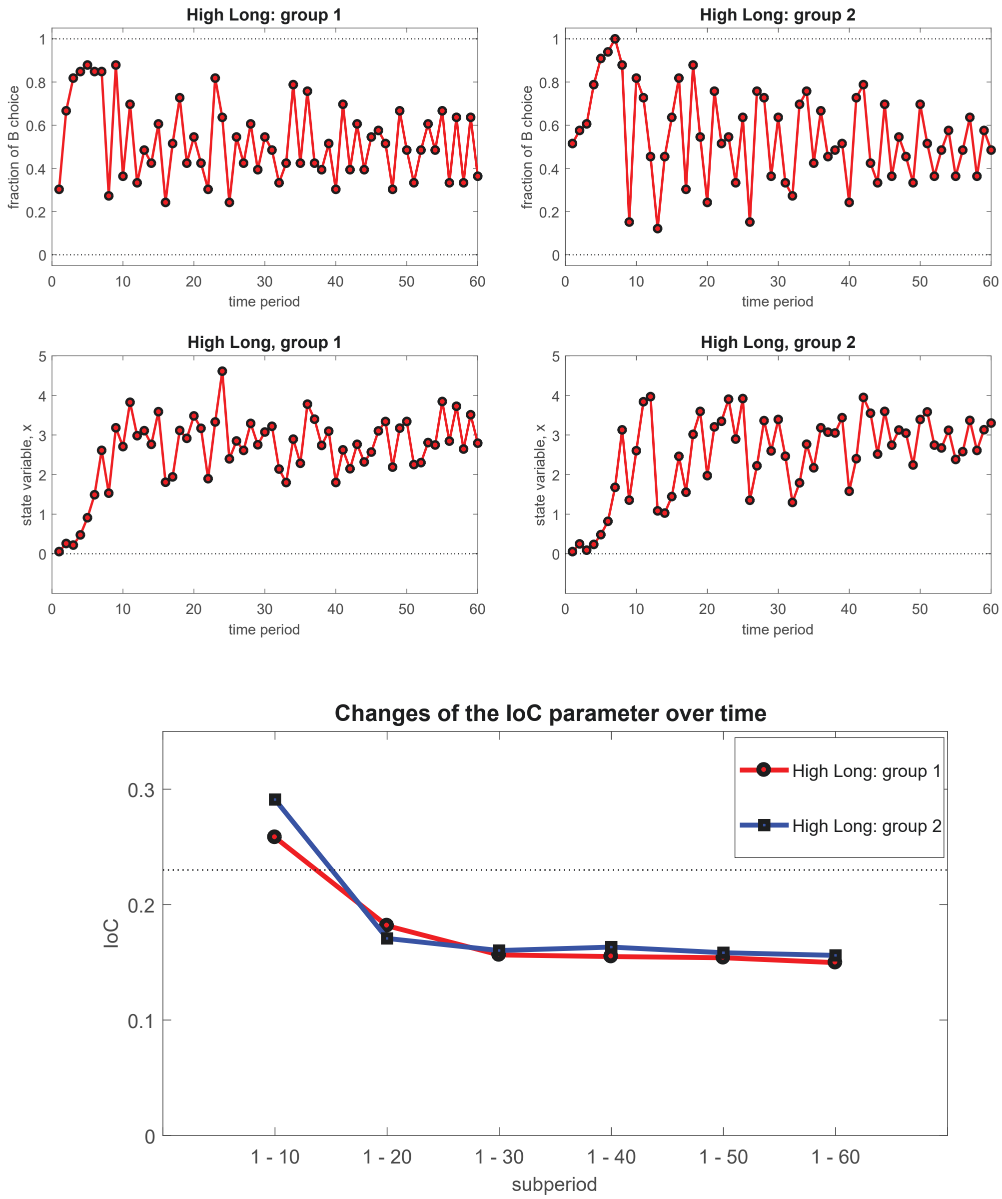

Figure 2.8: Experimental results of the two groups in the additional *High Long* treatment. *Upper panels*: Dynamics of fractions of B -choices, n_B. *Left lower panel:* Dynamics of the state variable, x. *Right lower panel:* Changes in the intensity of choice parameter, estimated on the 6 extending subsamples of the data from the *High Long* treatment. The dotted line represents the boundary above which the non-zero steady state is unstable.

that the increase in the number of periods allows us to better study the adaptation of the intensity of choice parameter to the decision environment over time, and the increase in the number of participants gives us a less coarse approximation of the fraction $n_{B,t}$ of the original model. We ran two sessions of this treatment, with one group in each session.

The upper panels of Fig. 2.8 show the evolution of the fraction $n_{B,t}$ in the two groups. The larger group size hardly smooths out the dynamics: The fractions are still quite volatile over the full 60 periods of the experiment. In fact, in quite some periods at least around half of the subjects switch to another heuristic. However, the larger group size contributes to the fact that the fraction $n_{B,t}$ rarely approaches its boundary values of 0 and 1 closely – only in one period in session 2 all 33 participants choose the same heuristic. The middle panels of Fig. 2.8 show the dynamics of the state variable x_t for the two groups. Again, both groups start out with the familiar pattern of a slowly increasing bubble which crashes after a couple of periods, after which the state variable increases again. Eventually – as with the shorter High blocks studied above – the dynamics fluctuate in an erratic manner around some positive fixed value of the state variable (note that the amplitude of the fluctuations in the state variable is less than that in the High blocks, again due to the increased group size). These results suggest that our earlier findings for the High cost blocks are robust to increasing the number of periods and the group size. This is confirmed by the descriptive statistics for this *High Long* treatment in Table 2.1 and the estimation of the discrete choice model for this treatment presented in Table 2.2 and in Fig. 2.1, which are consistent with those for the High blocks.

The larger number of periods in the *High Long* treatment allows us to investigate in more depth how subjects adapt their behavior to the dynamics of payoffs. To that end, we split our sample into six subsamples of increasing length, with the first subsample consisting of the first 10 time periods, the second subsample consisting of the first 20 time periods, and the last 'subsample' corresponding to the full time series of 60 periods. The discrete choice model is estimated on each of these six subsamples, and the estimated value of the intensity of choice parameter β is presented in the bottom panel of Fig. 2.8. We see that the estimated intensity of choice parameter indeed decreases over time. This parameter lies above the stability threshold (assuming $\alpha = 0$) of the non-zero steady state for the first subsample, but below it from the second subsample onwards. Moreover, the estimated values show very similar patterns for the two different groups.

Why do subjects become less sensitive to past performance over time in the High cost environment? This may be because in that environment payoffs are very erratic and volatile, and may not predict future payoffs very well. Subjects may realize this after the first ten to fifteen periods, and then start to rely to a lesser extent upon past performance when choosing their heuristic. This in turn brings down volatility. Indeed, the standard deviation of profit differences is equal to 4.76 and 5.21 in the first 20 periods in the two *High Long* groups, but for periods 21–40 the standard deviation is down to 4.25 and 4.13, respectively, and it decreases to 3.01 and 2.95 for the last 20 periods.

2.6 Conclusion

After its introduction in [17, 18], the heuristic switching model has become a workhorse model in the field of heterogeneous agents and agent-based modeling. The standard discrete choice framework, according to which agents choose between heuristics – typically corresponding to different forecasting rules – on the basis of their past performance, is a central element of this model. The resulting changes in the distribution of agents over heuristics influence the evolution of the state variable (typically, the asset price), which feeds back into the performance of the heuristics. This interaction between the dynamics of the state variable, and the dynamics of the distribution of heuristics, is capable of generating endogenous bubbles and crashes, excess volatility, and other stylized facts of financial markets. It therefore presents a natural extension to the work on the dynamics arising from the coexistence of fundamental and trend-following rules that started with [35] and [25].

In this chapter we present a laboratory experiment to test the heuristic switching model. Similar to the experiment described in [2], the only task of the subjects is to choose one of two heuristics, and subjects are paid according to the performance of the heuristic they choose. Contrary to the previous experiment, however, in the experiment presented here the subjects' aggregate choices determine, through a hidden state variable, the payoffs generated by the heuristics – an important feature of the standard heuristic switching model as well. In particular, our experiment reproduces, in a stylized form, the interaction between a costly sophisticated and a simple cheap heuristic, where the latter is destabilizing when used by many subjects, and the former is stabilizing.

We vary the (implicit) cost of using the stabilizing heuristic between different blocks in the experiment. Theoretically, the heuristic switching model will generate dynamics that are relatively stable when costs are low (in the Low blocks) and give rise to endogenous bubbles and crashes when costs are high (in the High blocks). Although the theoretical prediction is confirmed in the Low blocks, the results from the High blocks are ambiguous. In particular, we do observe endogenous bubbles and crashes in the first half of the High blocks. However, this characteristic pattern disappears in the second half of the High blocks, and the dynamics become more stable. We estimate the discrete choice model on the experimental data and find that the intensity of choice parameter is much lower for the High blocks, suggesting that subjects adapt their behavior and become less sensitive to payoff differences in a less stable environment – this is consistent with the findings in [2]. The reason that subjects adapt their behavior might be driven by the fact that, in the high cost environment, payoff differences tend to be highly volatile and unpredictable, and therefore may not perform well as a predictor of future success. Upon realizing this, subjects' response to past payoffs becomes weaker, which brings down this volatility in payoff differences endogenously. This is confirmed in the *High Long* treatment, which features more decision periods. Here the estimated values of the intensity of choice parameter indeed decrease over time, inducing a reduction in volatility. An interesting extension for future research would be a treatment with a fixed group of subjects, each of whom has to choose between the two alternatives again, but where – without informing the subjects directly – the costs associated with alternative A change at several instances during the experiment. Such a change in costs will effect the volatility of payoff differences, which may lead subjects to adapt their behavior.[23] A practical difficulty of running such a treatment is that it requires many decision periods and will therefore take a long time, with the risk that boredom or lack of concentration on the part of the subjects will effect the results. Nevertheless, based upon our *High-Low* and *Low-High* treatments, where such a change in costs was implemented, but accompanied by a reshuffling of the groups and a restart of the time series of payoffs, we conjecture that subjects will respond to volatility of profit differences and the estimated intensity of choice parameter will change over time.

Our results have important implications for the way choice behavior is modeled in heuristic switching models. They suggest that a model in which the intensity of choice parameter is endogenous, and depends positively upon some measure of volatility of payoff differences, potentially provides a better description of choice behavior. Adapting the benchmark model in this way may turn out to be quite relevant since the assumption that the intensity of choice parameter is exogenously given can impose a bias in the conclusions derived from theoretical heuristic switching models. Consider, for example, a volatile financial market that is described well by a particular heuristic switching model. On the basis of that model the financial regulator may want to implement a policy that – based upon numerical simulations – stabilizes market dynamics. However, if traders react to the increased stability and predictability of profits in this market by starting to respond more strongly to profit differences – as suggested by our experimental results – this may strongly mitigate the effect of the policy.

[23]Recall that in our experiment subjects do not observe the costs for the stabilizing heuristic directly, nor do they observe the evolution of the state variable x_t. That is, their choice has to be solely based on past payoffs.

Chapter 3

A Self-Tuning Model of Adaptive Choice

3.1 Introduction

Models of adaptive discrete choice are used to describe the behaviour of economic agents in situations where agents do not have full knowledge of the environment. The logit model is often used as a building block of adaptive choice modelling. That is, a probabilistic prediction of the next choice between several options is modelled as a logistic map from past performances. Different papers fitted the model to both experimental and real world data and reported good explanatory power. The problem, which motivates this research, lies in the fact that estimates of the logit parameter in the model exhibit high unexplained heterogeneity. Particular values play a pivotal role and generate qualitatively different model predictions and associated dynamics, and related calibration difficulties—for instance, which value to use for analysis—motivates a search of more robust logit model of adaptive choice. This chapter proposes a self-tuning modification of the logit model in the form of a scaling, where a maximum recently observed payoff difference serves as a scaling factor and allows the model to "self-tune" to the environment. This chapter uses data from two binary choice experiments, where participants were not aware of the payoff-generating processes and behaved adaptively. The data on their choices allows running a model contest between different model specifications using a Model Confidence Set (MCS) approach. Results indicate significantly stronger explanatory power of the proposed self-tuning model. The estimates of a single logit parameter in this model no longer exhibit high heterogeneity across different environments indicating that the chosen scaling solves calibration issues.

The logit specification is widely used in modelling adaptive choice due to its simple closed-form expression, micro-foundations, and good fit to the data. A good fit is generally achieved at the expense of a high heterogeneity in the estimated values of the logit parameter, both within and between the studies. This observation may not only suggest that the simple logit model suffers from misspecification, but more importantly that it challenges the external validity of the results. The question about which value of the logit parameter should be chosen for calibration or whether an interval for estimates can fit different data sets remains open. The natural research question motivated by these observations involves identification of a logit model specification that could account for the main drivers of differences in the logit parameter and stabilise the estimates around the value that is suggested for calibration.

I use data from laboratory experiments on adaptive binary choice with a design which allows to control for several possible compounding sources of heterogeneity in the logit parameter. The first possible feature of the data, which may drive differences in estimates, is different magnitudes of payoffs in different data sets. This explanation does not apply to the chosen data where payoffs in all experimental sessions are of a comparable scale. The second possible source of heterogeneity lies in the interplay between adaptive and strategic behaviour. It is often hard to distinguish those two features of behaviours, but this is not an issue in the experiments considered here. No information about the payoff-generating process was revealed to the participants. Strategic interaction was presented in a fashion not clear to the participants, making their task adaptive. The third explanation attributes heterogeneity to a possible misspecification of constructed beliefs over counter-factual payoffs—payoffs that participants could potentially receive in case of alternative choices. In the experiments considered, participants were informed about counter-factual payoffs, and therefore no risk of misspecification of payoffs is involved.

Despite the absence of the effects identified above, heterogeneity is still present in the estimates of the logit parameter based on the experimental data. This motivates construction of models which can effectively endogenise identified differences. I consider several modifications of the logit that satisfy three criteria: (i) the model is simple, (ii) the model is motivated by "stylised facts" of observed heterogeneity, and (iii) the model has a behavioural interpretation. All the models are fitted to the data using Maximum Likelihood Estimation, and both the Akaike and Bayesian information criteria are used for model selection. Still, discriminating models by their performance is challenging. First, as the number of models becomes large, a tractable way of evaluating relative performances using pair-wise comparisons becomes infeasible. Second, many models fit the data very well as they are built to capture the main features of the data. For these reasons I use the MCS approach that successfully solves both issues: it iteratively narrows the initial set of models by excluding models that are significantly worse in explanatory performance.

The results suggest that scaling the logit parameter with the largest payoff difference recently observed by an agent fits the data best. This model has a clear behavioural interpretation: agents channel less attention to the choice task if stakes are low compared to previously observed. Additionally, estimated values of the logit parameter in the self-tuning model turn out to have close values across experimental sessions, suggesting the external validity of the model and to a large extent solving the problem of heterogeneity. This finding can have important policy implications once the self-tuning model is incorporated into existing Heterogeneous Agent Models (HAM), which study price stabilisation government interventions.

The chapter is organised as follows. Section 3.2 introduces related literature that supplies "stylised facts" about the heterogeneity. Section 3.3 discusses the laboratory experiments that generated the data used in this study. Section 3.4 introduces the logit model of adaptive choice and several behavioural modifications. Section 3.5 discusses the estimation strategy, the MCS approach, and the results. Section 3.6 provides illustrative simulations and a discussion of the effects of the self-tuning feature of the model on the stability of the dynamics. Section 4.6 provides concluding remarks..

3.2 Related literature

Evidence of the heterogeneity of the logit parameter estimates is spread across different branches of the economics literature. There is no previous research that systematically collects, reports, and discusses the heterogeneity of within-studies estimates and indirect between-studies indicators of the heterogeneity. This chapter does not aim to perform such a meta-analysis either, but instead focuses on several influential models that incorporate logit choice to capture adaptive behaviour. Those models were estimated on the clean experimental data that controls for a number of the possible sources of heterogeneity and allows for the comparison between estimates.

The logit model of probabilistic choice can be derived from different assumptions and hence there are several interpretations of the logit parameter. The classical Random Utility Model (RUM) is built on the assumption that payoffs associated with possible choices have a publicly observed component and a privately unobserved stochastic component.[1] A choice has a logistic value distribution if an agent has the private component distributed as Type I extreme value.[2] In this interpretation, the logit parameter stands for the inverse of the variance of the private component distribution. Large values of the logit parameter are associated with small private shocks. From the perspective of the observer, who only knows the public component of payoffs, the large logit parameter makes the choices less noisy.

Alternatively, the logit model can be motivated by introducing inattention to the choice task. If the cost associated with attention—careful

[1] Main contributions include [106], [80] and [82].

[2] This distribution is also referred to as the Gumbel or the log-Weibull distribution. It is defined as a distribution of the maximum values of a sample of a random value realisation. Details of the logit choice derivation can be found for example in [107].

consideration of all available choices—is high, it is optimal for an agent to make choices probabilistically.[3] A large value of the logit parameter relates to a small cost of attention. In a similar spirit of the "rational inattention" model, the logit choice model is derived to be an optimal choice when information[4] about the payoffs or processing[5] this information is costly. Large values of the logit parameter are interpreted as cheap information acquiring or processing.

Flexibility in the logit choice interpretation contributes to the popularity of the model. I consider the following strands of economic literature that utilise the logit specification to model adaptive choice: learning (in games) models, quantal response models, and heuristic switching models. A number of studies fitted these models to experimental data and provided various estimates of the logit parameter. Figure 3.1 depicts collection of estimates for a number of models discussed further in the text.

Learning models describe well non-equilibrium behaviour observed in games that are played in experimental laboratories. These models specify how behaviour is adjusted as participants accumulate experience in the game and they often allow players to eventually "learn" equilibrium play. One of the most prominent examples—the Experience Weighted Attraction (EWA) model of [22] —accommodates both payoff-based reinforcement learning and belief-based learning. In EWA, both learning approaches contribute to updating the "attractors" of strategies:their fitness measures. Based on the values of the attractors, the agent makes a probabilistic choice following to the logit model. Camerer and Ho found that the logit model better fits the

[3] For more details see [109] and [110].

[4] For static version see [81], and for dynamic – [101]. Experimental evidence is discussed in [24].

[5] See [115] for details.

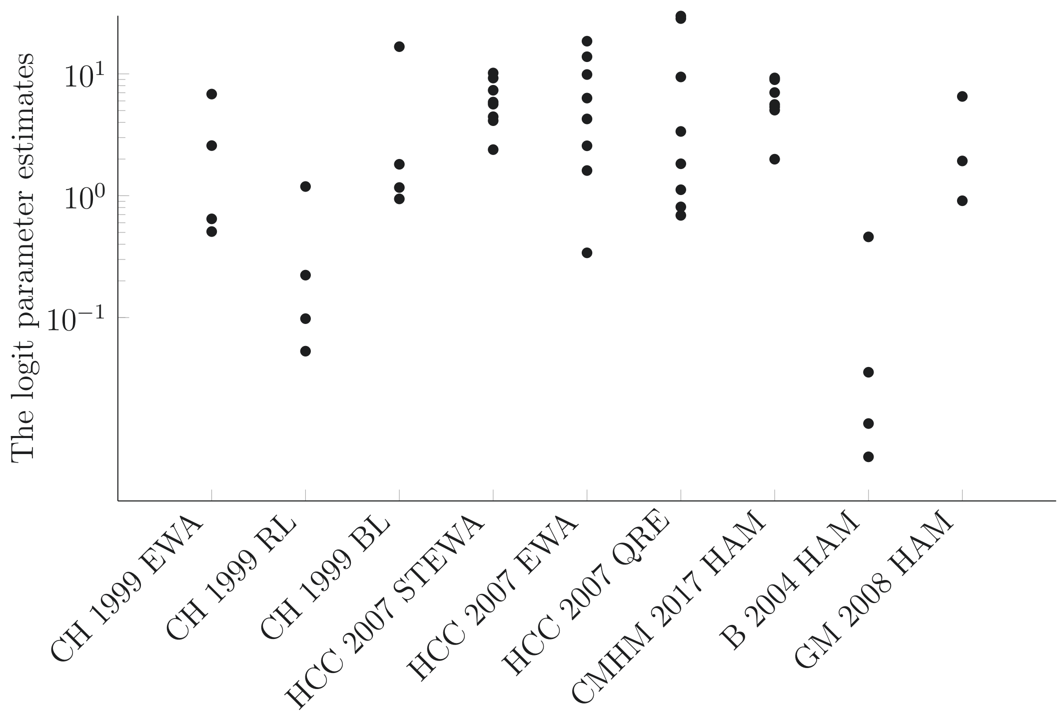

Figure 3.1: Examples of the logit parameter estimates for different models: Experience Weighted Attraction (EWA), Reinforcement Learning (LR), Belief-based Learning (BL), Self-Tuning Experience Weighted Attraction (STEWA), Quantal Response Equilibrium (QRE), and Heterogeneous Agent Model (HAM). Source: [22], [61], [34], [15] and [52].

data than alternative specifications. However, their estimates of the logit parameter, which here measures "sensitivity of players to attractions", vary significantly across games. Differences in the logit parameter generally affect the speed of learning process convergence[6], therefore different values predict qualitatively different convergence patterns. A large number of parameters in the original model motivated a self-tuning version of the EWA introduced in [61], where most parameters were replaced with functions of experience. In this version of the model, the logit parameter is the only "game specific" non-tuned parameter and the estimates exhibit high heterogeneity. Seemingly, the logit parameter is the most challenging parameter to be fitted to the data.

[6]In [11] an importance of flexibility in the logit parameter values for the convergence to the equilibrium and its speed in the EWA was demonstrated.

Quantal Response Equilibrium (QRE), introduced by [83], does not accommodate learning but offers a static version of the "noisy" Nash equilibrium.[7] Players calculate expected payoffs and choose actions probabilistically following the logit model. To calculate expected payoffs, players incorporate "noisiness" in the choices of other players, assuming they share the same logit parameter value. This model generally fits the experimental data well, however, that comes at the expense of an unexplained variation in the logit parameter between games. To capture learning effects, the model incorporates a gradual increase in the logit parameter over time. For example, in [83], estimates of the logit parameter increase from 0.17 to 0.59 based on the data from the [75] experiment, and from 1.95 to 4.64 for [86]' game. Different values of the logit parameter define how far from the Nash equilibrium the QRE predictions are.

In Heuristic Switching Models (HSM), agents make a choice between a fixed number of forecasting strategies framed as "heuristics", which differ in generated payoffs.[8] The setup of these models and associated dynamics is the closest to this chapter, and the logit parameter is referred in this literature as the Intensity of Choice (IoC) parameter. The interpretation is similar to the RUM: as the IoC parameter increases, choices get closer to best response. HSMs were fitted to a number of laboratory experiments that study forecasting strategies.[9] In [5] and [7], the IoC parameter varies significantly between different sessions. Heterogeneity of the logit parameter is high both within and between these studies. In HSM, an increase in the IoC

[7]For more details please see [87].

[8]For early theoretical contributions see [17] and [18]. The logit parameter plays a central role in this framework: the model dynamics vary from stability to chaos depending on the particular value of the logit parameter.

[9]Estimations of HAM with the use of financial market data, surveys on inflation expectations and mutual fund allocation decisions indicate heterogeneity in the IoC parameter estimates.

parameter leads to qualitatively different predictions about the stability of the dynamics; more detailed discussion is provided in Section 3.6. Heuristic switching models have a complicated interplay between the IoC parameter and other model components, which makes identification of regularities in differences of estimates infeasible. To address this issue, several laboratory experiments were explicitly designed to study adaptive choice, or “switching mechanism”, and the data from these experiments are used in this study.

3.3 Data

I use data from two laboratory experiments designed to study adaptive choice behaviour: [2] and [3], which I refer to as the ABT16 and the ACT18 respectively. The data collected in these experiments are ideal for studying the adaptive choice behaviour, including the estimation of the logit parameter.

The two experiments share a number of features in their design that effectively make the combined data set a meta-experiment on adaptive binary choice.[10] Both ABT16 and ACT18 experiments were organised as individual task experiments where participants were asked to choose between two options over 40 periods.[11] Participants were not informed of the exact process

[10] There are several close in spirit experiments, which were designed to study adaptive choice behaviour in a limited information environments, for example, [42], [49], and [84]. The data from the binary choice experiments were not included in this study because at least one of the two important requirements were not presented in the experiments. The first requirement is a non-disclosure to participants and information regarding payoff-generating rules. This requirement ensures adaptivity of the choice, which is based only on payoffs. The second requirement is non-convergence of the payoffs to fixed values during the experiment. This requirement guarantees sufficient variation in the choice task and attention of participants to payoffs during the experiment.

[11] There are several variations in the basic setup. For two groups of the ACT18 experiment, a number of periods was extended to 60. In the ABT16 experiment, participants had access to the initial 10 periods of payoffs’ history.

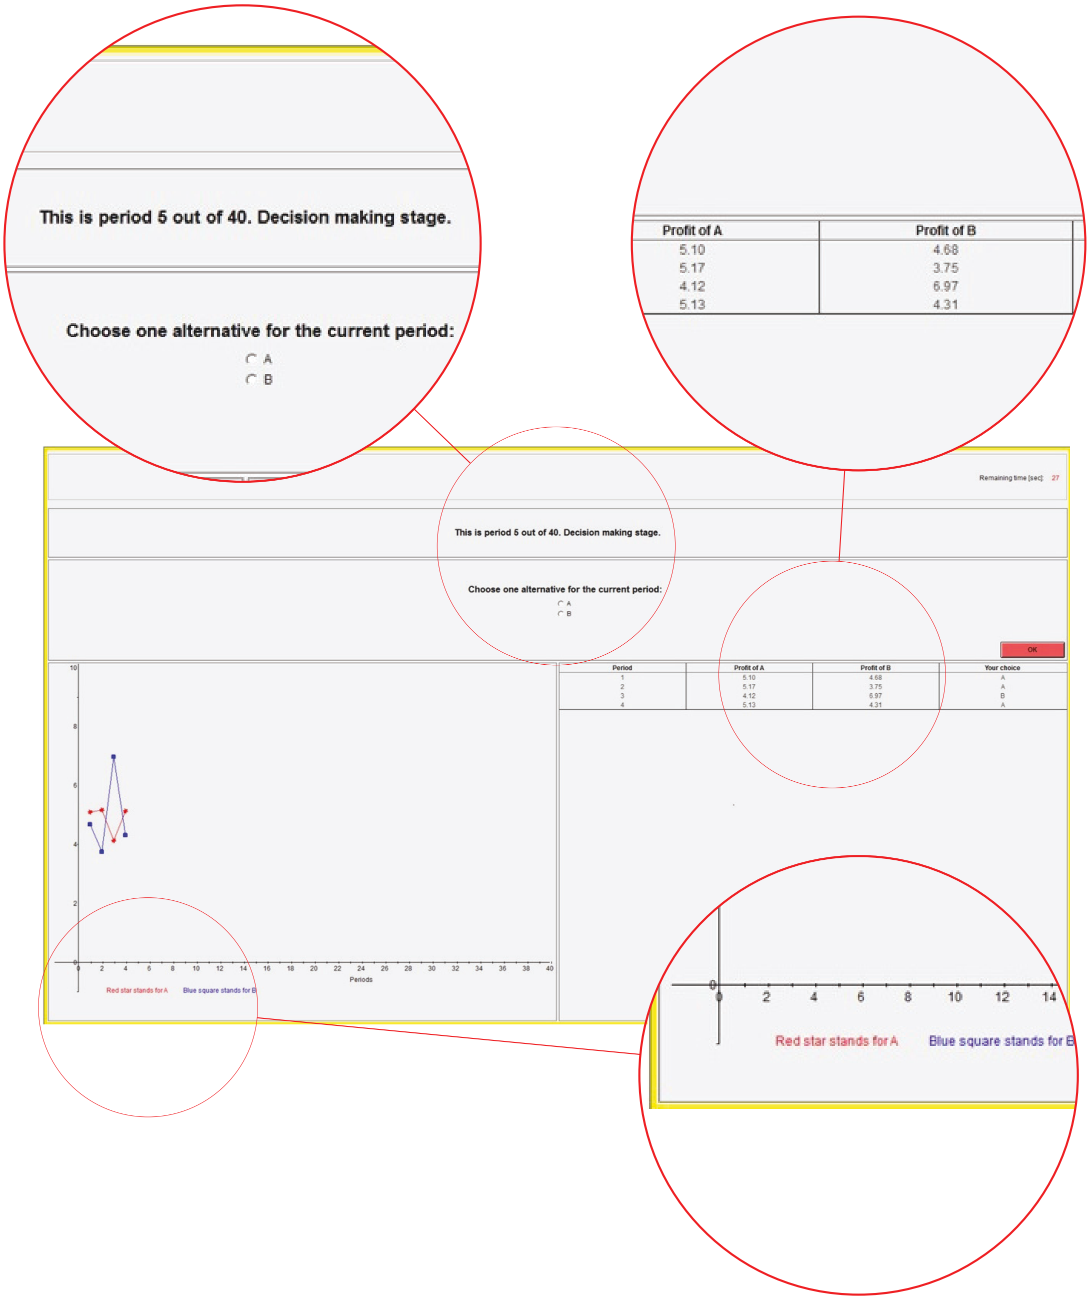

Figure 3.2: An example of the screen used in experiments. Participants make their decision by clicking on one of two options: A or B. In order to make their choice, participants observe information on profits of A and B from all previous periods in two formats: a graph and a table.

that generates payoffs for the two options. These processes were exogenous and endogenous in these two experiments, respectively, as explained below. Participants were told to choose between two neutrally labelled options, “A” and “B”, based on the history of payoffs from the earlier periods. Information was provided on past payoffs of both options. This experimental design feature distinguishes this setting from other similar experiments on repetitive binary choice and substantially simplifies analysis of the choice, releasing from the necessity to model beliefs over counter-factual payoffs.

The final monetary payments for participation in both experiments were calculated based on the realised payoffs of the options that had been chosen during the experiment. Payoffs were denominated in experimental points with a fixed exchange rate that determined monetary payment in local currency. In the ABT16 experiment, the final payment consisted of a cumulative sum of payoffs, while participants in the ACT18 experiment were paid for two randomly chosen periods for the option they chose.

The available options were described to participants as “investment funds” in the ABT16 experiment instructions, and as “investment alternatives” in ACT18.[12] Information on payoffs from previous periods was provided to participants in the form of a table and a graph – the example of the screen layout is provided in Figure 4.2. Each experiment included several treatments, which differed in the payoff generating process, and were designed to study adaptive choice of participants facing different series of payoffs.

I use data from three treatments of the ABT16 experiment.[13] In these treatments, participants chose between two options for which payoffs were pregenerated, and three treatments differed in the payoff generating rules: white noise, Brock-Hommes model simulation, and stock indexes. I refer to

[12]For the details of qualitative description of the tasks, please, refer to experiment instructions available in the appendixes of corresponding papers.

[13]Other treatments had participants choose between three and more options.

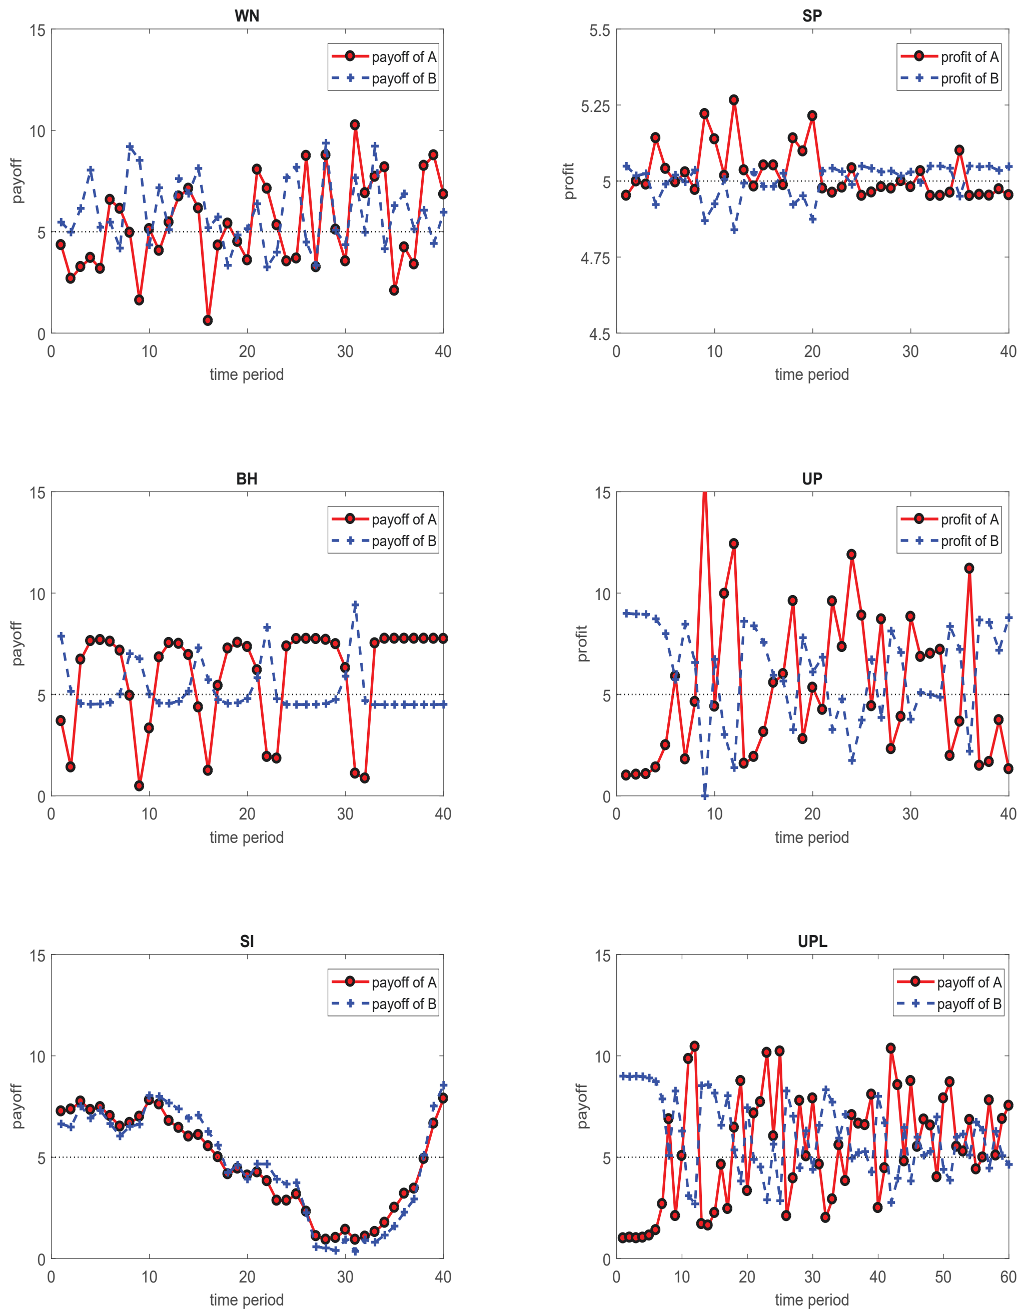

Figure 3.3: Examples of payoffs series in different treatments. Left panels depict treatments from ABT16 experiment: White Noise (WN), the Brock-Hommes model simulations (BH) and Stock Indexes (SI). Right panels depict treatments from ACT18 experiment: dynamics of stylized version of the Brock-Hommes model parametrised to generate stable payoffs (SP), unstable payoffs (UP), and unstable payoffs in long session (UPL). At each period of time only realised history of payoffs was available to participants.

them as WN, BH, and SI, respectively. See Figure 3.3 for illustration, where, for the ABT16 experiments, I excluded the first 10 periods when participants did not have to make the choices. In WN treatment, every payoff was a sum of a fixed value of 5 and a noise term distributed as independently and identically distributed standard normal. In the BH treatment, payoffs followed quasi-cyclical patterns that were generated by numerical simulations of the [17] model, which generates chaotic fluctuations of payoffs due to non-linearity in the system dynamics. In the SI treatment, two payoff streams were generated by adding 5 to the normalised average yearly returns for each month from 2008 to 2012 of two stock indexes: the Austrian Trade Index (ATX) and the Belgium 20 Stock Index (BFX). All three treatments were designed to have similar average payoffs but have different structure of payoff dynamics, ranging from no autocorrelation in WN to strong autocorrelation in SI. Examples of payoffs series used in the experiment are presented in the left panels in Figure 3.3.

I use all the data from the ACT18 experiment. The design of this experiment is very similar to the ABT16 experiment; the primary difference lies in the payoff generating rules. In the ACT18 experiment, payoffs were not pregenerated but evolved during the experiment depending on the choices of participants. The rules that determined payoffs were based on a stylised version of the Brock-Hommes model, which effectively combined the main features of [17] and [18] versions of the model. The mechanics of this model can be summarised as follows. Agents choose between two options to form their expectations over future price. One option is "free" to use, but is a destabilising naive prediction under which the price deviates from the "fundamental price" benchmark. Another option is "costly", but a stabilising rational expectation under which the price quickly returns to the "funda-

Table 3.1: Descriptive statistics of payoffs of two options in the ABT16 and AC18 experiments.

Block	Option	Mean	Median	Min	Max	Range	Variance
WN	A	5.33	5.12	0.60	10.25	9.65	5.04
	B	5.94	5.47	3.26	9.37	6.11	2.92
BH	A	5.96	7.36	0.47	7.75	7.28	6.08
	B	5.20	4.60	4.50	9.42	4.92	1.46
SI	A	4.65	4.68	0.93	7.89	6.96	5.73
	B	4.68	4.89	0.36	8.55	8.19	6.92
SP	A	5.02	4.98	4.95	5.60	0.65	0.02
	B	5.00	5.03	4.61	5.05	0.44	0.01
	A	5.02	4.99	4.95	5.27	0.31	0.01
	B	5.00	5.03	4.84	5.05	0.21	0.00
	A	4.99	4.98	4.95	5.15	0.20	0.00
	B	5.02	5.03	4.91	5.05	0.14	0.00
	A	5.01	4.98	4.95	5.35	0.39	0.01
	B	5.01	5.03	4.79	5.05	0.26	0.00
	A	5.04	5.00	4.95	5.84	0.89	0.02
	B	4.99	5.02	4.46	5.05	0.59	0.01
	A	4.99	4.98	4.95	5.15	0.20	0.00
	B	5.02	5.03	4.92	5.05	0.13	0.00
	A	5.03	5.01	4.95	5.36	0.41	0.01
	B	5.00	5.01	4.78	5.05	0.27	0.00
	A	5.02	5.00	4.95	5.46	0.51	0.01
	B	5.00	5.02	4.71	5.05	0.34	0.00
UP	A	5.89	3.92	1.00	32.96	31.96	38.40
	B	6.18	7.05	0.00	9.00	9.00	7.30
	A	5.41	4.28	1.00	19.44	18.44	20.47
	B	6.14	6.81	0.00	9.00	9.00	7.75
	A	6.05	4.79	1.00	41.90	40.89	45.57
	B	6.13	6.47	0.00	9.00	9.00	5.65
	A	5.36	4.41	1.00	16.06	15.06	14.20
	B	6.12	6.73	0.00	9.00	9.00	5.96
	A	5.20	4.32	1.00	16.80	15.79	14.97
	B	6.26	6.79	0.00	9.00	9.00	5.72
	A	5.03	4.45	1.00	16.34	15.34	10.15
	B	6.34	6.70	0.00	9.00	9.00	4.07
	A	4.68	3.39	1.00	18.31	17.31	18.41
	B	6.61	7.41	0.00	9.00	9.00	7.16
	A	4.96	3.34	1.00	16.17	15.17	16.29
	B	6.39	7.44	0.00	9.00	9.00	6.84
UPL	A	5.61	5.58	1.00	13.76	12.76	6.26
	B	5.93	5.95	0.49	9.00	8.50	2.78
	A	5.39	5.40	1.00	10.45	9.45	7.59
	B	6.07	6.07	2.70	9.00	6.30	3.37

mental" value. The dynamics of the payoffs and their stability depend on the cost of rational expectations: the dynamics are stable with similar payoffs of two strategies for low cost, and they is unstable with large differences in payoffs for high cost. The ACT18 experiments were run with two treatments: sessions with a low cost of rational expectation option and resulting stable payoffs, denoted by SP, and a high cost of rational expectation option and resulting unstable payoffs, denoted by UP and UPL. The sessions of UPL differ in two dimensions: larger number of participants in the session and longer sessions. In the standard SP and UP sessions, there were 40 periods and 10 participants in each of 8 groups per treatment, while in two sessions of UPL there were 60 periods and 33 participants in each group. Participants of the same group observed the same payoffs, and decisions of all group participants affected the dynamics of payoffs in the following way: as an option became more popular, it generated smaller payoffs, while a less popular option tended to have increasing payoffs. Just as in the ABT16 experiment, participants were neither informed about the payoff generating rules, nor about the effects of their decisions on subsequent payoffs. The right panels in Figure 3.3 demonstrate examples of the payoffs series in the ACT18 experimental sessions.

Table 3.1 contains descriptive statistics of payoffs observed by participants in different treatments. To avoid possible effects of the payoffs scale on the choice task, all treatments were designed to have similar values of payoffs: the mean and median values of both options were approximately equal to 5 experimental currency units. At the same time, the treatments exhibit variation in stability and predictability of payoffs: from a small variance in SP sessions to a high variance in UP. In the next section, I discuss the standard model of adaptive choice which demonstrates a good fit within each of the treatments. However, this good fit is achieved using treatment-specific values of the logit parameter. Therefore, the main challenge of building a portable version of the model of adaptive choice is to find an extension of the model which effectively endogenises identified between-treatments differences in subjects' behaviour.[14]

3.4 The Model

Presented research utilises the logit to model binary choice for several reasons. First, the logit specification is commonly used in various models from different strands of literature, and therefore a modified version of the logit can be easily incorporated into existing studies.[15] Additionally, the logit specification, as opposed to the probit, has a simple closed form expression for choice probabilities, which substantially simplifies the analysis of logit-based models. Finally, comparisons of the previously estimated models suggest that the logit form outperforms its alternatives of probit and power distributions at various instances, see [22] for discussion and further references.

3.4.1 The Logit model

In the basic logit discrete choice model, the probability to pick option k out of N alternatives, P_k, is a function of expected payoffs for all available options:

[14] The idea of portability of the model was inspired by the notion of Portable Extension of Existing Models (PEEM) introduced in [90]. In the course of presented research, psychological realism is introduced in the model by specifying the laws of the logit parameter adjustments.

[15] All structural models of HAM literature use logit specification, see for example [14]. It is also true for QRE and EWA models.

$$P_k = \frac{\exp(\beta \pi_k)}{\sum_{n=1}^{N} \exp(\beta \pi_n)}, \tag{3.1}$$

where π_k is the payoff of the option k, and β is the logit parameter.

The logit model has intuitive interpretations which relate to the adaptive choice: the probability to choose an option is positively affected by an increase in payoff of that option and negatively affected by an increase in payoffs of the alternative options. Additionally, while agents are modelled as making their choices imperfectly, by choosing an inferior option with a positive probability, agents make better choices with a higher probability if the differences between payoffs are more apparent. The value of the logit parameter determines the accuracy of the decisions: for $\beta = 0$, an agent makes choices at random and equally likely picks each of the options; for $\beta > 0$, an agent tends to pick an option with a larger payoff; for $\beta \to \infty$, an agent makes best-response choice and picks the option with the largest payoff.

Payoffs are modelled differently depending on the context. In QRE models, payoffs are calculated using equilibrium noisy best responses of other players in the game, and the variance of errors is proportional to the inverse of the logit parameter. In EWA models, payoffs are modelled as "attractors", which are based on both reinforcement and belief-based learning, and agents in the model "logistically" respond to the values of attractors. In HAMs, large markets are considered, and participants do not have strategic power, with only negligible effect of individual actions on the aggregate distribution of payoffs. Agents also do not have a knowledge of the environment and behave adaptively in the model. Additionally, payoffs not only from the chosen option, but from all options, are assumed to become common knowledge once realised. In this chapter, I follow the HAM literature approach, which utilises adaptive choice modelling: expected payoffs are replaced by performance measures $U_{k,t}$, which are calculated based on payoffs from previous periods. This adaptive version of the logit choice model defines the probability to pick option k in period t, $P_{k,t}$, as a function of performance measures for all options $1,..,N$ from period 1 to $t-1$ as follows:

$$P_{k,t} = \frac{\exp(U_{k,t})}{\sum_{n=1}^{N} \exp(U_{n,t})}, \tag{3.2}$$

where $U_{k,t}$ is the performance measure of option k at period t.

Performance measures, which are used to define choice probability in equation (3.2), are generally constructed using realised payoffs of the options (e.g., past profits of strategies). Updating performance measures over time allows incorporating all the history of previously observed payoffs, but in practice, this option is rarely used and agents are modelled as having limited memory.[16] This approach is consistent with the so-called "recency effect", which captures the observed tendency of people to discount past information and make decisions based on the recently observed outcomes (see [41] for an example). Consider the following performance measure that incorporates information on L lags of observed payoffs:

$$U_{k,t} = \alpha_k + \beta_{k,1}\pi_{k,t-1} + ... + \beta_{k,L}\pi_{k,t-L}, \tag{3.3}$$

where α_k is the predisposition effect towards option k, $\beta_{k,l}$ is the intensity of choice with respect to payoff $\pi_{k,t-l}$.

[16] Standard way to incorporate history of previous payoffs on choices is using geometrically declining weights of past payoffs. This approach was used for theoretical modelling, for example in [67], but the empirical evidence is mixed, see [2] for discussion.

The predisposition effect was estimated to be a significant determinant of the choice in [2]. The choices observed in all sessions of the ABT16 experiment can be explained by model (3.2)-(3.3), if performance measures include the predisposition effect[17] and up to three lags of payoffs with symmetric for different options and identical across agents' intensities of choices.[18] Following the existing results, I set logit parameters to be equal for the same lag of payoffs, that is $\beta_{k,l} = \beta_{n,l}$ for every option k and n for a fixed lag l.

To model choice of participants in the ABT16 and ACT18 experiments, I combine the logit choice defined by equation (3.2), and performance measure defined by equation (3.3), and apply to the experimental setting where choices are binary and are made between options A and B. A probability to pick option A at time t is then defined as follows:

$$P_{A,t} = \frac{\exp(\beta_1 \pi_{A,t-1} + ... + \beta_L \pi_{A,t-L})}{\exp(\beta_1 \pi_{A,t-1} + ... + \beta_L \pi_{A,t-L}) + \exp(\alpha + \beta_1 \pi_{B,t-1} + ... + \beta_L \pi_{B,t-L})}, \tag{3.4}$$

[17]The logit model in its canonical specification does not accommodate any behavioural bias towards any option. At the same time, only slight modification by adding a constant to performance measures introduces predisposition effect, which has also been found to be a significant determinant of choice in [15].

[18]In this chapter I also assume identical logit parameters for all participants, since the prime interest lies in time-variation of the parameter. In [114] importance of possible heterogeneity in logit parameter across participants is demonstrated with the use of simulations, and possible biases in estimation of weighting parameter, which shows relative importance of actual and simulated effects. Indeed, some evidence of heterogeneity of the logit parameter was identified in [62], where individual estimations were clustered in two distinct groups. On the contrary, individual estimations on the data utilised in this chapter are close to aggregate estimation of the logit parameter. I attribute this result to more accurate estimation of the choice function parameter since both factual and counter factual payoffs were available to participants, see Section 3.3 for details.

where L is a maximum lag form, the history of payoffs included in performance measurement, $\alpha > 0$ reflects predisposition effect towards option B[19], and $\alpha < 0$ reflects predisposition effect towards option A.

Equation (3.4) defines a very simple–yet very powerful in terms of explanatory power–model of adaptive binary choice. The simplest version of the model, which includes only one lag of payoffs in the choice decision, outperforms in explaining experimental data and more sophisticated models with larger number of parameters, for example, and those with weighted averaged histories of payoffs or asymmetric intensities of choice. Similarly to previously discussed empirical studies, the good fit in different experimental sessions is achieved at significantly different values of the logit parameters β. To address this issue, I consider several modifications that allow the logit parameter in the model to self-tune to existing conditions, meaning that the logit parameter varies given observed payoffs.

3.4.2 Self-tuning models

Estimated logit parameter values in the model appear to depend on the environment experienced by participants. Figure 3.4 illustrates the estimates which are clustered together for different treatments. Treatments are ordered in accordance with stability of the payoff dynamics measured by the number of significant lags in the autocorrelation function, which increases from the side of the graph to the right side. This observation suggests that the differences in the dynamics of payoffs in different treatments of the

[19]Intuition of the predisposition effect can be demonstrated via dividing both numerator and denominator of equation (3.4) by the numerator. Obtained expression clearly shows that increase in α reduces the probability to pick option A, and vice versa. Because of the issue of jointly unidentifiable constants, predisposition towards option A in corresponding performance measure is set to zero.

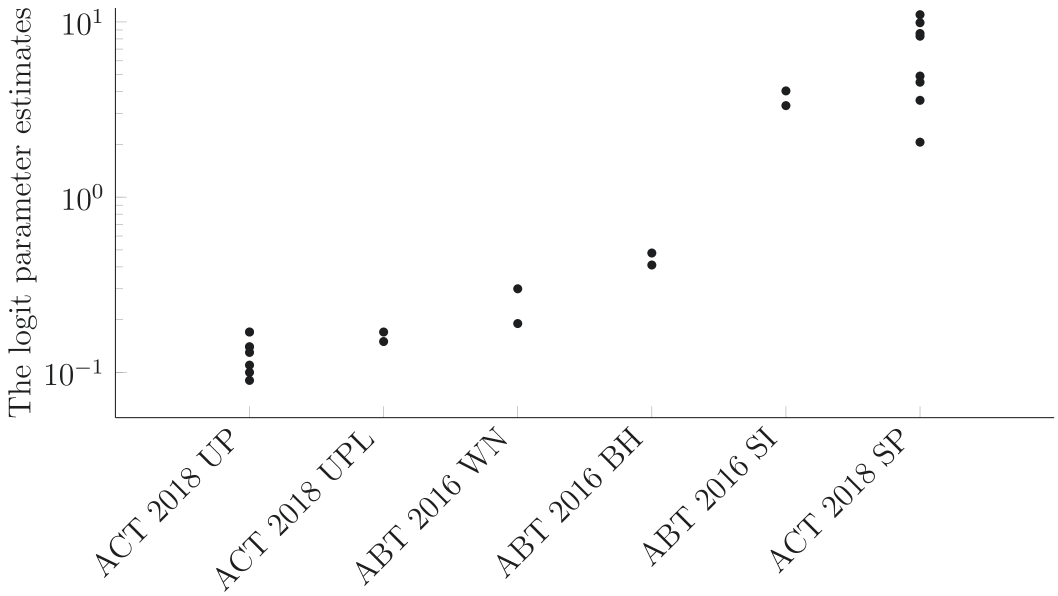

Figure 3.4: Examples of the logit parameter estimates for the model with previous period payoff as a performance measure on data from six treatments of ABT16 and ACT18 experiments.

ABT16 and ACT18 experiments affects the parameter estimates. Allowing agents in the model to adjust the logit parameter was previously acknowledged as an important factor of dynamics convergence, for example, in [11]. A model with changing sensitivity to reinforcing, introduced in [40], can serve as an example of the self-tuning mechanism. Self-tuning mechanisms of the logit parameter adjustment in this study are specified in the form of a scaling: performances are divided by the normalisation factor, similarly to the approach introduced in [105]. Because of linearity in transformation, one might treat scaling as the logit parameter adjustment.

Normalisation in the form of a scaling has a number of advantages. First, normalisation unifies the data by making estimates unit-free. It increases the portability of the model since empirical studies combine data from different markets, such as commodity, equity, or foreign exchange. Second, normalisation reduces sensitivity of the estimations to the extreme values, which are often presented in the data. Finally, scaling factors that are considered in this study have clear economic interpretations, such as accounting for predictability of time-series or relative intensity of financial stimulus. In such a manner, the self-tuning version of the model captures behavioural aspects of adaptive choice.

Construction of the self-tuning model that fits different experimental treatments requires specification of the scaling factor. This factor will reflect the main features of payoffs dynamics that potentially drive observed differences in the intensity of choices. To specify different models of scaling, I use the following "stylised fact" based on existing estimates: in the experimental sessions, where a higher level of instability in the dynamics of payoffs was observed, the lower estimated values of the logit parameter were obtained.[20]

Consider a simple version of the model (3.4) with no predisposition effects and one lag from payoff history for the sake of simplicity. The expression for choice probability in this case can be simplified in the following way:

$$P_{A,t} = \frac{\exp(\beta\pi_{A,t-1})}{\exp(\beta_1\pi_{A,t-1}) + \exp(\beta\pi_{B,t-1})} = \frac{1}{1+\exp(\beta(\pi_{B,t-1} - \pi_{A,t-1}))}.$$

This version of the model illustrates the importance of differences in payoffs rather than absolute values of payoffs for the choice. The self-tuning

[20] This relation was pointed out in [2] and supported by the analysis in [3]. Additionally, in [23], a footnote to Table 2 indicates a gradual adjustment of the logit parameter, and in particular, an increase over time, which can be associated with stabilisation of the payoffs due to convergence towards the equilibrium. One might expect that in a highly volatile environment subjects do not pay attention to payofss and do not try to chase them. On the contrary, in stable environment it is easier for subjects to focus attention and make more accurate choices based on the observed payoffs.

model can be constructed by introducing the scaling to payoff difference, that is, dividing by normalisation factor, in the following way:

$$P_{A,t} = \frac{1}{1 + \exp\left(\frac{\beta(\pi_{B,t-1} - \pi_{A,t-1})}{F_{t-1}}\right)}, \tag{3.5}$$

where F_{t-1} is a scaling factor, which is updated over time depending on recently experienced conditions.

Equation (3.5) applies a linear transformation to the payoff difference. It allows for an interpretation of the self-tuning mechanism not only as a payoff scaling procedure, but also as the logit parameter adjustment. When the scaling factor F_{t-1} is high, the "effective" value of the logit parameter becomes smaller, indicating that agents would pay less attention to the past payoff differences. Different scaling factors are considered below. All factors have behavioural interpretation and appear to capture important characteristics of the payoff series in different treatments. Scaling brings context dependency to the choice, while different scaling factors capture different aspects of the contexts, which can be important for the choice. In this sense, all the following specifications satisfy three criteria outlined in the introductory section: simplicity, behavioural interpretation, and empirical support. For several particular functional forms of the scaling factor, we use "recent history" of observations, which is the number of lags k of previous payoffs.[21] Four different factors are considered. Two factors capture possible effects of predictability of payoffs on "intensity" of choice, and two factors capture effects of relativity of payoffs.

[21] In Section 3.5, which describes estimation procedure, the length of the recent history of observations is fixed to 10. This length appears to be a reasonable compromise between adaptivity of the choice function and memory constraints of the agents. Robustness of the results and importance of particular length fixing are discussed in Section 3.5.4.

Standard Error Factor. The mean value and the standard error are often used for normalisation in econometrics and statistics.[22] Consider the following factor to scale the profit difference:

$$F_{t-1} = \sqrt{\frac{1}{k-1}\sum_{l=1}^{k}\Big((\pi_{B,t-l} - \pi_{A,t-l}) - \overline{(\pi_B - \pi_A)}\Big)^2}, \qquad (3.6)$$

where k is a number of lags in the history payoffs, $\overline{(\pi_B - \pi_A)}$ is an average difference in payoffs for the last k observations.

The standard errors serve as a measure of (in)predictability of the payoff difference observed by the agent. Thus, this normalisation procedure can be interpreted as the process of attention adjustment: as standard errors increase, the logit parameter is adjusted down, and time-series of past payoff differences become less important for the future choices.

Range Factor. An alternative statistical measure, which can reflect the dispersion of the past payoffs, is the range of payoff differences. The range is calculated as a difference between the largest and the smallest observed values. The range factor for scaling the profit difference takes the following form:

$$F_{t-1} = \max_{1\leq l\leq k}(\pi_{B,t-l} - \pi_{A,t-l}) - \min_{1\leq l\leq k}(\pi_{B,t-l} - \pi_{A,t-l}). \qquad (3.7)$$

The behavioural interpretation of this normalisation factor is similar to standard error factor: agents tend to channel less attention to payoffs if the

[22] Using standard error as a scaling factor is similar to standardisation and subtraction of the mean is dropped since the differences are considered.

observed dynamics are unstable, which is reflected in the large value of the range.

Maximum Absolute Value Factor. To capture the possible effects of relativity of payoffs, the maximum of the absolute values of past payoff differences is used as a factor. This scaling factor is defined as:

$$F_{t-1} = \max_{1 \leq l \leq k} |\pi_{B,t-l} - \pi_{A,t-l}|. \tag{3.8}$$

The interpretation of this self-tuning model is the following. Agents pay less attention to the observed difference in payoffs, which is captured by lower values of the logit parameter, if this difference is relatively small in comparison to the largest difference recently observed. The largest difference serves as a proxy for the importance of the current choice: the logit parameter is scaled down for large values of this factor.

Relative Performance Factor. Using the sum of payoffs for scaling was previously applied in [105] and [104]. This specification is motivated by the attempt to combine data from different sources and use unified data in the range of $[-1; 1]$ for estimation. The scaling factor is defined as follows:

$$F_{t-1} = \pi_{A,t-1} + \pi_{B,t-1}. \tag{3.9}$$

This normalisation suggests using a relative profit difference rather than absolute value. A behavioural interpretation is that agents make more accurate choices if the level of payoffs is smaller. This self-tuning model serves as a natural benchmark, as it has been used in the literature already. Alternative specifications introduced above will be compared to this benchmark.

All self-tuning models of choice combine different specifications of performance measures with different scaling factors. In the next section, devoted to estimation, attention is focused on performance measures that include predisposition and up to 3 lags of payoffs, and combines it with 5 scaling factors, including a non-scaled model. The set of competing models consists of 31 different models.[23] Building a competitive set of models insures that the model, which survives the model competition, has superior explanatory power in both absolute and relatively terms.

3.5 Estimation

Models are estimated using Maximum Likelihood. As one might expect, models fit to the data and significant parameter estimates are obtained for most specifications. The main challenge is to identify the model which explains the data better than the others. A general approach for how to choose from competing models would be to select the best model in terms of value of log-likelihood function.[24] The problem with this approach lies in the fact that generally, behavioural models are built to accurately capture observed stylised facts, and therefore all models fit to data soundly. In this case, an

[23] This number is obtained as follows. Modification with solely predisposition effect is not affected by the scaling procedures. There are six different models which accommodate different combinations of predisposition effect and payoff lags inclusion. Initial model and four scaling rules allow to estimate $1 + 6 \times (1 + 4) = 31$ distinct models.

[24] In fact, this approach is popular tool for competitions between learning models, see [22] or [6] for examples. Compared models are usually non-nested, and therefore, likelihood-ratio tests are not applicable. Alternative approaches include: non-parametric tests for mean squared distances as in [96] and [21], tests with the use of posterior odds criterion as in [45] and [56], and comparisons of the simulated dynamics as in [32].

additional selection procedure is required to identify the model which fits significantly better than the others.

3.5.1 Model contest

In order to obtain insight into whether some models perform significantly better than others, I follow the Model Confidence Set (MCS) approach proposed in [55]. This process of iterative elimination of poorly performing models generally allows for a reduction of the number of models under consideration.[25] There are several advantages that make the MCS an ideal instrument to use in the model contest.[26] First, the MCS approach allows for testing of the relative predictive power of a large number of models at once. To test the relative performance of a particular model by using standard techniques such as the Likelihood Ratio test, every model must be evaluated against every other available model. In a situation with 31 competing models[27], this approach will lead to results from 465 tests[28], which is infeasible to systematically interpret afterwards. On the contrary, the MCS approach identifies a subset of models from the initial set, which contains models of equally good preformance, and generates a list of excluded models, which are ranked in accordance with diminishing performance. Second, the MCS approach is suitable for testing the relative performance of non-nested

[25]Model contest is a usual approach in situations where estimations of different models give mixed results. For example, in [47] seven competing specifications estimated with the method of simulated moments are compared for several separate criteria.

[26]The MCS approach was mainly motivated by the applicaion in macroeconomics, which was provided in [55], and finance context, for example, competition between models of different copula specification for forecasts presented in [37]. To my best knowledge, this chapter is the first application of the MCS approach to selecting between behavioural models.

[27]This number is not large in comparison with other empirical papers employing the MCS approach, for instance, 600 models are compared in [76].

[28]$C(31,2) = \frac{31!}{29!2!} = 465$.

models, which is the case in the presented research. Finally, the robustness check–which is provided by the MCS–prevents picking a wrong model and appears to be the best performing approach due to some unfortunate coincidence in the noisy data, which is usually the case for data for behavioural models.

3.5.2 Maximum Likelihood Estimation

I apply the MCS approach to model selection according to two information criteria based on log-likelihood measure. Parameters of the model are obtained by maximisation of the value of the log-likelihood function, which is calculated based on one-period ahead forecasts on probability to choose a particular option by every participant of the experiment. Hereafter, I restrict the attention to the probability to choose option “B” since the probability to choose option “A” is always uniquely determined as a complimentary probability. The number of participants in each session is denoted by N and the number of periods by T. For each experimental session block, I use the data on choices $y_{i,t}$, where $y_{i,t} = 1$ if option “B” is chosen, and $y_{i,t} = 0$ otherwise. The information on choices is available for all participants with $i = 1, .., N$ and all time periods $t = 1, .., T$. The joint likelihood function is defined in the following way:

$$L = \prod_{i=1}^{N} \prod_{t=1}^{T} \left(P(y_{i,t} = 1)^{y_{i,t}} P(y_{i,t} = 0)^{(1-y_{i,t})} \right).$$

I illustrate the mechanics of this approach by considering the simple model without scaling defined by equation (3.2). The corresponding log-likelihood function, as a function of parameter β, is:

$$LL(\beta) = \sum_{i=1}^{N}\sum_{t=1}^{T}(y_{i,t}ln(\frac{1}{1+e^{\beta(\pi_{A,t-1}-\pi_{B,t-1})}}) + (1-y_{i,t})ln(1-\frac{1}{1+e^{\beta(\pi_{A,t-1}-\pi_{B,t-1})}})).$$

I use an unconstrained numerical minimisation procedure **fminunc** implemented in Matlab to estimate β, which minimises the value of the negative log-likelihood function.[29] I use the Fisher Information matrix to obtain standard errors for the estimated values. The estimation procedure will be analogous for all models considered further in the text.

Table 3.3 contains results of the estimation of the model defined by equation (3.2). Estimated values of the logit parameter β for each experimental session, as well as for the pooled data, are reported under the "Standard Model" title. Estimates are significant in all sessions. Reported standard errors indicate that the logit parameter estimates are similar within treatments but have statistically significant differences across treatments. This observation relates to the calibration issues raised earlier, and it motivates the search of the model which endogenises identified differences and gives a close estimate.

3.5.3 The Model Confidence Set

The basic intuition of the MCS approach can be described as follows.[30] For each model in the initial set, an explanatory power is measured, for example,

[29]Switching to dual problem does not affect the procedure and in fact is rather technical as there is no unconstrained numerical maximisation procedure available.

[30]The procedure is described in details in Appendix B.2.

Table 3.2: Results of the MCS procedure for AIC and BIC. For each model the order of elimination from the set and corresponding probability of false rejection are reported.

		AIC		BIC	
№	**Scaling model and performance measure model**	**Order of exclusion**	**Probability of a false exclusion**	**Order of exclusion**	**Probability of a false exclusion**
	No scaling				
1	Predisposition	2	0.000	2	0.000
2	Predisposition and one lag	18	0.000	18	0.000
3	Predisposition and two lags	17	0.004	17	0.000
4	Predisposition and three lags	16	0.001	14	0.001
5	One lag	14	0.000	15	0.000
6	Two lags	13	0.005	13	0.001
7	Three lags	12	0.003	12	0.000
	Standard error scaling				
8	Predisposition and one lag	4	0.000	4	0.000
9	Predisposition and two lags	5	0.000	5	0.000
10	Predisposition and three lags	7	0.000	7	0.000
11	One lag	1	0.000	1	0.000
12	Two lags	3	0.000	3	0.000
13	Three lags	6	0.000	6	0.000
	Range scaling				
14	Predisposition and one lag	9	0.002	9	0.000
15	Predisposition and two lags	15	0.001	16	0.000
16	Predisposition and three lags	19	0.000	19	0.000
17	One lag	8	0.001	8	0.000
18	Two lags	10	0.005	10	0.002
19	Three lags	11	0.002	11	0.000
	Maximum absolute value scaling				
20	Predisposition and one lag	MCS	0.243	MCS	0.116
21	Predisposition and two lags	MCS	0.243	MCS	0.116
22	Predisposition and three lags	MCS	0.243	MCS	0.116
23	One lag	MCS	0.243	MCS	0.116
24	Two lags	MCS	0.243	MCS	0.116
25	Three lags	MCS	0.243	MCS	0.116
	Relative performance scaling				
26	Predisposition and one lag	25	0.000	24	0.000
27	Predisposition and two lags	24	0.000	23	0.000
28	Predisposition and three lags	22	0.000	20	0.000
29	One lag	23	0.000	25	0.000
30	Two lags	21	0.000	22	0.000
31	Three lags	20	0.000	21	0.000

MSC denotes the final model confidence set with 95% level of significance

by the value of the log-likelikood function, denoted by LL. For each model, a confidence interval for relative performance is constructed with the use of the bootstrap method, which means that the values of the log-likelihood function are randomly subsampled. Iterative search for the worst model in terms of relative differences in LL value is then performed. For the worst model, the test for significant differences with other models in terms of performance is run. If the hypothesis of equality in LL can be rejected at an initially fixed confidence level, then the worst model is excluded and the process continues with the new, reduced initial set. The process stops once the worst model can not be excluded, and the current set of models is selected as the model confidence set, which contains the true model with a fixed level of confidence.

In this chapter, I follow the bootstrap implementation with a fixed block length, which is used for subsampling. I generate $B = 5000$ different resamples of the original data set for different values of the blocks' length parameter b, with $b = \{2, 4, 8\}$. I restrict the attention to the case of $b = 4$ and consider the results obtained for the remaining values as the robustness check.[31] As a measure of performance, I use both the Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC), which penalise increasing the number of parameters in the model.

3.5.4 Results

The results are presented in Table 3.2. For each of the 31 models, which combine scaling with performance measures, order of exclusion from the initial set and corresponding probability of false rejection are reported. These values should not necessarily coincide for different information criteria. Results suggest that the final set of models, which contain true model with 5% confidence level, consist only of different modifications of performance measures with absolute value scaling. Two conclusions can be made based on this observation. The first result relates to the fact that only the family of models with absolute value scaling performs substantially better than

[31] Results are confirmed for other parameter values which I attribute to the persistence in relative performance of the models.

Table 3.3: Estimates of the logit parameter in the Standard model and Self-Tuning model with maximum absolute value factor scaling.

	Standard Model		Self-tuning Model	
Block	β	***S.E.***	β	***S.E.***
WN	0.19***	(0.03)	0.87***	(0.15)
	0.30***	(0.04)	1.37***	(0.18)
BH	0.41***	(0.04)	2.33***	(0.21)
	0.48***	(0.04)	2.90***	(0.24)
SI	3.33***	(0.27)	3.19***	(0.26)
	4.04***	(0.32)	3.83***	(0.31)
SP	3.57***	(0.95)	1.39***	(0.21)
	4.91***	(0.93)	1.31***	(0.23)
	11.00***	(1.46)	1.44***	(0.19)
	9.90***	(1.50)	1.93***	(0.28)
	2.06***	(0.66)	1.53***	(0.24)
	8.28***	(1.40)	0.93***	(0.18)
	4.53***	(1.03)	0.92***	(0.21)
	8.59***	(1.40)	1.77***	(0.26)
UP	0.09***	(0.02)	1.54***	(0.22)
	0.13***	(0.02)	1.72***	(0.23)
	0.11***	(0.02)	1.86***	(0.25)
	0.17***	(0.02)	1.90***	(0.24)
	0.10***	(0.02)	1.05***	(0.20)
	0.14***	(0.02)	1.23***	(0.21)
	0.14***	(0.02)	1.58***	(0.21)
	0.17***	(0.02)	1.86***	(0.21)
UPL	0.15***	(0.01)	1.04***	(0.09)
	0.17***	(0.01)	1.28***	(0.09)
Pooled	0.17***	(0.01)	1.53***	(0.09)

Standard errors are reported in the column S.E.

*, **, *** indicates significance at the 90%, 95% and 99% level, respectively.

all other models. This result suggests that the process of tuning of the logit parameter to the environment, which is introduced in this chapter, is a statistically significant driver of subjects' adaptive behaviour.

The second important finding is that no model with an absolute value scaling can be excluded from the model confidence set. This finding suggests

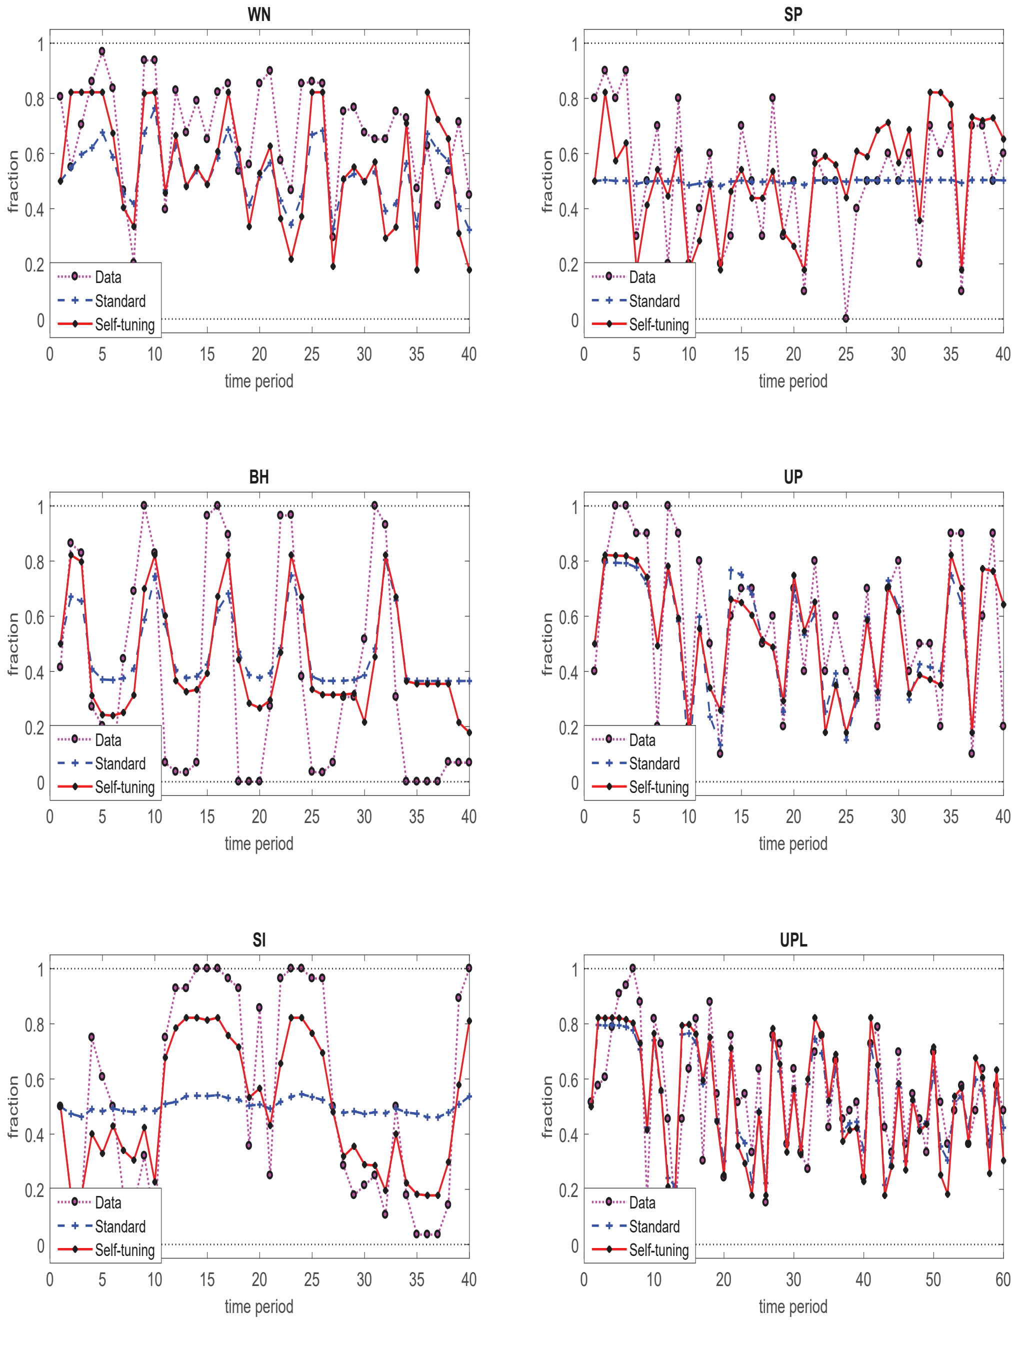

Figure 3.5: Fraction of subjects choosing B, denoted as Data, and two model predictions: the standard logit model, denoted as Standard, and the self-tuning model with absolute value scaling, denoted as Self-tuning. Both models are parametrised with value of the logit parameter which was obtained as best-fit for the pooled data: $\beta = 0.17$ for the standard model and $\beta = 1.53$ for the self-tuning model.

using additional steps to determine a preferable model of adaptive choice. I use uniformity of the estimated values across different sessions as an additional eligibility criteria. With this extra step, the simple discrete choice model with one lag of payoffs as performance measures and absolute value scaling is favoured (model number 23 in Table 3.2). Estimated values of parameters for this model are presented in Table 3.3.[32] Obtained values of the logit parameter are close to each other and often lie in confidence intervals of estimates from other sessions. Stability of estimates across sessions can be compared to, for example, estimates of the model discussed in Section , which are reported in Table B.1. Self-tuning model values for the logit parameter vary from 0.87 to 3.83 and are centred around the estimate of 1.53 obtained for the pooled data, and reported in the last line of Table 3.3.

3.5.5 Robustness

To illustrate the predictive power of the model, I provide a comparison between experimental data and the model predictions on choices in each of the six sessions presented as an example in Section 3.3. To stress the importance of normalisation, I also add predictions of the simple discrete choice model without normalisation. The corresponding graphs are presented in Figure 3.5. It can be observed that the model with absolute value scaling predicts experimental data the best and performs substantially better. This is especially true in SI and SP treatments, as the scaling is able to capture such features of the data as small magnitude in differences in payoffs.[33]

[32] Estimations for five other models are available in Appendix B.3.

[33] This is also captured by close values of RMSE of the simple model and model with scaling in WN (0.21 vs 0.22), BH (0.32 vs 0.3), UP (0.17 vs 0.17) and UPL (0.13 vs 0.13). But in the SI (0.34 vs 0.19) and SP (0.22 vs 0.16) RMSE is much larger for the model without normalisation.

Table 3.4: Values of RMSE for predictions of a fraction of subjects choosing B. Specifications of the self-tuning model differ with respect to the memory length k, which is used for scaling.

	Length of scaling factor memory k						
Block	7	8	9	10	11	12	13
WN	0.22	0.22	0.22	0.22	0.22	0.22	0.22
BH	0.29	0.29	0.29	0.30	0.30	0.30	0.30
SI	0.17	0.18	0.18	0.19	0.19	0.19	0.19
SP	0.16	0.16	0.16	0.16	0.16	0.16	0.16
UP	0.17	0.17	0.16	0.17	0.17	0.17	0.17
UPL	0.14	0.14	0.14	0.14	0.14	0.14	0.14

As an additional robustness check, consider the importance of the memory length, which reflects how many lags of payoffs are used to construct scaling factors. For estimation and model contest, the length of the "recent history" was fixed as 10 periods. Consider the explanatory power of the model for different values of $k = 7, .., 13$. Table 3.4 contains RMSE values for alternative specifications[34]. It can be observed that the model has almost the same explanatory power for the values which are close to 10, and falls as k becomes more distant from 10. This observation suggests that results are not an artefact of particular memory length parameter fixing.

3.6 Applications

To understand the policy implications of the scaling in the self-tuning model, the effects of incorporating the new model of adaptive choice into existing studies are investigated. I refer to the logit parameter as the IoC parameter to be consistent with the HAM literature, which was discussed in Section

[34]Root Mean Square Error is computed for one period ahead predictions for the fractions of choices for each model.

3.2. I start by recalling the seminal model[35] from [17], which utilises a simple model of adaptive choice without scaling. In that model, the IoC parameter is constant and the sensitivity of the dynamics, with respect to variation in the value of the IoC, is studied. The paper demonstrates that the IoC parameter plays a pivotal role in the dynamics, and the importance of particular values of IoC for ensuring convergence is established. Several papers that study policy implications of price stabilising interventions, for example, [111] and [113], were built upon the framework of the Brock-Hommes model. To illustrate the importance of the self-tuning model for the dynamics, the Brock-Hommes model will be augmented with the absolute value scaling.

3.6.1 The Brock-Hommes Model

Consider a "cobweb" model of economy, where a continuum of firms operates, each optimising production volume given its price expectations. Two prediction rules are available to form expectations over the next period price: rational expectations (or fundamental), which provide costly perfect foresight, and free naive prediction, which uses the current price as a best predictor for the future price. The cost of rational prediction is denoted by $C > 0$. The demand for goods produced and supplied to the markets is $A - Bp_t$, where p_t is the market clearing price of the good at time t. The supply curve is derived from the quadratic cost function $c(q_t) = q_t^2/2b$, where q_t denotes the production volume of a firm. Firms are assumed to adaptively switch between prediction rules–rational and naive–based on their performances. These monetary performances include the cost and the accuracy

[35]This paper effectively laid the foundations for HSM literature, utilising simple model to generate complicated dynamics. Therefore, it is an ideal candidate for micro-based modifications.

of price prediction, which is measured as the squared forecasting error of the previous period prediction. Fraction of firms, which uses naive prediction, f_t is then defined by the logit choice model, analogously to previously introduced equation (3.2):

$$f_{n,t} = \frac{\exp(\beta \pi_{n,t-1})}{\exp(\beta \pi_{r,t-1}) + \exp(\beta \pi_{n,t-1})}, \tag{3.10}$$

where $\pi_{n,t}$ and $\pi_{r,t}$ are performance of naive and rational prediction rules respectively, which are calculated as squared forecasting error minus cost, and β is the IoC parameter, which is assumed to be constant.

Combining this price predictor choice rule with optimal production decisions and market clearing conditions results in an Adaptive Rational Equilibrium Dynamics (ARED) (see Appendix B.4 for details of the derivation). The ARED is described by the following system of non-linear equations:

$$\begin{cases} p_t = \dfrac{A - p_{t-1} b f_{n,t-1}}{B + b(1 - f_{n,t-1})}, \\ f_{n,t} = \dfrac{1}{1 + \exp\left(\beta\left(\frac{b}{2}\left(\frac{A - p_{t-1}(B+b)}{B + b(1 - f_{n,t-1})}\right)^2 - C\right)\right)}, \end{cases} \tag{3.11}$$

Complicated dynamics for high values of the IoC parameter may arise in this model due to increasing attractiveness of free naive price prediction when the price converges towards its fundamental value. Prediction errors become smaller and the naive predictor attracts “free riding” firms, which prefer to save on costly rational prediction. This drives the price away from the fundamental value, which in the end forces firms to switch back to the rational prediction, as errors of naive predictors are too large. That stabilises the price and reinforces the story to repeat itself.

The IoC parameter plays a crucial role in generating these complicated dynamics, and the original [17] paper shows that, as the IoC parameter value increases, the dynamics of the model change from convergence to the stable rational expectation equilibrium, to cycles of different periods, to topological chaos. Intuitively, complicated dynamics arise in this framework due to non-linearity of the dynamical system, where initially small deviations can lead to qualitatively different predictions over a long horizon.

3.6.2 Simulations

Note that the model (3.11) describes non-linear dynamics that substantially complicate its stability analysis. In this section, two instruments to study the dynamics–which are commonly used in the literature–are applied: simulations and bifurcation diagrams. Once the model (3.11) is augmented with scaling (3.8), the following system of non-linear equations defines dynamics:

$$\begin{cases} p_t = \dfrac{-p_{t-1} b f_{n,t-1}}{B + b(1 - f_{n,t-1})}, \\ f_{n,t} = \dfrac{1}{1 + \exp\left(\frac{\beta}{F_{t-1}}\left(\frac{b}{2}\left(\frac{A-(B+b)p_{t-1}}{B+b(1-f_{n,t-1})}\right)^2 - C\right)\right)}, \\ F_{t-1} = \max\limits_{1 \leq l \leq 10} [|(\frac{b}{2}(\frac{A-(B+b)p_{t-l}}{B+b(1-f_{n,t-l})})^2 - C)|]. \end{cases} \tag{3.12}$$

The IoC parameter, which was constant in (3.11) is now scaled by the factor F_t. Equation (3.12) defines F_t as a function of two state variables: previously realised fractions $f_{n,t}$ and prices p_t.

By means of simulations, I address two major questions regarding the robustness of the original model (3.11):

1. How does the introduction of the new self-tuning switching mechanism, with the IoC parameter adjusting over time, affect model dynamics?

2. Is the augmented model robust to different IoC parametrisations of the switching mechanism?

I address these questions by comparing simulated dynamics in the original and modified models. Throughout the analysis, I fix the parameters of the model in the following way: $A = 4$, $B = 0.5$, $b = 1.35$, and $C = 1$. Similar parametrisation was used for illustrative numerical examples in [17]. For three different values of the IoC parameter β, the original model predicts qualitatively different price dynamics: stable fundamental zero-price for $\beta = 0.75$, two-period cycle for $\beta = 3.8$, and unstable for $\beta = 5$. Indeed, the left panels of Figure 3.6 replicate predictions of the model (3.11). The right panels of Figure 3.6 present dynamics of the modified model (3.12) with self-tuning switching mechanism for the same value of β. It can be observed that the modified model qualitatively replicates predictions while making the dynamics more noisy, for instance, by breaking two-period cycles in the middle row panels.[36]

I further investigate if this result of similarity between the models' dynamics is valid for a wider range of parameter values. I build a bifurcation diagram that illustrates changes in long-term dynamics with respect to changes in the IoC parameter. Bifurcation is a qualitative change in the dynamics and will be reflected on the diagram with transitions, for example, from the steady state to the two-period cycle. Consider bifurcation

[36] More simulations and bifurcation diagrams available in Appendix B.6 for different parametrisations.

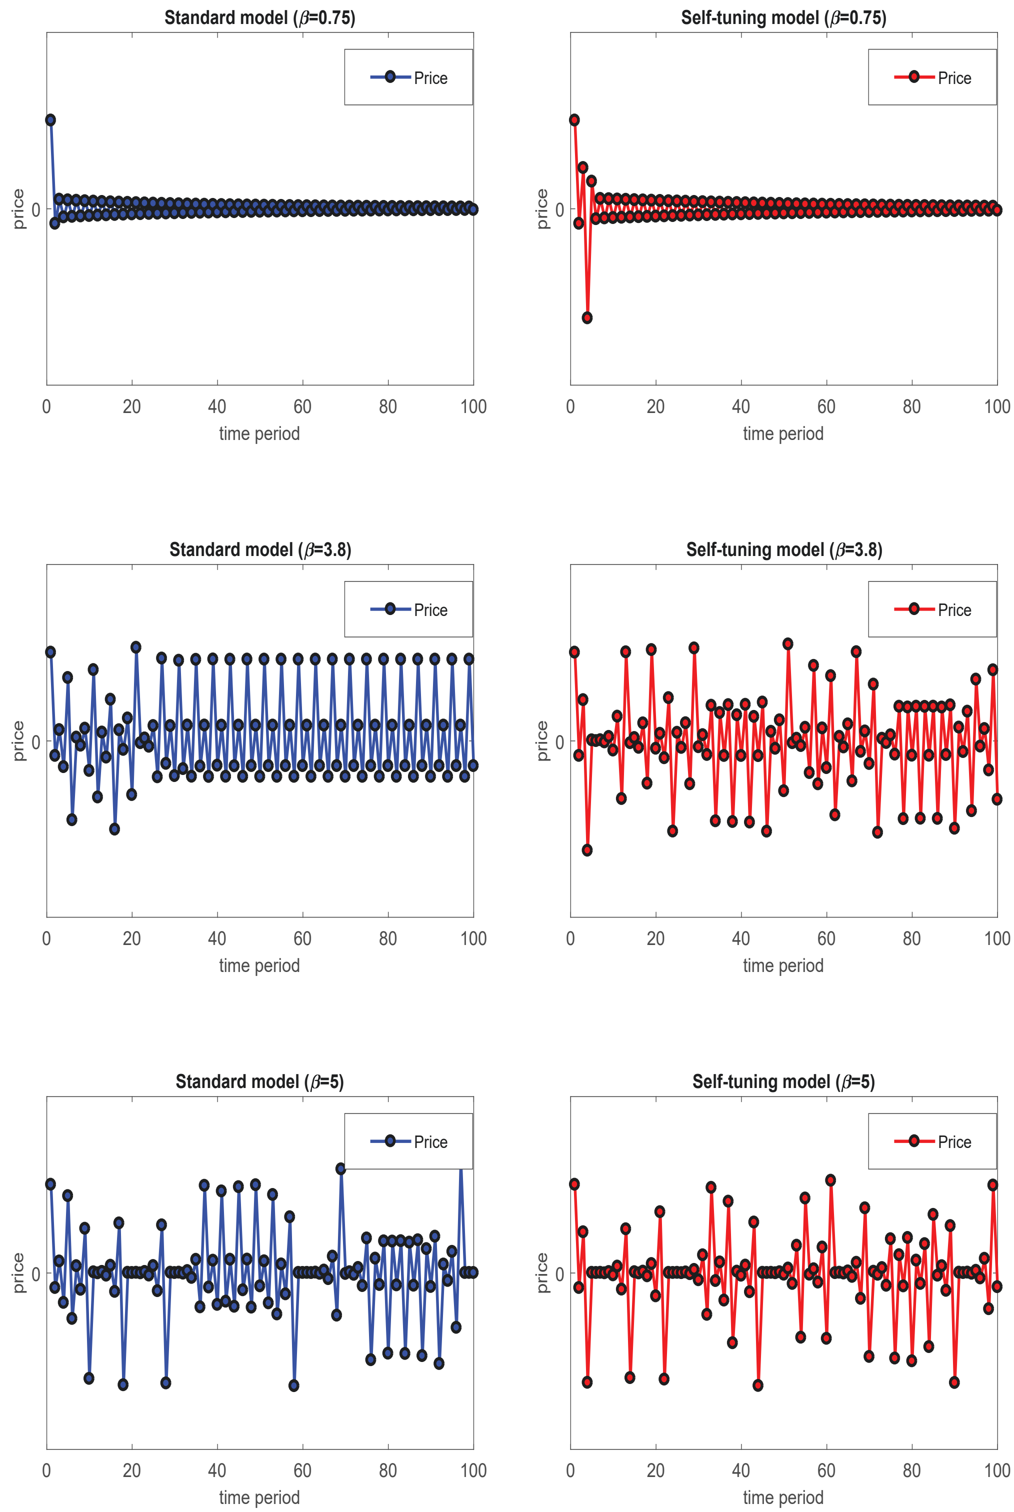

Figure 3.6: Simulations of the price dynamics for different values of the logit parameter. Model parameters are fixed to $A = 0$, $B = 0.5$, $b = 1.35$ and $C = 1$. *Top row*: $\beta = 0.75$. *Middle row*: $\beta = 3.8$. *Bottom row*: $\beta = 5$. *Left panels*: Dynamics of the standard model. *Right panels*: Dynamics of the self-tuning model.

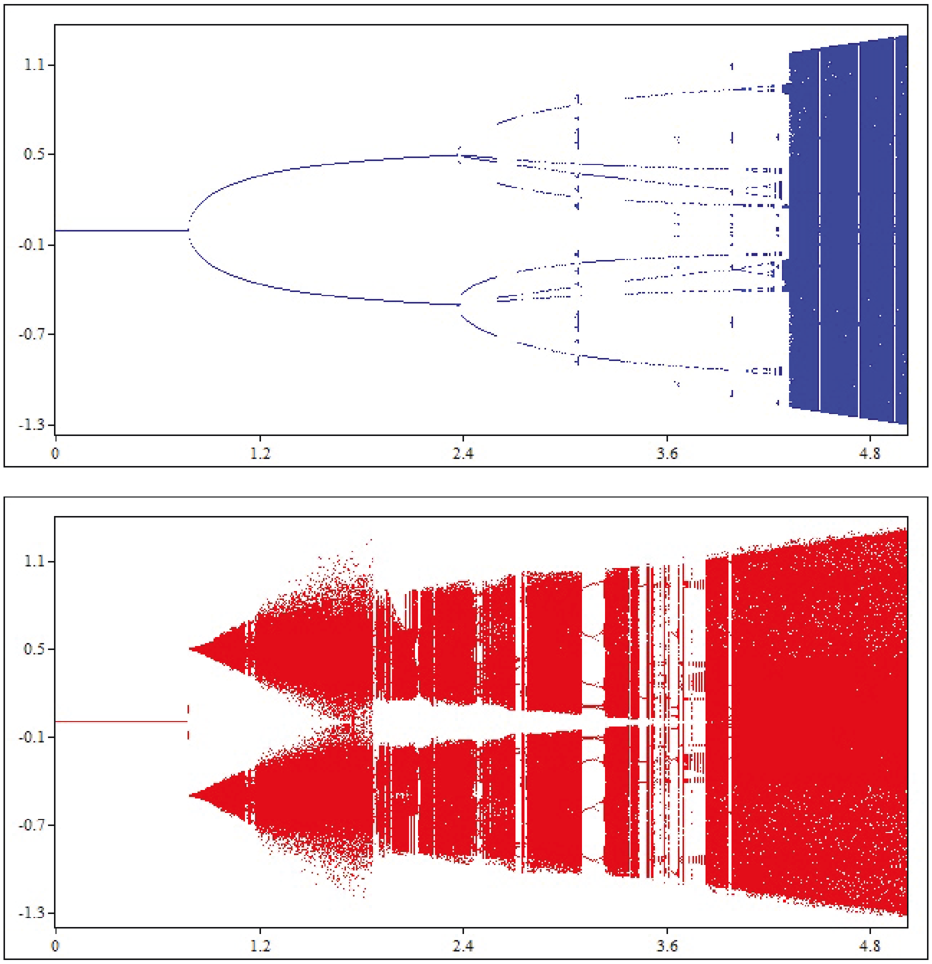

Figure 3.7: Bifurcation diagrams for the long-run behaviour of prices as a function of IoC parameter. *Top panel*: standard model. *Bottom panel*: self-tuning model. Other parameters are fixed at $A = 0$, $B = 0.5$, $b = 1.35$ and $C = 1$.

diagram[37] for the price dynamics of the original model (3.11) as a function of the IoC parameter depicted on the upper panel of Figure 3.7. For each IoC parameter value on the x-axis, 2000 price points–defined by the system (3.11)–are depicted on the y-axis, while the first 2000 periods are dropped to ensure that long-term convergence is captured on the diagram. The following results can be observed: stability of the steady state for low values of the IoC parameter up to 1, 2-period cycle for larger IoC values, and, in the end, unstable dynamics for values larger than 4 (which can be theoretically

[37]All bifurcation diagrams were generated in E&F Chaos program for non-linear systems analysis, see [36] for the software details.

proven to be deterministic chaos). We can compare these results with the bifurcation diagram for the augmented model (3.12), depicted in the lower panel of Figure 3.7. For the augmented model, we observe the same point of bifurcation at IoC equal to 1, after which the dynamics lose stability with a transition straight to unstable dynamics and not to a 2-period cycle.[38]

The numerical analysis suggests that the introduction of the new self-tuning mechanism eliminates a large part of the original model predictions in parts of the cycles with different lengths. In the modified model, only two regimes of the dynamics can be observed: steady state convergence and instability. Moreover, for the logit parameter estimate of 1.53, which was based on the pool data, the modified model predicts non-convergence. In a broad sense, it is consistent with the results of the ACT18 experiment, where different parametrisations of the stylised version of the Brock-Hommes model fluctuations of the price were observed.

3.7 Conclusion

This chapter addresses the issue of heterogeneity of the logit parameter, which is a behavioural parameter reflecting intensity or subjects' attentiveness. In different strands of the literature, the logit parameter in models of adaptive choice plays a crucial role and generates qualitatively different predictions for different values of the parameter. The new self-tuning model of adaptive choice endogenises identified differences in the logit parameter estimates by using a scaling mechanism. In this model, agents make their

[38]Indeed, these chaotic dynamics can be traced by computing the Lyapunov exponent, which is available in Appendix B.6. Positive values of the Lyapunov exponent, which indicates chaos, can be observed for the values of the IoC parameter larger than 1.

choice based on the observed payoffs, but they scale their attention if the current stakes are relatively small in comparison to the previously observed magnitude of fluctuations. The self-tuning model of choice performs well in explaining experimental data from different studies, and the MCS approach supports robustness of this result in the model contest with other tuning mechanisms.

The estimates of the logit parameter in the self-tuning model do not exhibit high heterogeneity, which solves the calibration issue. I incorporate the self-tuning model into the setup of the seminal Brock-Hommes model to investigate the effects of the logit parameter scaling on the price dynamics. In the Brock-Hommes model, large values of the logit parameter generate instability of the price dynamics. Computational analysis suggests that the self-tuning model replicates existing results and generates transition from the steady state to instability, but moves away knife-edge prediction regarding the cycles of different length. The new self-tuning model of agents' adaptive choice could potentially enrich understanding of the dynamics after the price stabilisation policies, introduced in a number of studies based on the Brock-Hommes model framework, and presented in [95].

There are several more directions of potential application for the self-tuning model. First, incorporating the new model of adaptive choice into existing empirical studies could improve the fitness of the learning and agent-based model. Previously, these models did not allow the logit parameter to adjust over time; the scaling mechanism of adjustment can bring flexibility to the model. Second, the self-tuning model can be used to introduce a dynamical aspect to the static concepts that fix the logit parameter at an exogenously given value. The scaling mechanism specifies the evolution of the logit parameter, which can be applied to the logit quantal response equilibrium.

Chapter 4

Price volatility and forecasting horizons: An experimental investigation

4.1 Introduction

Phenomena of the excess volatility, which is one of the most renowned stylized facts of financial markets, remains an open issue of finance literature.[1] Models of excessive instability of price dynamics, which can not be explained by fluctuation of fundamental values, are presented in both theoretical and empirical studies. We contribute to the literature by running a controlled laboratory experiment, which is designed to study the effects of increasing the forecasting horizon on volatility of the market price. We demonstrate that an increase in the forecasting horizon may have different effects de-

[1]One of the earliest contributions was made by the Nobel Prize laureate Robert J. Shiller, see [97].

pending on the stability of price history. In stable markets, increasing the forecasting horizon does not affect the stability of the price discovery process. In unstable markets, the following effect is observed: an increase in the forecasting horizon destabilizes the price and increases volatility, but the effect is diminishing. The latter finding might be attributed to the lack of coordination on non-fundamental expectations.

The Efficient Market Hypothesis (EMH) is a cornerstone of modern financial economics and stipulates that the market price effectively aggregates available information and reflects the value of the traded good. Nonetheless, various stylized facts of financial markets, including excess volatility, fat tails, or volatility clustering, suggest that the EMH may not hold, whereas a number of behavioural models successfully explained different market phenomena. In fact, the currently observed prosperity of behavioural economics, and in particular behavioural finance, illustrate the potential of this approach.[2] One of the crucial assumptions that is used to derive the EMH is the "rationality" of the market participants' expectations. We utilize a behavioural approach to modelling expectations and examine the effects of changes in the forecasting horizon on the volatility in the asset markets populated by agents with boundedly rational expectations. In the absence of perfect foresight, agents, induced by rational expectation, have only a limited understanding about future price movements. With the increase in the horizon length the precision of forecasts falls, which could potentially lead to an increase in price volatility and long periods of an asset mispricing. At the same time, theoretically, an increase in the forecasting horizon may smooth short-term deviations of the price from the fundamentals due to a

[2] Awarding Richard H. Thaler with the Nobel Prize in 2017 for his contributions to behavioural economics stresses the current importance of behavioural economic models.

weakened short-term speculative motive.[3] The relative strength of these two divergent effects determines the price (in)stability. Dedicated experimental investigation is required since intensities of both effects are unknown.[4]

Two described effects, which affect market price (in)stability, rely on important features of real markets: mutual feedback effects. Expectations affect current prices through corresponding trading decisions, and realized prices affect the expectation through corresponding observational learning and adaptation. In the case of speculative asset markets, this positive feedback can be responsible for driving prices away from the fundamental values. If investors expect prices to grow in the future, they increase current demand for the asset by speculative buying. This increase in demand today drives prices up, therefore reinforcing expectations of further price increases in the future. On the contrary, if the market is stable and price expectations coincide with price observations, trader behaviour does not destabilize the market, therefore securing its stability. Consequently, the importance of mutual feedback effects depends on observed price fluctuations. This observation motivates two separate experimental studies of the forecasting horizons length in two different environments: unstable, where current mispricing reinforces fluctuations, and stable, where price dynamics do not diverge from the fundamental value.

[3] Intuitively, it follows that the speculative part of an excess return is smaller in the case of a longer investing period, which relates to a larger forecasting horizon. More details are given in Section 4.2, where a formal asset-pricing model is provided.

[4] Currently, only an indirect comparison of effects is available based on data from previously run experiments. Possible positive relation between the length of the forecasting horizon and price volatility is suggested by the results of two experiments: [68] and [58]. Experimental markets in these papers differ only in the dimension of the forecasting horizon: one-period ahead prediction markets versus two-periods ahead prediction markets. Significantly higher volatility is observed in markets with longer horizons (see Table C.11 for numerical values).

Laboratory experiments on expectation formation have proven to be a powerful instrument to study forecasting strategies (see [93], [39], [60]). At the same time, these studies on individual behaviour were run in the absence of the mutual feedback effects–forecasting was performed based on the pregenerated price sequences. Alternative studies, which explicitly introduced feedback effects in group experiments on expectations formation are referred to as Learning-to-Forecast (LtF) experiments (see [65] for a summary of existing experimental results). In these experiments, participants repetitively provide forecasts for the price, and their predictions are used to compute an optimal trading decision for the firm that they are "advising". Two types of experiments were run: (i) negative feedback markets, where the underlying model is "supply-driven" by the decisions of producers on the market, and (ii) positive feedback markets, where the speculative asset market model is used. The LtF asset-pricing experiments with positive feedback performed in the laboratory have been quite successful in replicating bubbles and crashes–prolonged deviations of the market price from the fundamental value. Participants in the experimental markets coordinated their expectations on non-fundamental predictions (see [68], [71], and [58] for details). We study the effects of the forecasting horizon length by extending the framework of the LtF asset-pricing experiments.[5] To do this, we modify the model to allow for investors with different forecasting horizons, and formulate testable predictions based on simulations with previously estimated behavioural forecasting rules. We program experimental markets, which differ in forecasting horizon lengths by one- , two-, or three-periods ahead. Each market has six participants, whose only task is to predict the

[5]Experiments with forecasting horizons of different lengths were previously run (see [33]), but these experiments lack important feedback effects. Long-run forecasts were elicited but they did not play a role in the dynamics.

price for 50 consecutive periods. To induce stable and unstable markets, we show price sequences from previous experiments, which differ in their level of stability. We fix the initial price history that is available to participants to be stable in one group of treatments, and unstable in the other. Resulting price series from different experimental markets suggest that initially stable markets irrespective of forecasting horizon length do not exhibit volatility, whilst initially unstable markets demonstrate bubbles and crashes patterns less often as the forecasting horizon increases.

The chapter is organized as follows. Section 4.2 introduces the asset-pricing model that is used to run experimental markets. Section 4.3 provides numerical simulations of the market dynamics and details of parametrisation. Section 4.4 presents the experimental design and hypotheses.. Section 4.5 contains results and discussion. Section 4.6 concludes.

4.2 Model

Existing studies of LtF experiments with positive feedback use a standard asset-pricing model, which is referred to as a Present Discounted Value (PDV) model in [18]. In this model, a higher expected future price increases a speculative demand for this asset, which increases the current price. This relationship between expectations and the price reproduces a main feature of a positive feedback system. This model can generate bubbles and crashes if a market is populated by boundedly rational agents with potentially heterogeneous expectations. Our main interest is whether a larger forecasting horizon can bring volatility to the price sequence.

Consider a market that is populated by many investors, such as pension funds. These institutional investors maximize the wealth by investing in two available assets: risky and risk-free. The risk-free asset guarantees a return of r per period. The risky asset pays stochastic dividend y_t at every period t. The difference in horizons translates to the period for which the final wealth is subject to maximization.

Let us denote the wealth of investor i in period t as $W_{t,i}$. Then an investor with horizon h will maximize $W_{t+x,i}$ whose evolution is given by

$$\begin{aligned} W_{t+h,i} &= (1-x_{t,i})W_{t,i}R^h + \frac{x_{t,i}W_{t,i}}{p_t}\left(p_{t+h} + \sum_{s=1}^{h} R^{h-s}y_{t+s}\right) = \\ &= W_{t,i}R^h + \left(p_{t+h} + \sum_{s=1}^{h} R^{h-s}y_{t+s} - R^h p_t\right) z_{t,i}\,, \end{aligned} \tag{4.1}$$

where $x_{t,i}$ is the share of wealth invested to the risky asset so that $z_{t,i} = x_{t,i}W_{t,i}/p_t$ are the *holdings* of the risky asset bought at time t, and $R = 1+r$. In expression (4.1), we assume that all the dividends are automatically reinvested into the riskless asset.

Assume that all investors are myopic mean-variance maximizers with risk aversion parameter a and shared beliefs about the conditional variance of excess returns to be equal to σ_h^2. The optimal amount of the risky asset to be purchased by any investor with horizon h is calculated, given expectations over the returns as follows:

$$z^*_{h,t,i} = \frac{\mathbb{E}_{h,t,i}\left(p_{t+h} + \sum_{s=1}^{h} R^{h-s}y_{t+s} - R^h p_t\right)}{a\sigma_h^2}\,,$$

where we keep individual index i to account for possibly heterogeneous expectations of investors with the same horizon. Let N_h be the number of investors with horizon h. The total demand for the risky asset is a sum of individual demands by all investors in the market:

$$z_{D,t} = \sum_h \frac{1}{a\sigma_h^2} \sum_i \left(\mathbb{E}_{h,t,i}[p_{t+h}] + \sum_{s=1}^{h} R^{h-s} \mathbb{E}_{h,t,i}[y_{t+s}] - R^h p_t \right) . \qquad (4.2)$$

Assume that supply of the risky asset is exogenously fixed to the constant value of z_S. We can set $z_S = 0$ so that the market clearing condition $z_{D,t} = z_S$ reads as

$$\sum_h \frac{1}{a\sigma_h^2} \left(\sum_i \mathbb{E}_{h,t,i}[p_{t+h}] + N_h \sum_{s=1}^{h} R^{h-s} \mathbb{E}_{h,t,i}[y_{t+s}] - N_h R^h p_t \right) = 0 ,$$

where N_h stands for the number of investors with horizon h. Dividing this equation by the total number of investors N and denoting the "adjusted" fraction of investors with horizon h, $f_h = N_h/(a\sigma_h^2 N)$, we obtain

$$\frac{1}{N} \sum_h \frac{1}{a\sigma_h^2} \sum_i \mathbb{E}_{h,t,i}[p_{t+h}] + \sum_h f_h \sum_{s=1}^{h} R^{h-s} \mathbb{E}_{h,t,i}[y_{t+s}] - \sum_h f_h R^h p_t = 0 .$$

We then can derive the price dynamics equation:

$$p_t = \frac{1}{\sum_h f_h R^h} \left(\sum_h f_h \frac{\sum_i \mathbb{E}_{h,t,i}[p_{t+h}]}{N_h} + \sum_h f_h \sum_{s=1}^{h} R^{h-s} \mathbb{E}_{h,t,i}[y_{t+s}] \right) . \qquad (4.3)$$

Independently and Identically Distributed (IID) Dividends. Let us now assume that the dividend process is IID with the mean value $\bar{y}$ and that this is known to all investors. Then we can calculate the stream of expected dividend payments as follows:

$$\sum_{s=1}^{h} R^{h-s} \mathbb{E}_{h,t,i}[y_{t+s}] = \bar{y} \, \frac{R^h - 1}{R - 1} .$$

To derive the *fundamental price*, let us assume that all investors have rational expectations $\mathbb{E}_{h,t,i}[p_{t+h}] = p_{t+h}$. Then the price equation (4.3) becomes

$$p_t = \frac{1}{\sum_h f_h R^h} \left(\sum_h f_h p_{t+h} + \sum_h f_h \frac{\bar{y}}{r} (R^h - 1) \right).$$

Let us look for the constant price solution, $p_t = p_{t+h} = p^*$. Then we have

$$p^* = \frac{\bar{y}}{r} \frac{\sum_h f_h (R^h - 1)}{\sum_h f_h R^h} \Big/ \left(1 - \frac{\sum_h f_h}{\sum_h f_h R^h} \right) = \frac{\bar{y}}{r}.$$

Thus, the mean dividend divided by the risk-free rate is the only constant fundamental price independent of the horizon length.

If all investors in the market have the same horizon h, then the "adjusted" fraction boils down to $f_h \equiv 1/(a\sigma^2)$ and the pricing equation (4.3) simplifies to

$$p_t = \frac{1}{R^h} \left(\frac{1}{N} \sum_i \mathbb{E}_{t,i}[p_{t+h}] + \bar{y} \, \frac{R^h - 1}{R - 1} \right). \tag{4.4}$$

For the simplest case, when investing and forecasting horizons are equal to one, $h = 1$, we obtain the standard equation which governs price dynamics in the majority of LtF experiments:

$$p_t = \frac{1}{R} \left(\frac{1}{N} \sum_i \mathbb{E}_{t,i}[p_{t+1}] + \bar{y} \right).$$

As we increase the horizon of holding and forecasting to two periods, $h = 2$, we have the pricing equation as follows:

$$p_t = \frac{1}{R^2} \left(\frac{1}{N} \sum_i \mathbb{E}_{t,i}[p_{t+2}] + \bar{y}(1 + R) \right).$$

Finally, for three-periods ahead forecasting markets, $h = 3$, the pricing equation is defined as follows:

$$p_t = \frac{1}{R^3}\left(\frac{1}{N}\sum_i \mathbb{E}_{t,i}[p_{t+3}] + \bar{y}(1 + R + R^2)\right).$$

The following conclusion can be made based on comparison of these three equations: as the horizon length increases, the feedback strength of the forecasts falls. This observation relates to the previously given theoretical prediction regarding the role of horizon length–with an increase in length, the effect of potentially non-fundamental forecasts falls.

4.3 Simulations

We run simulations to gain insight to how the price behaves in the experimental markets defined by the equation (4.4), depending on the horizons length and forecasting rule. Our experimental design is a "2 x 3 design": treatments differ in initial price history (stable versus unstable) and forecasting horizon length (one-, two- or three-periods ahead prediction). We run a simulation for a particular parametrisation of market fundamentals (dividend and interest rate), and initiate dynamics on the price histories that are demonstrated to participants in experimental markets. Forecasting strategies are either textbook benchmarks or behavioural strategies estimated on individual data from previous LtF experiments.

Parametrisation. We fix parameters of the model at the following values: $R = 1.05$ and $\bar{y} = 3$. Following the LtF literature, we add a small noise

$\epsilon_t \sim N(0, 0.25)$ to the equation (4.4). This noise represents some small idiosyncratic shocks on the supply side of the market. We set number of investors to be equal to 6 in each market. Let us denote the average expected price by $\overline{p_{t+H}}^e$. Three different forecasting horizons' lengths, denoted by H, give three pricing equations:

$$\text{For } H = 1 : p_t \simeq 0.95 \left(\overline{p}^e_{t+1} + 3\right) + \epsilon_t,$$

$$\text{For } H = 2 : p_t \simeq 0.9 \left(\overline{p}^e_{t+2} + 6.15\right) + \epsilon_t,$$

$$\text{For } H = 3 : p_t \simeq 0.86 \left(\overline{p}^e_{t+3} + 9.46\right) + \epsilon_t.$$

It is easy to check that the fundamental price is the same for every horizon length and is equal to 60.[6] What may drive the differences between markets is difference in forecasting rules and associated average forecast $\overline{p_{t+H}}^e$, which determines the market price.

Initial price history. From the previous LtF experiments, it is known that the initial sequence of observed prices in the experiment plays an important role in the resulting price dynamics in the experimental market. Moreover, the recent Overlapping Generations (OLG) experiment in [10] utilizes a training phase on a two-period cycle, which induces the coordination of subjects' forecasts on observed "learned" cyclical pattern. We acknowl-

[6]Equations are presented in approximated forms. Discounting factor and dividend were rounded up to two decimal points for the sake of tractability. To solve for the fundamental price of 60, non-rounded values are required.

edge the importance of the initial conditions and fix initial prices to control the sequence of initially observed price histories. Initial prices help to avoid the spontaneous rise of irregular price patterns, which would inevitably create obstacles in the comparison of different experimental sessions. In order to initialize experimental markets with the same price series for the first 10 periods, we use the price sequences from two previous experimental sessions for the LtF experiments. In order to initiate markets with different price histories, we use the data from two experimental markets from the experiment studied in [68]. We utilize the data from the first 10 periods from the group 6 of this experiment for stable initial price markets, and the data from the group 7 for unstable markets. We simulate subsequent dynamics for these two initial histories using a different expectation formation process available for market participants. We consider the general form forecasting rules and the behavioural rules estimated on the data from previous experiments.

General form rules. We start computational analysis of the dynamics by considering archetypical prediction rules: rational, naive, and sophisticated. For the simulated dynamics in each treatment, we study the mean price and variance of the price on the market. Deviation of the mean price from the fundamental value of 60 indicates mispricing. Variance of price, in case it is higher than the variance of the noise added, indicates excess volatility of the price discovery process.

Denote the price expected for the period $t + H$ at the period t as p^e_{t+H}, where H is the forecasting horizon length. Prediction rules are then defined as follows:

- Rational expectation:

$$p^e_{t+H} = 60.$$

  This rule corresponds to the standard rational expectations prediction of the fundamental price value. Such behaviour of the market participants generates a price sequence, which fluctuates around 60 due only to the small noise added to the price generating process.

- Naive prediction:

$$p^e_{t+H} = p_{t-1}.$$

  Naive prediction, which uses the last observed price as a price forecast, converges to small fluctuations around the fundamental value of 60.

- Sample average learning:

$$p^e_{t+H} = \frac{1}{t-1}(\sum_{s=1}^{t-1} p_s).$$

  This sophisticated forecasting rule allocates equal weights for all the past observations of the price. Although this rule utilises more information and intuitively produces more efficient markets price, sometimes the dynamics converge to the fundamental value slower than with a very simple naive prediction.

There are two important observations that can be made with respect to the benchmark forecasting rules. First, price on markets, which are populated by agents with benchmark forecasting rules, quickly converges to the fundamental value and stays there. This prediction contradicts the experimental evidence from the previous LtF experiments. Second, all benchmark predicting rules generate forecasts, which do not depend on the forecasting

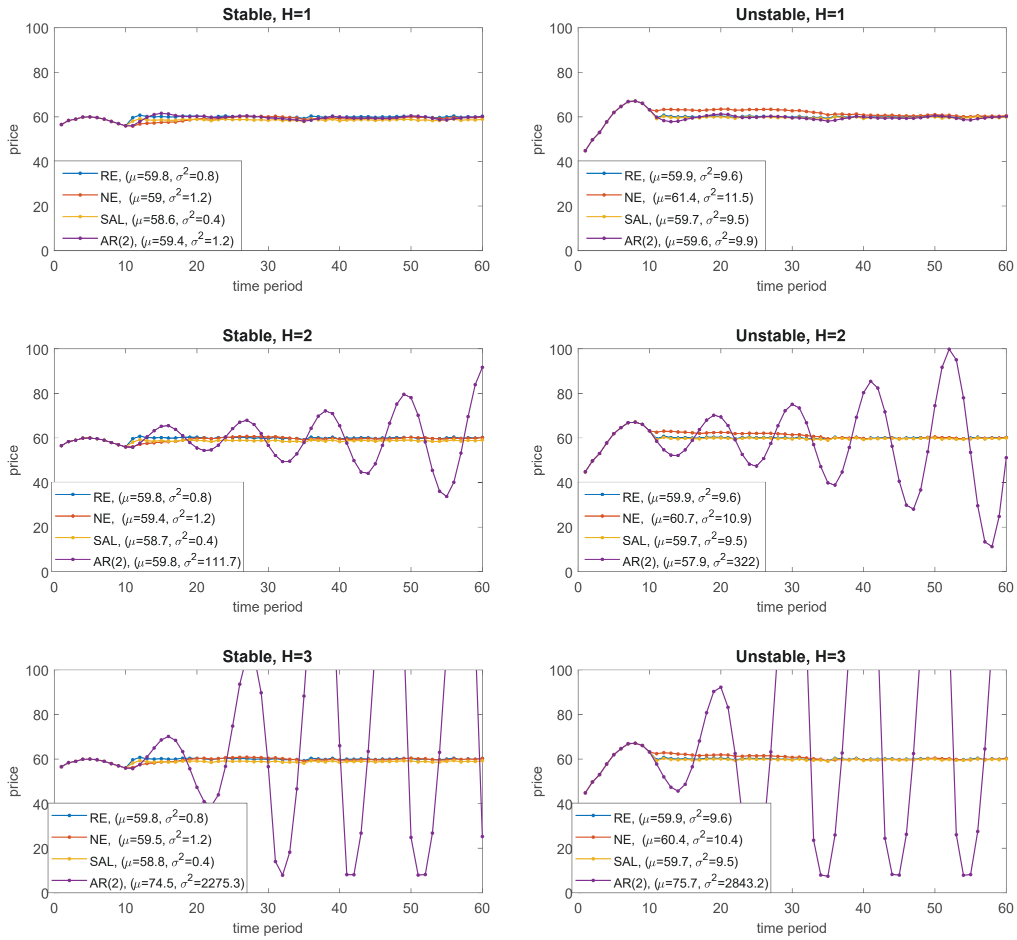

Figure 4.1: Simulated price dynamics for market with different combinations of the initial price history (stable and unstable) and different forecasting horizon length($H = 1, 2, 3$). Simulated market participants have one of the four forecasting rules: Rational Expectation (RE), Naive Expectation (NE), Sample Average Learning (SAL) and behavioural AR(2). Resulting prices are displayed in different colours. For each rule, the resulting mean and variance of the price after 50 periods of simulations are reported. *Left panels*: stable initial price history. *Right panels*: unstable initial price history. *Top row*: one-period ahead forecasting. *Middle row*: two-periods ahead forecasting. *Bottom row*: three-periods ahead forecasting.

horizon length. For these reasons, both the observed price fluctuations and the differences in time series of price between experimental markets with different horizons would suggest that participants have trend-extrapolative expectations.

AR(2) rules. We refer to the trend extrapolation predictions with two lags as AR(2) rules. This behavioural rule utilizes the information in the last two prices observed and has a number of parameters: the parameter of anchoring α, the weight of the last observation β, and the trend extrapolation parameter δ. Prediction of the price in H periods from period t is calculated as follows:

$$p^e_{t+H} = \alpha + \beta p_{t-1} + (\sum_{s=1}^{H} \delta^s)(p_{t-1} - p_{t-2}).$$

This forecasting rule can be interpreted as follows: the agent uses his anchor and the last observation as the current level of prices, but also extrapolates trends from recent prices and adds them to his expected value. The AR(2) rule explicitly depends on the forecasting horizon length H: the trend extrapolation component increases with the length of the forecasting horizon. For the simulation purposes, we parametrise AR(2) based on the estimations provided in [68]. We use the value of parameters estimated on the data from group 4 and fix AR(2) as follows: $\alpha = 15$, $\beta = 0.7$, and $\delta = 0.7$. Simulations are presented in Figure 4.1. Based on the observed price dynamics, we formulate several testable hypotheses:

Hypothesis 4. *There is no substantial difference in average prices between all experimental treatments.*

This hypothesis may be seen as a weak form of informational efficiency of the experimental markets. In our simulations, we observe either convergence to, or fluctuations around, the fundamental price.

Hypothesis 5. *There is a substantial difference in the volatility of price between markets with stable price history and unstable price history for each forecasting horizon length.*

This hypothesis relates to the importance of initial price histories. Simulations suggest that initial price fluctuations are reinforced in all experimental markets, and volatility is higher overall in initially unstable markets.

Hypothesis 6. *There is a more substantial increase in the volatility of price in markets with an unstable price history than for markets with a stable price history as the length of the forecasting horizon increases.*

This hypothesis directly relates to the role of the forecasting horizon's length. Theoretically, observed price fluctuations are reinforced as agents extrapolate the price trend more frequently when forecasting a price further in the future.

4.4 Experimental design

The experiment was conducted at the University of Technology Sydney's Behavioural Lab in May, August, and September 2018. A total number of 222 students, recruited in the ORSEE[7] system, participated in 6 main experimental treatments and 1 extra session, discussed in Section 4.5.4. No student participated in more than one session. Each session lasted approximately 90 minutes and subjects earned 30 Australian dollars (including the show-up fee) on average. The experiment was programmed and conducted using the z-Tree software.[8]

At each experimental treatment, the session participants were randomly divided in fixed groups of 6 people to operate in the same market. In the

[7] The Online Recruitment System for Economic Experiments (ORSEE) (see [54] for details).

[8] Zurich Toolbox for Readymade Economic Experiments (z-Tree) (see [46] for details).

Figure 4.2: An artificial example of the screen layout used in experiments. Participants can submit their prediction for the future price. Participants are shown information on previous period earnings and cumulative earnings. Past prices and predictions are available in two formats: a graph and a table.

instructions, participants were given the information on their task of acting as "professional forecasters", as well as qualitative and quantitative features of the market and the remuneration procedure.[9] At the beginning of the experiment, participants were shown the initial history of prices: an either stable (we denote these treatments with an "S") or unstable (we denote these treatments with a "U") sequence of the first 10 prices. Based on observed history, participants made their prediction in period 11 for the price one-period ahead (we add "1" to the treatment name), two-periods ahead, or three-periods ahead (reflected in the treatment name accordingly). This task of predicting the future price was performed for 50 consecutive periods.[10] Table 4.1 summarizes information on the number of experimental markets in each treatment.

[9]Instructions are available in the appendix.

[10]We count only predictions, which are elicited with monetary incentives, therefore the actual number of periods in the sessions slightly increases with an increase in horizon length.

Price history	**Horizon length**	**Market label**	**Number of markets**
Stable	H=1	S1	4
	H=2	S2	4
	H=3	S3	3
Unstable	H=1	U1	8
	H=2	U2	6
	H=3	U3	6
Unstable with strong feedback	H=3	SU3	6

Table 4.1: Information on a number of experimental markets for each treatment. Each treatment is a combination of the initial price history and the forecasting horizon length. Basic experimental design is 2 by 3, whilst one extra treatment with strong feedback was added to investigate coordination effects.

In all sessions, the price-generating process followed the dynamics defined by the equation (4.4). The participants collected points for the accuracy of their forecasting at each period t in accordance with the following rule:

$$E = \max\{1300 - \frac{1300}{49}(p_t - p_t^e)^2, 0\}. \tag{4.5}$$

This trimmed quadratic scoring rule is widely used in LtF experiments. The payoff function can be derived from the profit of the firm, which is advised by the “professional forecaster”. This function is trimmed to avoid inquiring about losses and possible bankruptcy of the subject during the experiment.

Points earned during the experiment were converted by the end of the experiment to Australian dollars, with the following exchange rate: 0.5 dollar for 1300 points. A participant could earn up to a half-dollar for each prediction, and remuneration linearly decreased with an increase in quadratic error of prediction, providing incentives to give accurate forecasts.

4.5 Results

In this section, we start by discussing the aggregate results of our experimental markets. We consider price dynamics and compare results in different treatments using various measures of price stability and efficiency. We further investigate the origins of identified differences with examination of individual forecasting behaviour. Results from an extra session are used to highlight the importance of coordination of expectations for generating price fluctuations.

4.5.1 Aggregate price dynamics

Aggregate price dynamics observed in different treatments overall are in line with preliminary simulation–price fluctuated around its fundamental value of 60.[11] Figure 4.3 shows price sequences in different markets in each of the 6 treatments. Based on observations of these graphs, several preliminary conclusions can be made.

[11] There are several notable outlying observations. For one of the groups in one-period ahead forecasting with stable price history, we can observe cyclic fluctuations of the price around non-fundamental value. This out of the box behaviour is explained by forecasting behaviour of one participant, who had been predicting the price to be around 5 throughout the duration of the experiment (see group 3 in Figure C.1). We could also observe two outliers in prices: in two-periods ahead forecasting with stable price history and in three-periods ahead forecasting with unstable price history. Price jumps in that markets were caused by the typos in submitted individual forecasts (see group 2 in Figure C.2 and group 2 in Figure C.6). These distortions could substantially affect quantitative comparison of the different treatments. Therefore, all discussions and tests will be based on the data, which do not include these 3 markets.

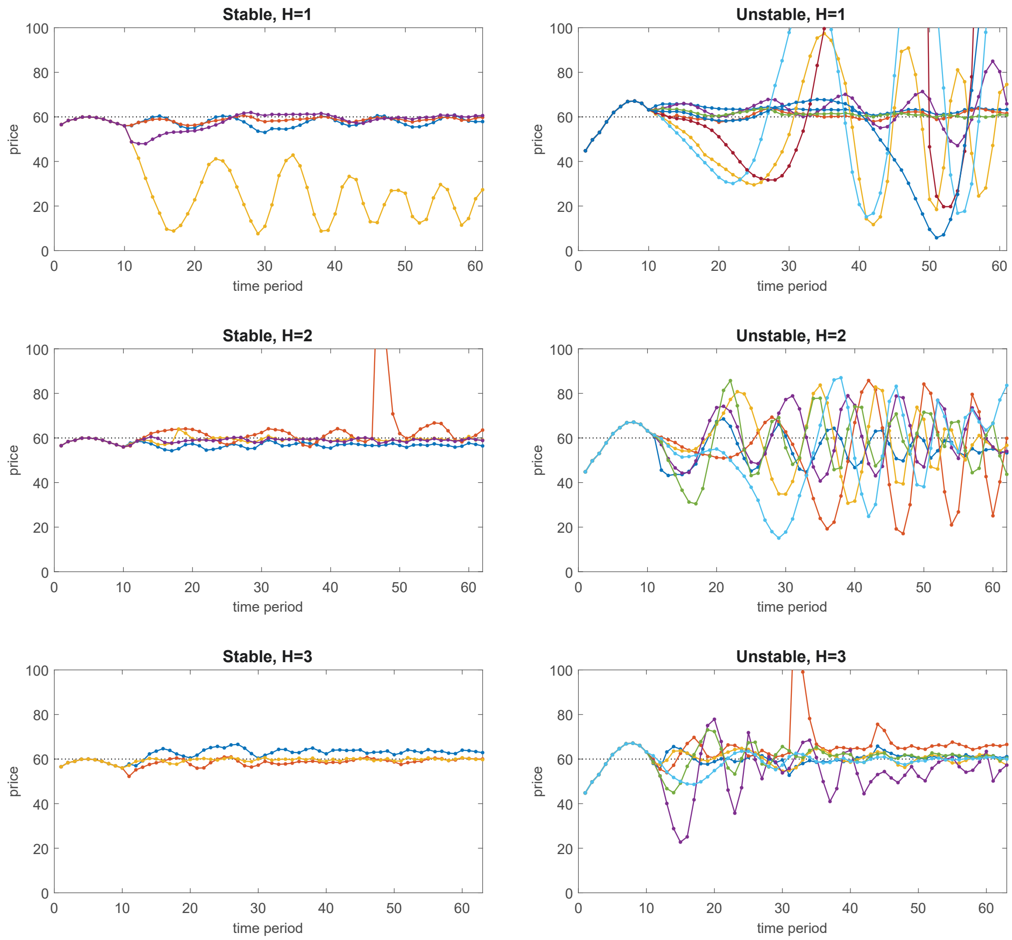

Figure 4.3: Price dynamics in main 6 experimental treatments. Different groups are captured with different colours on each graph. Constant fundamental price of 60 is captured by the dotted black line. *Left panels*: stable initial price history. *Right panels*: unstable initial price history. *Top row*: one-period ahead forecasting. *Middle row*: two-periods ahead forecasting. *Bottom row*: three-periods ahead forecasting.

First, if we consider observations of price dynamics in markets with different initial price history, we can conclude that, for all forecasting horizon lengths, instability of initial price history leads to higher instability in the market. Markets with an initially stable price demonstrate only slight fluctuations of the price around the fundamental value. Much of this variation can also be attributed to the noise term added to the pricing equation 4.3. By contrast, most of the markets with unstable initial price history display periodic fluctuations with substantial divergence of the price from the fundamental value. This sensitivity of dynamics to initial conditions is in line with the previous LtF experiments, but initial prices previously had never been explicitly controlled. Therefore, this first result is partly in line with initial predictions, but more importantly, it can shed light on several previously unexplained results of the LtF experiments.

Second, we can observe that increasing the length of the forecasting horizon has different effects in different environments. In markets with stable initial history, an increase in the length of the forecasting horizon has a negligibly small stabilizing effect which nevertheless smoothes away even a small fluctuation observed in the one-period ahead forecasting markets. A clearer picture is provided by the markets with unstable price histories: an amplitude of cyclic price fluctuations in one-period ahead markets is the largest; prices in two-periods ahead markets have smaller periodic fluctuations; and most of the three-period ahead market cycles vanish shortly after the start.

We formalise the preliminary results and test several hypotheses for the effects of horizon length by comparing volatility of price dynamics in markets of different treatments. Additionally, several measures of bubbles used in the literature are reported: Relative Absolute Deviation[12] (RAD), Relative Deviation (RD), and Price Amplitude[13] (PA). Results are presented in Table 4.2. The average price on most of the markets is close to the fundamental value of 60, but markets differ in terms of volatility and bubble measures. We test these differences for statistical significance and report the following results with regard to the hypotheses introduced in Section 4.4.

We summarize aggregate experimental results and conclude that, when looking at the average prices of all treatments, Hypothesis 4 cannot be re-

[12]Both RAD and RD were introduced in [102] and later used in [74].
[13]For examples of PA treatment see, for example, [85] and [72].

Price history	Horizon length	Market	Mean price	Variance	RAD	RD	PA
Stable	H=1	S1G1	57.6678	4.27	0.04	-0.04	0.13
		S1G2	58.6889	0.98	0.02	-0.02	0.07
		S1G3*	23.2870	106.63	0.61	-0.61	0.59
		S1G4	58.2308	13.33	0.04	-0.03	0.23
	H=2	S2G1	56.6320	0.98	0.06	-0.06	0.08
		S2G2*	64.3277	205.65	0.09	0.07	1.61
		S2G3	59.3702	1.17	0.02	-0.01	0.12
		S2G4	59.0789	0.49	0.02	-0.02	0.05
	H=3	S3G1	63.4562	1.64	0.06	0.06	0.10
		S3G2	58.9571	1.53	0.02	-0.02	0.08
		S3G3	59.7626	0.38	0.01		0.05
Unstable	H=1	U1G1	63.2846	1.58	0.05	0.05	0.09
		U1G2	60.6354	2.68	0.02	0.01	0.10
		U1G3	52.2529	571.48	0.36	-0.13	1.43
		U1G4	63.8478	48.46	0.10	0.06	0.63
		U1G5	61.2963	1.14	0.02	0.02	0.07
		U1G6	72.2549	2039.10	0.62	0.20	2.87
		U1G7	231.6256	93474.39	3.17	2.86	15.25
		U1G8	67.6405	3253.69	0.42	0.13	5.71
	H=2	U2G1	54.8434	45.76	0.12	-0.09	0.42
		U2G2	51.9298	344.63	0.26	-0.13	1.14
		U2G3	58.4176	198.11	0.19	-0.03	0.88
		U2G4	60.3232	139.88	0.17	0.01	0.64
		U2G5	57.8411	171.40	0.18	-0.04	0.92
		U2G6	53.5109	362.53	0.27	-0.11	1.20
	H=3	U3G1	60.4672	4.76	0.03	0.01	0.22
		U3G2*	67.57	162.07	0.13	0.13	1.48
		U3G3	60.56	6.61	0.04	0.01	0.17
		U3G4	54.37	123.75	0.15	-0.09	0.92
		U3G5	61.33	21.19	0.05	0.02	0.47
		U3G6	58.5669	13.05	0.04	-0.02	0.25
Unstable with strong feedback	H=3	SU3G1	88.9184	4029.25	0.50	0.48	3.98
		SU3G2	60.7229	0.94	0.02	0.01	0.08
		SU3G3	59.7607	243.52	0.20		1.05
		SU3G4	55.7279	3.30	0.07	-0.07	0.17
		SU3G5	50.6042	381.77	0.27	-0.16	1.92
		SU3G6	45.0973	216.59	0.29	-0.25	1.04

Table 4.2: Different price statistics and bubble measures calculated for each of the 37 experimental markets: Mean, Variance, Relative Absolute Deviation (RAD), Relative Deviation (RD), and Price Amplitude (PA). The first ten observations of the price–which were available for participants as initial history–and non-incentivised prices are excluded, leaving 50 price observations per market. Three markets with outlying observations are marked with the asterisk: S1G3*, S2G2*, and U3G2*.

jected.[14] Hypothesis 5 can only be rejected for one-period ahead forecasting markets, while it is confirmed for two-periods ahead and three-periods ahead markets, where unstable initial history induced more volatile dynamics.[15] Hypothesis 6 is not supported by the data: all stable initial price markets have similar levels of price volatility. In markets with initially unstable prices, both the one- and two-periods ahead forecasting produces a high level of volatility, whilst the three-periods ahead forecasting market has smaller levels of volatility and the hypothesis of equality with another two treatments can be rejected.

4.5.2 Individual forecasting behaviour

Behavioural heterogeneity in individual forecasts appears to play an important role in aggregate fluctuations of the market prices. We investigate individual forecasting rules as a possible source of mismatch between observed experimental prices and preliminary simulations, which were run using $AR(2)$ behavioural rules. Consider market dynamics in Figure 4.4, where both market prices and individual forecasts are depicted.[16]

First, we do observe heterogeneity in forecasts, and it is higher for unstable initial price history treatments. This is in line with an intuition regarding the complexity of the forecasting task. Unstable history leaves room for participants to have significantly different beliefs about future dynamics. This observation was predicted by simulations provided in Section 4.3, where

[14] Table C.2 contains results of all pair-wise Wilcoxon rank sum tests. The only two treatments, which equality can be rejected at 10% level of significance, is stable history treatments with one-period ahead and three-periods ahead forecasting horizons.

[15] Table C.3 contains results of all pair-wise Wilcoxon rank sum tests for equality of the variances.

[16] More illustrative graphs for all experimental markets are available in the appendix.

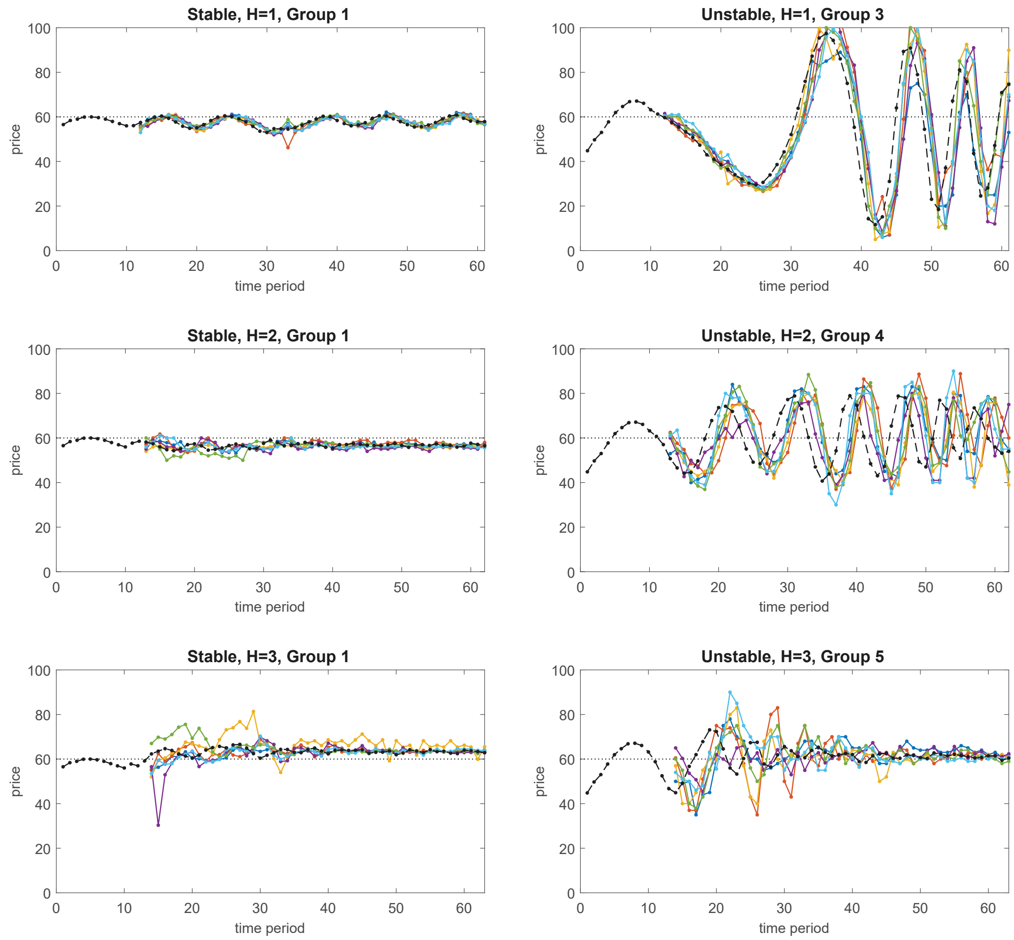

Figure 4.4: Examples of individual forecasts and market price dynamics in six basic treatments, one typical group was chosen for each treatment. Different colours capture different participants' forecasts on each graph. Market price fluctuations are captured by the dashed black line. Constant fundamental price of 60 is captured by the dotted black line. *Left panels*: stable initial price history. *Right panels*: unstable initial price history. *Top row*: one-period ahead forecasting. *Middle row*: two-periods ahead forecasting. *Bottom row*: three-periods ahead forecasting.

convergence was slower in initially unstable markets, even for sophisticated general-form rule such as sample average learning. At the same time, the high level of heterogeneity in the three-periods ahead forecasting markets with stable price histories may suggest that even in a stable environment, the coordination of the three-periods ahead forecasts is harder to achieve at the group level.

Second, we can observe an overall quick coordination in forecasting in all experimental markets, especially in those with shorter horizons. Coordination on the fundamental price forecasting had a stabilizing effect in all the markets with an initially stable price. At the same time, the coordination on non-fundamental predictions, especially in one- and two-periods ahead markets, amplified initial fluctuations of unstable price history markets and led to the aggregate fluctuations during the main body of market price observations in the markets.

Third, coordination of forecasts is overall lower in the three-periods ahead forecasting markets. It contributes to the price volatility in the stable price history markets, but not significantly. On the contrary, in unstable history markets, a longer horizon appears to be responsible for breaking down coordination on the price cycles observed in the two-periods ahead markets. Therefore, we observe a stabilization effect of an increase of the horizon length which does not relate to a discounting factor, but rather has coordination failure origins. We further investigate the importance of coordination in Section 4.5.3.

Data on individual forecasts of participants allow us to access whether the behavioural rules used for the simulation in Section 4.3 is a valid approach to modelling expectation formation. We estimate the following general adaptive forecasting rule, which uses information both on previously observed prices and previously submitted forecasts:

$$p_t^e = \alpha + (\sum_{s=1}^{k} \beta_s p_{t-s}) + (\sum_{s=1}^{l} \gamma_s p_{t-s}^e).$$

We estimate this rule for each participant and the results are reported in Tables C.7-C.10, available in the appendix. Although overall there is a substantial level of heterogeneity in individual forecasting strategies, several coefficients are significant substantially more often then the others: β_1, β_2, and γ_1. These coefficients are significant determinants of forecasting behaviour for 73%, 47%, and 35% of participants, respectively. This rule can be effectively rearranged to obtain an $AR(3)$ rule that is very similar to the one used for simulation in Section 4.3.

Since observed individual forecasting strategies are well described by behavioural rules that are similar to those used in the simulation, the experimental results on an aggregate level remain puzzling. We propose the coordination of expectations as one of the important factors that might be responsible for convergence in the three-periods ahead forecasting markets. Preliminary simulations were run with homogeneous expectations, whilst we observe a lot of heterogeneity in expectations in experimental markets, even among trend-followers.

4.5.3 Coordination

To quantify the coordination of individual forecasts, we first measure individual quadratic errors that measure the accuracy of submitted forecasts. Then we split an error into two components: dispersion and common errors. Dispersion reflects disagreement in submitted forecasts, while common error demonstrates coordination on non-fundamental forecasts. Consider the following measure of average individual forecasting squared error of all N

Price history	Horizon length	Market	Individual Error	Fraction	Dispersion	Fraction	Common Error	Fraction
Stable	H=1	S1G1	3.04	100%	0.80	26%	2.24	74%
		S1G2	1.04	100%	0.11	11%	0.93	89%
		S1G3*	158.47	100%	106.77	67%	51.69	33%
		S1G4	7.17	100%	6.04	84%	1.13	16%
	H=2	S2G1	3.81	100%	1.83	48%	1.98	52%
		S2G2*	1498.64	100%	1056.03	70%	442.61	30%
		S2G3	8.28	100%	5.50	66%	2.78	34%
		S2G4	2.48	100%	1.07	43%	1.41	57%
	H=3	S3G1	18.03	100%	10.73	60%	7.30	40%
		S3G2	13.17	100%	7.18	55%	5.99	45%
		S3G3	1.80	100%	0.76	42%	1.05	58%
Unstable	H=1	U1G1	1.28	100%	0.54	42%	0.75	58%
		U1G2	3.47	100%	2.31	67%	1.16	33%
		U1G3	251.67	100%	44.64	18%	207.03	82%
		U1G4	23.94	100%	5.36	22%	18.58	78%
		U1G5	1.32	100%	0.69	52%	0.63	48%
		U1G6	423.07	100%	98.27	23%	324.80	77%
		U1G7	13200.67	100%	1644.70	12%	11555.96	88%
		U1G8	373.91	100%	37.54	10%	336.37	90%
	H=2	U2G1	127.29	100%	39.64	31%	87.65	69%
		U2G2	684.72	100%	53.76	8%	630.96	92%
		U2G3	410.45	100%	50.60	12%	359.85	88%
		U2G4	331.19	100%	41.35	12%	289.84	88%
		U2G5	466.10	100%	117.50	25%	348.60	75%
		U2G6	1035.65	100%	172.05	17%	863.60	83%
	H=3	U3G1	35.98	100%	23.97	67%	12.01	33%
		U3G2*	1410.13	100%	998.19	71%	411.94	29%
		U3G3	23.63	100%	5.75	24%	17.88	76%
		U3G4	472.73	100%	111.42	24%	361.30	76%
		U3G5	237.59	100%	23.36	10%	214.24	90%
		U3G6	181.44	100%	2.82	2%	178.63	98%
Unstable with strong feedback	H=3	SU3G1	11122.26	100%	9749.27	88%	1372.99	12%
		SU3G2	4.04	100%	1.21	30%	2.83	70%
		SU3G3	677.85	100%	102.41	15%	575.44	85%
		SU3G4	21.34	100%	12.93	61%	8.41	39%
		SU3G5	1744.78	100%	919.55	53%	825.23	47%
		SU3G6	633.61	100%	73.58	12%	560.03	88%

Table 4.3: Mean Square Error (MSE) of individual forecasts and its decomposition into Dispersion and Common Error components calculated for each of the 37 experimental markets. Three markets with outlying observations are marked with the asterisk: S1G3*, S2G2*, and U3G2*.

agents participating in the same market for T periods:

$$MSE = \frac{1}{T}\left(\sum_{t=1}^{T}\frac{1}{N}(\sum_{i=1}^{N}(p_{t,i}^{e}-p_t)^2)\right),$$

where $p_{t,i}^{e}$ is participants i forecast of the price in period t.

This MSE measure of forecasting accuracy, which averages all quadratic errors across time periods and participants, shows how well participants completed their experimental task of "professional forecasters". At the same time, by inspecting graphs with individual forecasts in each market, we can observe strong coordination on non-fundamental predictions, for example, in Group 3 in one-period ahead markets and Group 4 in two-periods ahead markets with unstable initial history, which are depicted in Figure 4.4. To measure both dispersion of forecasts around non-fundamental forecasts and the common error of this group error, we use the following decomposition technique:

$$\begin{aligned} MSE &= \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{N}(\sum_{i=1}^{N}(p^e_{t,i}-p_t)^2)\right) = \\ &= \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{N}(\sum_{i=1}^{N}(p^e_{t,i}-\bar{p^e_t})^2)\right) + \frac{1}{T}\sum_{t=1}^{T}(\bar{p^e_t}-p_t)^2 = \\ &= Dispersion + CommonError, \end{aligned}$$

where $\bar{p^e_t} = \frac{1}{N}\sum_{i=1}^{N} p^e_{t,i}$ is the group average forecast of the price in period t.

Decomposition of individual errors MSE into *Dispersion* and *Common-Error* components is similar in spirit to decomposition of the MSE of an estimator into variance and squared bias; derivation is provided in the appendix. Based on estimates provided in Table 4.3, several notable conclusions can be made. Individual forecasting errors are smaller in the markets with stable initial price history when compared to markets with unstable initial conditions. This is in line with the overall aggregate price stability observed in these markets. Common error fractions in individual errors are much larger in markets with unstable initial price histories. This measure quantifies the coordination of forecasts on non-fundamental prediction strategies and demonstrates connection to observed price fluctuations.

4.5.4 Extra sessions with strong feedback

Results from the markets with initially unstable histories and long forecasting horizons provide somewhat intriguing evidence. Price dynamics in these markets are relatively stable on the aggregate level, however at the individual level, and especially in the initial rounds, forecasts are not concentrated around the fundamental value. Despite that, the price converges and promptly stabilises seemingly because of two effects. The first effect is a weak feedback in these markets. As discussed in detail in Section 4.3, an increase in a forecasting horizon's length results in a mechanical increase in the stabilization forces. The second effect is motivated by the differences in individual behaviour observed in one-period ahead, two-periods ahead, and three-periods ahead forecasting markets–coordination between forecasting behaviour, and in particular on some non-fundamental value, is harder to achieve if the information on accuracy of prediction arrives to participants with the longer delay. To test the hypothesis regarding validity of the second effect, we run extra sessions that are designed to isolate the effect of feedback strength.

We apply new parametrisation of the fundamentals to change the feedback strength in the 3-periods ahead forecasting market. We follow three criteria in search of parameter values for extra sessions: (i) fundamental price is equal to 60 similarly to all markets, (ii) feedback strength is close to 0.95 replicating the feedback strength of one-period ahead forecasting

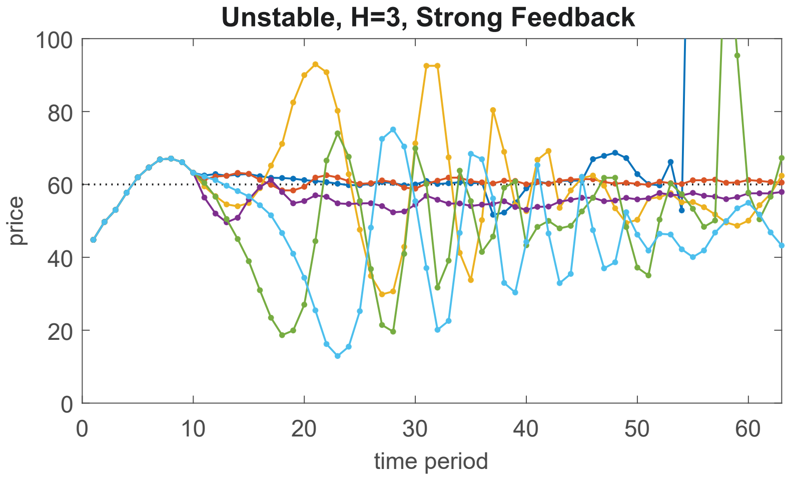

Figure 4.5: Price dynamics in all 6 experimental markets with unstable initial price history, 3-period ahead forecasting and strong feedback. Different groups are captured with different colours. Constant fundamental price of 60 is captured by the dotted black line.

markets, and (iii) values of the interest rate and the dividend are not overcomplicated by the number of decimals. The following parametrisation is utilised in extra sessions as a reasonable compromise between three conflicting criteria: gross interest rate $R = 1.016$ and dividend $\bar{y} = 0.96$. In the market with three-periods ahead forecasting, this parametrization induces the following pricing equation:

$$\text{For } H = 3 : p_t \simeq 0.95 \left(\overline{p}^e_{t+3} + 3\right) + \epsilon_t.$$

Note that this equation with approximately calculated values is identical to the pricing in a one-period ahead market with default parametrization of the fundamentals. Therefore, differences between market dynamics of these two markets can be associated with the information delay structure and related coordination difficulties.

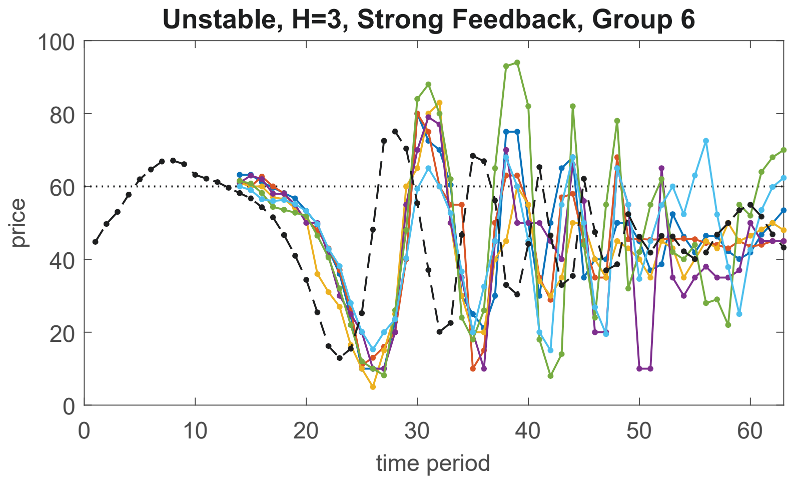

Figure 4.6: Forecasts and price dynamics in the Group 6 of 3-periods ahead forecasting with unstable price history and strong feedback.

We run six markets of long horizons forecasting with strong feedback. The resulting market price dynamics are depicted in Figure 4.5. Three markets clearly exhibit instability and the other three exhibit some fluctuations around the fundamental price similar to one-period ahead markets. Overall, strengthening the feedback brings instability to markets in line with mechanical model comparative static predictions. Observed convergence of price dynamics in several markets to non-fundamental value is another implication of the stronger feedback in line with model intuition.

At the individual level, we also observe behaviour consistent with the initial hypothesis. Figure 4.6 contains a graph with dynamics of the price in Group 6, as well as all individual forecasts. We can observe that coordination appears to be weak and initial concentration of all forecasts on the same trend vanishes over time.

We calculate the median of the standard deviation of predictions to measure the level of non-coordination in all markets. Figure 4.7 contains a graph with non-coordination measure for each market. All markets with initially

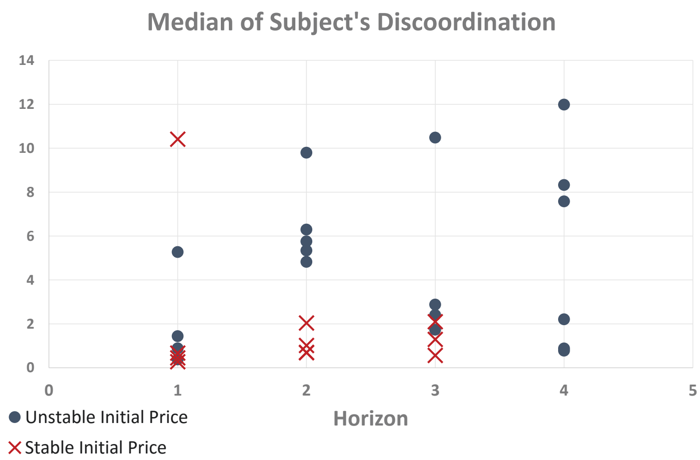

Figure 4.7: Level of discoordination of prediction in all markets measured as median of the standard deviation of predictions.

stable price history exhibit good coordination. For markets with unstable price history, we can observe a positive relationship between the length of the forecasting horizon and the level of non-coordination.

4.6 Conclusion

We run experimental markets with positive feedback effects to investigate the effects induced by an increase in forecasting horizon length on price volatility. We find that in experimental asset markets with initially stable price histories, the length of the forecasting horizon does not have any effect, and markets remain stable. On the contrary, increasing the horizon length on the markets with unstable price histories lead to stabilisation, and periodic fluctuations observed in short forecasting horizon markets with persistent divergence from the fundamental value become less salient.

Observed bubbles and crashes can be explained by simple behavioural forecasting rules that are used by participants in all markets. Subjects tend to extrapolate observe price patters, which reinforce and amplify initial price fluctuation if those are initially observed. We run extra experimental sessions to show that expectation coordination failure can play a role in the stabilising factor. In three-periods ahead forecasting markets, although participants still tend to have trend-following or adaptive expectations, it appears to be harder for participants to coordinate on non-fundamental predictions that can otherwise lead to periodic price fluctuations.

Presented results motivate further investigations about possible policy implication. In markets with short forecasting horizons, investors may be more prone to volatility contagion, which can be addresses by introducing lower limits on asset holding periods. This raises a policy trade-off problem between price stability and price efficiency.

The experimental study in this chapter was focused on markets with homogeneous forecasting horizons. Note that the asset-pricing model introduced in Section 4.2 allows us to derive price equations for markets populated by investors with heterogeneous forecasting horizons. We leave experimental investigation of possible effects induced by co-existence of investors with different forecasting horizons on the price dynamics for future research.

Additionally, we acknowledge several related ways to extend our understanding of the forecasting horizon's effects on price stability. The experiment in which participants forecast in both stable and unstable markets may identify effects of the volatility experience. It may shed light on how prone are different markets to volatility contagion. Further experiments may be run in order to investigate possible effects of the combination of trading task alongside the forecasting task, for example Learning-To-Optimize experiments (see [12]). Alternatively, a standard double auction auction with restrictions on holding period can be used in the experimental setup to study the effects of forecasting and trading periods in more sophisticated market environment.

Chapter 5

Conclusion and Future Work

All three chapters of this thesis generate insightful results, which are of theoretical and empirical importance. The presented papers extend our understanding of both the markets and those forces that may drive prices away from fundamental values. This knowledge is crucial for tailoring appropriate policy responses with price stabilisation interventions. Additionally, all three chapters generate results with insightful extensions and applications that could shape a promising research agenda.

In the second chapter of this thesis, I present results from the laboratory experiment on HSM. In the experiment, subjects choose between: (i) a sophisticated and stabilising, but costly, heuristic, and (ii) a destabilising, but cheap, naïve heuristic. Theoretical predictions are confirmed, and an increase in the costs for the stabilising heuristic generates instability and leads to endogenous fluctuations of price. These results have an important implication for agent-based models that were built on the basis of HSMs and applied to the study of policy implications in complex, nonlinear dynamics settings, e.g., Schmitt and Westerhoff (2015).

The third chapter studies behavioural models of attention adjustment to repetitive binary choice. Models from the different strands of the literature, including HSM, use a logit model to study adaptive choice. Usually, the logit parameter, which represents 'focus' or 'attention', is set to an exogenously fixed value. A behavioural model of attention adjustment, or a 'self-tuning' model where subjects scale down the attention if the payoff difference is not large, demonstrates the best fit to the data. This model is also portable across environments with close estimates of the parameter, and thus, the use of this model addresses important calibration issues discussed in the literature. A self-tuning model of adaptive choice can be incorporated into existing empirical studies to improve the goodness of fit of learning and agent-based models. Additionally, the provided results can be used to endogenise variation in the attention parameter in static equilibrium notions such as in the logit quantal response equilibrium, introduced in Mckelvey and Palfrey (1995).

The forth chapter of this thesis relates to behavioural finance literature and experimentally investigates one of the possible sources of the excess volatility in the markets: the length of the forecasting horizon. The experimental results show that markets with initially unstable price dynamics tend to reinforce price deviations more often as the horizon length increases. Participants extrapolate the price trends in medium-long horizons, which destabilises the market price and leads to bubbles and crashes. To confirm the robustness of the results, future work should focus on running new experiments with markets populated by investors with mixed forecasting horizons, and markets with alternative pricing mechanisms, e.g., a double auction.

Appendix A

Appendix to Chapter 2

A.1 Stylized HSM

In the experiment we use a stylized heuristic switching model. Below we provide a formal analysis of this model and show that the properties of the original Brock-Hommes model are preserved.

The dynamics of the stylized model are given by the 1D equation (2.5). We switch the noise off, $\epsilon_t = 0$, to focus on the analysis of the deterministic skeleton $x_{t+1} = f(x_t)$, with

$$f(x) = \frac{\lambda x}{1 + \exp[\alpha + \beta(x^2 - C)]}\,, \tag{A.1}$$

where $C = W_B - W_A$ can be interpreted as the costs for alternative A. Given the value of the state variable, x_t, we can easily recover the fraction of B-choices as

$$n_{B,t+1} = \frac{1}{1 + \exp[\alpha + \beta(x_t^2 - C)]} = \frac{x_{t+1}}{\lambda x_t}\,. \tag{A.2}$$

In the analysis below we assume $C > 0$ and $\beta \geq 0$, but we will not restrict the other parameters. Note that in the experiment presented in the paper we set $\lambda = 2.1$ and $C = 8$ (High blocks) or $C = 0.1$ (Low blocks). The behavioral parameters α and β, from the discrete choice model, describe subjects' choices.

The following result characterizes the dynamic properties of (A.1).

Proposition A.1.1. *Let f be defined in* (A.1) *with $C > 0$ and $\beta \geq 0$. Consider the system $x_{t+1} = f(x_t)$ and let $n^* = 1/(1 + \exp[\alpha - \beta C])$.*

1. *$x^* = 0$ is a steady state for all parameter values. In this steady state the fraction of agents choosing alternative B is given by $n^*_B = n^*$. This steady state is unique for $\lambda < 1$ and globally stable when $|\lambda n^*| < 1$.*

2. *For $\lambda > 0$, the system undergoes a pitchfork bifurcation when $\lambda = 1/n^* > 1$. At the bifurcation two non-zero steady states, $x^+ > 0$ and $x^- = -x^+$, are created, with corresponding steady state fraction $n^+_B = n^-_B = 1/\lambda$.*

3. *For $\lambda < 0$, a period doubling bifurcation occurs when $\lambda = -1/n^* < -1$. At the bifurcation a two-cycle $\{x^+, x^-\}$, with $x^+ > 0$ and $x^- = -x^+$, is created, with corresponding steady state fraction $n^+_B = n^-_B = -1/\lambda$.*

Proof. It is obvious that $x^* = 0$ is a steady state of the system. Other steady states are solutions of the equation

$$1 = \frac{\lambda}{1 + \exp[\alpha + \beta(x^2 - C)]} \quad \Leftrightarrow \quad \exp[\alpha + \beta(x^2 - C)] = \lambda - 1\,.$$

This equation has no solutions when $\lambda \leq 1$. For the case $\lambda > 1$ we find that

$$x^2 = \frac{\ln(\lambda - 1) - \alpha}{\beta} + C\,.$$

Therefore, for $\alpha < \ln(\lambda - 1) + \beta C$ (which is equivalent to $\lambda > 1/n^*$) the original system has two non-zero steady states

$$x^{\pm} = \pm\sqrt{\frac{\ln(\lambda - 1) - \alpha}{\beta} + C}\,. \tag{A.3}$$

In both steady states, n_B^*, the fraction of agents choosing B is equal to $1/\lambda$.

To investigate the local stability properties of the steady states we take the derivative of $f(x)$:

$$f'(x) = \frac{\lambda}{1 + \exp[\alpha + \beta(x^2 - C)]} - 2x^2\lambda\beta\frac{\exp[\alpha + \beta(x^2 - C)]}{(1 + \exp[\alpha + \beta(x^2 - C)])^2}.$$

In the zero steady state the slope is $\lambda/(1 + \exp[\alpha - \beta C]) = \lambda n^*$. Thus it is locally asymptotically stable for $|\lambda n^*| < 1$. Moreover, we can also write

$$|x_{t+1}| = |\lambda x_t|\frac{1}{1 + \exp[\alpha + \beta(x^2 - C)]} \leq |\lambda x_t|\frac{1}{1 + \exp[\alpha - \beta C]} = |\lambda||x_t|n^*\,.$$

Therefore, when $|\lambda n^*| < 1$ the map is a contraction and the zero steady state is globally stable.

The slope at the zero steady state is equal to 1 when $\lambda = 1/n^*$. Given that the map f is odd, system is symmetric with respect to x. Therefore at $\lambda = 1/n^*$ the system undergoes a pitchfork bifurcation. We have shown above that the two steady states for $\lambda > 1/n^*$ are x^+ and x^- as given by (A.3).

When $\lambda = -1/n^*$, the slope at the zero steady state is equal to -1 and, hence, the system undergoes a period-doubling bifurcation. For $\lambda < -1/n^*$ we can define x^+ and x^- in analogy with (A.3) as

$$x^{\pm} = \pm\sqrt{\frac{\ln(|\lambda| - 1) - \alpha}{\beta} + C}\,.$$

Then

$$f(x^+) = \frac{\lambda x^+}{1 + \exp[\alpha + \beta((x^+)^2 - C)]} = \frac{\lambda x^+}{|\lambda|} = -x^+ = x^- \quad \text{and} \quad f(x^-) = x^+\,.$$

Thus, $\{x^+, x^-\}$ is a two-cycle created at the moment of the bifurcation. It then follows from (A.2) that along this cycle $n_B \equiv -1/\lambda$. □

We illustrate the properties of the system in Fig. A.1 with six bifurcation diagrams. To build them we fix $\lambda = 2.1$ and consider two values of cost, as in the experiment. The diagrams in the left panels are built for High cost ($C = 8$) and the diagrams in the right panels are built for Low cost ($C = 0.1$). For each diagram we fix a specific value of α as follows: $\alpha = 0.4$ for the diagrams in the top panels, $\alpha = 0$ for the diagrams in the middle panels, and $\alpha = -0.4$ for the diagrams in the lower panels. Finally, for each diagram, we vary β between 0 and 5 for High cost and between 0 and 50 for Low cost, and show (against every value of β) 500 points of the dynamics of system (A.1), initialized at $x_0 = 0.25$, after 400 transitory periods.[1]

We note that the left and right diagrams look similar for every row, exhibiting the typical shape of the Brock-Hommes framework. When β increases, first, the zero steady state loses stability through the pitchfork

[1]Since the noise ϵ_t is off, the state variable x will stay positive on every trajectory for such initial value. The bifurcation diagrams for negative initial value look similar to those shown in Fig. A.1, but it is in the negative domain of the state variable.

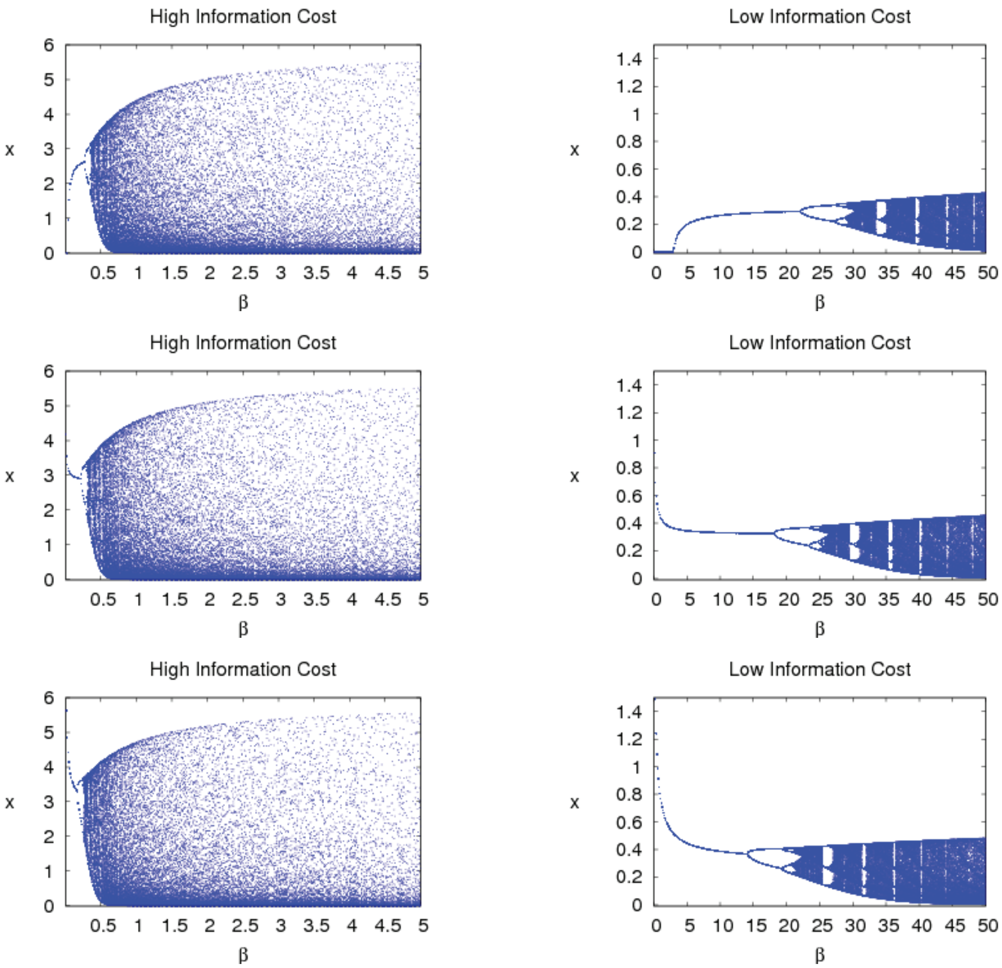

Figure A.1: Bifurcation diagrams with respect to the Intensity of Choice, β, for system (A.1) with $\lambda = 2.1$. *Left panels*: High cost, $C = 8$. *Right panels*: Low cost, $C = 0.1$. Predisposition parameter: $\alpha = 0.4$ (top panels), $\alpha = 0$ (middle panels) and $\alpha = -0.4$ (lower panels).

bifurcation, as proven in Proposition A.1.1.*2*. Two non-zero steady states are created, and we observe one of them, x^+. It is straightforward to check that the slope of f in each of the steady states is given by

$$f'(x^{\pm}) = 1 - 2\beta \left(\frac{\ln(\lambda - 1) - \alpha}{\beta} + C \right) \left(1 - \frac{1}{\lambda} \right).$$

Therefore, for any value of λ and the other parameters, when β increases the slope will decrease and the non-zero steady state will eventually lose stability through a period doubling bifurcation. A further increase in β, as shown in Fig. A.1, leads to complicated dynamics with increasing ranges of the state variable. The points of bifurcations occur much earlier, however, for the High cost case than for the Low cost case. Bifurcations diagrams do not change significantly also with changes in α (except that the zero steady state is never stable for relatively high α, but it is stable for low values of β, when α is low).

Our choice of parameter values in the experiment was motivated by these diagrams, as they predict different behavior in two treatments for very large range of values of behavioral parameters, β and α, under assumption of similar behavior in two treatments.

A.2 Experimental Instructions

Below we provide the short version of the instructions. For the full version which includes the examples of screens explained to the participants, quiz, and questionnaire, see the On-line Appendix.

General information

Today you will participate in an experiment which will require you to make economic decisions. During the experiment you will be able to earn a number of points. The better your decisions are, the more points you are likely to earn. These points will be converted into Australian dollars after the experiment.

Task overview

Consider a situation in which you are offered a choice between two investment alternatives. As in real markets, the profits of these alternatives depend on the decisions of all market participants (including yourself) and also on chance.

This experiment has **two parts**[2], each consisting of **40 decision periods**. In the beginning of each part you will be randomly matched to several other participants in this experiment that will be active in the same market. The same participants will be in the same market as you during this **first part** of the experiment. Every decision period you will have to choose one of the two alternatives, **A** or **B**. After you, and all other participants, made

[2]In additional sessions with High Long block Instructions were changed accordingly.

a decision, the profits of both alternatives are determined. Your decision will determine how many points you may get for this period. This number of points will be used to determine your final earnings. The exact procedure that determines your final earnings from the experiment is explained below (under **Remuneration**).

Example: if in some decision period you have chosen **A**, and then (on the basis of this decision and others' choices) the profits of the alternatives turned out to be 7.35 for **A** and 2.53 for **B**, then for this period you will be able to get 7.35 points.

Before you select one of the alternatives you will be shown the profits of both alternatives for all **previous** periods, as well as the history of your own decisions. The only exception is the very first period (for each of the two parts), when no history is available. Note that the prospective profits are not known to you at the moment of your decision. After you made your decision for a given period, you will get all the information on realized profits of both alternatives as well as your number of points for this period.

After the first part of the experiment is finished, i.e., after the first 40 decision periods, **the second part** of the experiment starts. For this part you will be assigned to another group with several other participants (some, but not all, of them may be from the market that you participated in in the first part). This new group of participants now participates in another market, which is similar but where profits of the two alternatives are determined in a different manner. You will perform the same task of choosing between the two alternatives for another 40 periods.

You will see a countdown timer at every decision period. You do not have to finalize your decision in that time interval, but we nevertheless ask you to make your decisions in a timely manner.

Remuneration

Your total earnings in this experiment will be determined after the experiment by the following procedure. Two different periods will be chosen randomly by the computer program from the 40 periods of the first part of the experiment. Another two periods will also be chosen randomly by the computer program from the 40 periods of the second part of the experiment.

The sum of your points from these 4 periods divided by 2 will constitute the final earnings in Australian Dollars. Your final earnings will be rounded to the nearest value with a half-dollar precision.

In addition you will receive 10 AUD as show-up fee.

Note that with this procedure your performance for **every period** of the experiment potentially matters for your final monetary payoff.

At the end of **each part** of the experiment you will be shown the final screen with information about all your decisions in this part as well as profits of the two alternatives. You will also be informed which decision periods were randomly chosen for payment purposes. If you wish, you can keep this information for yourself in the last page of these Instructions.

Additional information

- Before the experiment starts you will have to take a short **quiz** which is designed to check your understanding of the instructions.

- By the end of the experiment you will be asked to answer a **questionnaire**. Inserted data will be processed in nameless form only. Please fill in the correct information.

- During the experiment any communication with other participants, whether verbal or written, is forbidden. The use of phones, tablets or any other gadgets is not allowed. Violation of the rules can result in exclusion from the experiment without any remuneration.

- Please follow the instructions carefully at all the stages of the experiment. If you have any questions or encounter any problems during the experiment please raise your hand and the experimenter will come to help you.

Please ask any question you have now!

Appendix B

Appendix to Chapter 3

B.1 Estimations of basic non-tuning models

Table B.1: The logit parameter estimates for the standard model and performance measure model with and without predisposition.

Model		Simple DCM		DCM with predisposition			
№	**Block**	$\boldsymbol{\beta}$	***S.E.***	$\boldsymbol{\alpha}$	***S.E.***	$\boldsymbol{\beta}$	***S.E.***
1	WN	0.19	0.03	0.97	0.09	0.16	0.04
2	WN	0.30	0.04	0.42	0.10	0.28	0.04
3	BH	0.41	0.04	-0.22	0.11	0.38	0.04
4	BH	0.48	0.04	-0.39	0.12	0.42	0.04
5	SI	3.33	0.27	0.40	0.13	3.56	0.30
6	SI	4.04	0.32	0.66	0.15	4.66	0.41
7	SP	3.57	0.95	0.26	0.11	3.33	0.84
8	SP	4.91	0.93	0.03	0.11	4.93	0.93
9	SP	11.00	1.46	-0.13	0.13	11.75	1.69
10	SP	9.90	1.50	0.17	0.12	9.26	1.51
11	SP	2.06	0.66	0.23	0.11	2.23	0.66
12	SP	8.29	1.40	-0.38	0.13	10.56	1.74
13	SP	4.53	1.03	0.13	0.11	4.75	1.04
14	SP	8.59	1.40	0.28	0.11	8.60	1.36
15	UP	0.09	0.02	0.31	0.11	0.08	0.02
16	UP	0.13	0.02	0.41	0.12	0.12	0.02
17	UP	0.11	0.02	0.30	0.11	0.10	0.02
18	UP	0.17	0.02	0.26	0.12	0.16	0.02
19	UP	0.10	0.02	0.15	0.11	0.09	0.02
20	UP	0.14	0.02	0.05	0.11	0.13	0.03
21	UP	0.14	0.02	0.21	0.12	0.13	0.02
22	UP	0.17	0.02	0.38	0.12	0.16	0.02
23	UPL	0.15	0.01	0.08	0.05	0.15	0.01
24	UPL	0.17	0.01	0.13	0.05	0.17	0.01
	Pooled	0.17	0.01	0.14	0.05	0.17	0.01

Table B.2: The logit parameter estimates for the standard model and performance measure model with three lags.

№	Block	β_1	*S.E.*	β_2	*S.E.*	β_3	*S.E.*
1	WN	0.21	0.03	0.02	0.03	0.12	0.03
2	WN	0.32	0.04	-0.10	0.04	0.01	0.04
3	BH	0.84	0.08	-0.43	0.06	-0.08	0.05
4	BH	1.11	0.10	-0.59	0.08	-0.09	0.06
5	SI	3.80	0.39	-0.33	0.49	-0.25	0.38
6	SI	5.38	0.48	-1.35	0.55	-0.14	0.44
7	SP	3.45	0.95	0.08	0.68	0.78	0.61
8	SP	5.21	0.96	-0.84	0.82	-0.75	0.81
9	SP	10.93	1.59	1.55	1.35	-1.65	1.47
10	SP	9.79	1.52	2.09	0.90	-0.10	0.96
11	SP	1.91	0.67	0.29	0.49	0.22	0.47
12	SP	8.74	1.49	-1.35	1.40	-0.86	1.31
13	SP	4.79	1.02	-1.28	0.94	0.15	0.82
14	SP	8.82	1.42	-0.93	0.74	-0.63	0.74
15	UP	0.08	0.02	0.02	0.02	-0.01	0.01
16	UP	0.12	0.02	0.01	0.02	0.01	0.02
17	UP	0.11	0.02	-0.02	0.01	0.00	0.01
18	UP	0.17	0.02	0.01	0.02	-0.02	0.02
19	UP	0.08	0.02	0.01	0.02	0.05	0.02
20	UP	0.12	0.03	0.01	0.02	0.02	0.02
21	UP	0.13	0.02	0.01	0.02	0.03	0.02
22	UP	0.18	0.02	-0.01	0.02	0.01	0.02
23	UPL	0.15	0.01	0.01	0.01	0.02	0.01
24	UPL	0.16	0.01	0.03	0.01	0.01	0.01
	Pooled	0.17	0.01	0.00	0.01	0.00	0.01

Table B.3: The logit parameter estimates for the standard model and performance measure model with predisposition and two lags.

№	Block	α	*S.E.*	β_1	*S.E.*	β_2	*S.E.*
1	WN	1.01	0.10	0.17	0.04	-0.05	0.04
2	WN	0.54	0.11	0.31	0.04	-0.15	0.04
3	BH	0.01	0.15	0.86	0.09	-0.48	0.06
4	BH	-0.08	0.18	1.07	0.12	-0.61	0.08
5	SI	0.38	0.13	3.91	0.41	-0.44	0.33
6	SI	0.59	0.16	5.78	0.54	-1.33	0.36
7	SP	0.27	0.11	3.20	0.88	0.60	0.57
8	SP	0.00	0.11	5.19	0.97	-0.98	0.81
9	SP	-0.15	0.13	11.31	1.72	1.39	1.33
10	SP	0.17	0.12	9.15	1.52	2.10	0.89
11	SP	0.25	0.11	2.10	0.68	0.52	0.45
12	SP	-0.36	0.13	10.65	1.75	-0.61	1.38
13	SP	0.12	0.11	4.97	1.04	-1.12	0.93
14	SP	0.26	0.12	8.63	1.36	-0.64	0.76
15	UP	0.31	0.11	0.07	0.02	0.01	0.01
16	UP	0.40	0.12	0.12	0.02	0.02	0.02
17	UP	0.32	0.11	0.11	0.02	-0.03	0.02
18	UP	0.26	0.12	0.16	0.02	0.01	0.02
19	UP	0.12	0.11	0.09	0.02	0.02	0.02
20	UP	0.04	0.12	0.13	0.03	0.01	0.02
21	UP	0.18	0.12	0.12	0.02	0.02	0.02
22	UP	0.43	0.13	0.17	0.02	-0.03	0.02
23	UP	0.08	0.05	0.15	0.01	0.01	0.01
24	UP	0.12	0.05	0.16	0.01	0.03	0.01
	Pooled	0.14	0.05	0.17	0.01	0.00	0.01

B.2 The Model Confidence Set approach

I closely follow the notations introduced in the original paper and summarise the MCS approach in general terms as follows. Using the initial set of competing models $\mathcal{M}^0$, empirical criterion of evaluation, for instance value of Log-Likelihood function, and confidence level α, the MCS approach allows to identify a set of equally good models $\mathcal{M}^*$, which contains the best model with a given level of confidence.

Denote the initial set of model by $\mathcal{M}^0$, which contains m different models indexed by i. Performance of each model is evaluated by the loss function value $L_{i,t}$, which represents models performance in predicting observed data at time period t. Define relative performance for every pair of models $d_{i,j}$ as follows:

$$d_{ij,t} \equiv L_{i,t} - L_{j,t}.$$

Alternative models are compared by the expected loss $\mu_{ij} = E(d_{ij})$. In such a way model i is preferred to the model j if $\mu_{ij} < 0$. Accordingly, the set of superior models $\mathcal{M}^*$ is defined as follows:

$$\mathcal{M}^* \equiv \{i \in \mathcal{M}^0 : \mu_{ij} \leq 0 \text{ for all } j \in \mathcal{M}^0\}.$$

The set $\mathcal{M}^*$ is identified as a result of the sequence of tests. Each test is design to check if any particular element in the set is inferior to elements. If hypothesis of equal fitness is rejected the corresponding element is eliminated:

$$H_{0,\mathcal{M}} : \mu_{ij} = 0 \text{ for all } i, j \in \mathcal{M}^0.$$

Iterative process of elimination in the MCS approach is build upon two essential components. The first component is equivalence test $\delta_{\mathcal{M}}$, such that $\delta_{\mathcal{M}} = 0$ corresponds to $H_{0,\mathcal{M}}$ acceptance, and $\delta_{\mathcal{M}} = 1$ corresponds to rejection. The second component is elimination rule $e_{\mathcal{M}}$ which identifies the element to be excluded from the set at the next iteration. Setting initially $\mathcal{M} = \mathcal{M}^0$, the following algorithm leads to the α-confidence set of models $\mathcal{M}^*_{1-\alpha}$:

1. Test $H_{0,\mathcal{M}}$ with $\delta_{\mathcal{M}}$ at the level α.
2. If $\delta_{\mathcal{M}} = 0$, set $\mathcal{M}^*_{1-\alpha} = \mathcal{M}$. If $\delta_{\mathcal{M}} = 1$, use $e_{\mathcal{M}}$ to identify inferior model m_j, define $\mathcal{M} = \mathcal{M} \backslash m_j$, and repeat the iteration.

Equivalence test $\delta_{\mathcal{M}}$ and elimination rule $e_{\mathcal{M}}$ are based on constructed t-statistics. Denote the relative sample loss statistics by $\bar{d_{ij}} \equiv n^{-1} \sum_{t=1}^{n} d_{ij,t}$ and $\bar{d_i} \equiv m^{-1} \sum_{j \in \mathcal{M}} \bar{d_{ij}}$, t-statistics are constructed as follows:

$$t_{ij} = \frac{\bar{d_{ij}}}{\sqrt{v\hat{a}r(\bar{d_{ij}})}} \text{ and } t_i = \frac{\bar{d_i}}{\sqrt{v\hat{a}r(\bar{d_i})}} \text{ for } i, j \in \mathcal{M}.$$

Given this t-statistics, test statistics for hypothesis and elimination rulesare constructed as follows:

$$T_{R,\mathcal{M}} = \max_{i,j \in \mathcal{M}} |t_{ij}| \text{ and } T_{max,\mathcal{M}} = \max_{i \in \mathcal{M}} t_i.$$

Asymptotic distributions are obtained with bootstrap methods, for instance stationary bootstrap scheme proposed in **(author?)** [89]. In this case the MCS algorithm has straightforward implementation, where $T_{R,\mathcal{M}}$ is used to test if $\mathcal{M}$ is a confidence set, and model with maximum value of $T_{max,\mathcal{M}}$ is eliminated in case hypothesis is rejected.

B.3 Estimations of absolute value scaling models

Table B.4: The logit parameter estimates for the self-tuning model with the maximum absolute values scaling and performance measure model with 2 lags and without predisposition.

№	Block	β_1	*S.E.*	β_2	*S.E.*
1	WN	0.82	0.15	0.24	0.15
2	WN	1.47	0.19	-0.36	0.18
3	BH	4.61	0.42	-2.73	0.34
4	BH	6.45	0.58	-3.77	0.46
5	SI	3.67	0.37	-0.56	0.30
6	SI	5.17	0.46	-1.46	0.33
7	SP	1.16	0.24	0.46	0.23
8	SP	1.39	0.24	-0.33	0.22
9	SP	1.58	0.24	-0.22	0.24
10	SP	1.76	0.29	0.45	0.26
11	SP	1.49	0.24	0.60	0.23
12	SP	1.03	0.19	-0.27	0.19
13	SP	0.97	0.23	-0.11	0.23
14	SP	1.80	0.26	-0.22	0.22
15	UP	1.42	0.23	0.42	0.21
16	UP	1.61	0.25	0.25	0.23
17	UP	1.75	0.27	-0.33	0.24
18	UP	1.89	0.25	0.02	0.21
19	UP	0.95	0.23	0.19	0.23
20	UP	1.22	0.23	0.02	0.21
21	UP	1.45	0.22	0.32	0.21
22	UP	1.91	0.24	-0.10	0.22
23	UP	1.03	0.09	0.14	0.09
24	UP	1.22	0.09	0.28	0.08
	Pooled	1.51	0.09	0.05	0.09

Table B.5: The logit parameter estimates for the self-tuning model with the maximum absolute values scaling and performance measure model with 3 lags without predisposition.

№	Block	β_1	*S.E.*	β_2	*S.E.*	β_3	*S.E.*
1	WN	0.92	0.16	0.08	0.16	0.52	0.15
2	WN	1.49	0.19	-0.40	0.19	0.11	0.17
3	BH	5.29	0.47	-3.61	0.45	0.81	0.23
4	BH	6.50	0.58	-3.87	0.51	0.15	0.32
5	SI	3.66	0.37	-0.34	0.45	-0.23	0.36
6	SI	5.17	0.46	-1.28	0.50	-0.20	0.42
7	SP	1.12	0.24	0.33	0.25	0.33	0.24
8	SP	1.40	0.24	-0.30	0.22	-0.23	0.22
9	SP	1.76	0.27	0.04	0.26	-0.64	0.26
10	SP	1.78	0.30	0.56	0.27	-0.37	0.27
11	SP	1.51	0.24	0.57	0.23	0.15	0.24
12	SP	1.01	0.20	-0.18	0.20	-0.24	0.18
13	SP	0.98	0.24	-0.17	0.25	0.14	0.23
14	SP	1.82	0.27	-0.21	0.22	-0.10	0.25
15	UP	1.39	0.23	0.37	0.22	0.09	0.13
16	UP	1.62	0.26	0.26	0.24	-0.03	0.11
17	UP	1.72	0.28	-0.38	0.26	0.06	0.11
18	UP	1.87	0.25	0.00	0.22	0.04	0.12
19	UP	0.86	0.23	-0.01	0.24	0.30	0.16
20	UP	1.12	0.24	-0.06	0.22	0.14	0.12
21	UP	1.47	0.23	0.35	0.22	-0.05	0.10
22	UP	1.89	0.24	-0.23	0.24	0.19	0.15
23	UP	0.98	0.09	0.09	0.09	0.10	0.05
24	UP	1.22	0.09	0.27	0.09	0.01	0.05
	Pooled	1.51	0.10	0.04	0.09	0.01	0.05

Table B.6: The logit parameter estimates for the self-tuning model with the maximum absolute values scaling and performance measure model with predisposition and one lag.

№	Block	α	*S.E.*	β_1	*S.E.*
1	WN	0.98	0.09	0.74	0.17
2	WN	0.44	0.10	1.29	0.18
3	BH	-0.24	0.11	2.13	0.22
4	BH	-0.38	0.12	2.58	0.24
5	SI	0.41	0.13	3.42	0.29
6	SI	0.65	0.15	4.39	0.39
7	SP	0.01	0.12	1.38	0.22
8	SP	-0.09	0.11	1.31	0.23
9	SP	-0.29	0.13	1.68	0.23
10	SP	0.08	0.12	1.88	0.28
11	SP	0.03	0.11	1.52	0.24
12	SP	-0.35	0.13	1.18	0.21
13	SP	0.05	0.11	0.93	0.22
14	SP	0.24	0.11	1.77	0.25
15	UP	0.22	0.11	1.47	0.22
16	UP	0.37	0.12	1.66	0.23
17	UP	0.23	0.11	1.55	0.25
18	UP	0.21	0.12	1.86	0.24
19	UP	0.11	0.11	1.00	0.20
20	UP	0.03	0.11	1.21	0.22
21	UP	0.18	0.12	1.49	0.21
22	UP	0.34	0.12	1.73	0.21
23	UPL	0.07	0.05	1.04	0.09
24	UPL	0.14	0.05	1.26	0.09
	Pooled	0.11	0.05	1.51	0.09

Table B.7: The logit parameter estimates for the self-tuning model with the maximum absolute values scaling and performance measure model with predisposition and 2 lags.

№	Block	α	*S.E.*	β_1	*S.E.*	β_2	*S.E.*
1	WN	1.01	0.10	0.78	0.17	-0.20	0.17
2	WN	0.54	0.11	1.45	0.20	-0.63	0.20
3	BH	-0.23	0.13	4.25	0.44	-2.66	0.34
4	BH	-0.29	0.17	5.82	0.63	-3.57	0.46
5	SI	0.38	0.13	3.75	0.39	-0.41	0.31
6	SI	0.58	0.16	5.51	0.51	-1.31	0.35
7	SP	-0.02	0.12	1.17	0.24	0.46	0.24
8	SP	-0.10	0.11	1.40	0.24	-0.34	0.22
9	SP	-0.28	0.14	1.76	0.27	-0.14	0.24
10	SP	0.07	0.12	1.72	0.29	0.44	0.26
11	SP	-0.02	0.12	1.50	0.24	0.61	0.23
12	SP	-0.33	0.13	1.21	0.21	-0.12	0.19
13	SP	0.05	0.11	0.97	0.23	-0.11	0.23
14	SP	0.23	0.12	1.80	0.26	-0.16	0.22
15	UP	0.19	0.12	1.37	0.23	0.37	0.21
16	UP	0.36	0.12	1.58	0.25	0.18	0.24
17	UP	0.25	0.11	1.70	0.27	-0.39	0.25
18	UP	0.21	0.12	1.86	0.24	-0.01	0.22
19	UP	0.10	0.11	0.93	0.23	0.15	0.23
20	UP	0.03	0.12	1.21	0.24	0.01	0.22
21	UP	0.14	0.12	1.40	0.22	0.25	0.22
22	UP	0.39	0.13	1.89	0.25	-0.32	0.24
23	UPL	0.07	0.05	1.02	0.09	0.13	0.09
24	UPL	0.12	0.05	1.21	0.09	0.26	0.09
	Pooled	0.11	0.05	1.50	0.09	0.03	0.09

Table B.8: The logit parameter estimates for the self-tuning model with the maximum absolute values scaling and performance measure model with predisposition and 3 lags.

№	Block	α	*S.E.*	β_1	*S.E.*	β_2	*S.E.*	β_3	*S.E.*
1	WN	1.01	0.10	0.78	0.17	-0.21	0.18	0.03	0.18
2	WN	0.57	0.11	1.42	0.20	-0.58	0.20	-0.19	0.19
3	BH	-0.06	0.15	5.14	0.57	-3.54	0.48	0.77	0.24
4	BH	-0.35	0.20	5.64	0.70	-3.39	0.57	-0.22	0.41
5	SI	0.40	0.13	3.76	0.39	-0.09	0.45	-0.35	0.37
6	SI	0.63	0.16	5.65	0.55	-0.81	0.51	-0.59	0.45
7	SP	-0.03	0.12	1.13	0.24	0.33	0.25	0.34	0.24
8	SP	-0.11	0.11	1.43	0.25	-0.32	0.22	-0.24	0.22
9	SP	-0.22	0.14	1.88	0.28	0.06	0.26	-0.55	0.26
10	SP	0.08	0.12	1.73	0.30	0.56	0.27	-0.38	0.27
11	SP	-0.03	0.12	1.53	0.25	0.58	0.24	0.16	0.24
12	SP	-0.31	0.14	1.19	0.22	-0.09	0.20	-0.11	0.19
13	SP	0.05	0.11	0.99	0.24	-0.17	0.25	0.15	0.23
14	SP	0.23	0.12	1.81	0.26	-0.15	0.23	-0.09	0.25
15	UP	0.18	0.12	1.35	0.23	0.33	0.23	0.08	0.13
16	UP	0.37	0.12	1.63	0.26	0.23	0.25	-0.08	0.11
17	UP	0.25	0.12	1.69	0.27	-0.42	0.26	0.03	0.11
18	UP	0.20	0.12	1.85	0.25	-0.02	0.23	0.02	0.12
19	UP	0.05	0.12	0.86	0.23	-0.02	0.25	0.29	0.16
20	UP	0.01	0.12	1.12	0.25	-0.06	0.22	0.13	0.12
21	UP	0.16	0.13	1.43	0.23	0.29	0.23	-0.07	0.10
22	UP	0.37	0.13	1.88	0.25	-0.38	0.25	0.11	0.14
23	UPL	0.06	0.05	0.98	0.09	0.09	0.09	0.09	0.05
24	UPL	0.13	0.05	1.21	0.09	0.26	0.09	-0.01	0.05
	Pooled	0.11	0.05	1.50	0.10	0.03	0.09	0.00	0.05

B.4 The Brock-Hommes Model

Consider economy, which consists of producers of non-storable good which is supplied to a competitive market. Given expected market price producers supply optimal quantity of good to the market. Equilibrium price is determined by the market clearing condition. Realized profits affect choice of price prediction rule for the next period. Given chosen prediction expectation about next period price is formed and the process is repeated.

Producers are assumed to have identical quadratic cost function:

$$c(q_i) = q_i^2/2b,$$

where q_i is produced quantity and $b > 0$ is production function parameter.

This production function is later substituted into the following profit function:

$$\pi_{i,t} = q_{i,t}p_t - c(q_{i,t}) = q_{i,t}p_t - q_i^2/2b,$$

where $\pi_{i,t}$ is profit, p_t is market price.

From solving profit maximization problem by each agent i, given expected market price $p_{i,t}^e$, one derives linear supply function:

$$s_i(p_{i,t}^e) = q_{i,t} = bp_{i,t}^e.$$

Demand function for the good is set to be linear be definition:

$$d(p_t) = A - Bp_t,$$

where A and B are demand function parameters.

Equilibrium price is determined by the market clearing condition:

$$D_t = S_t,$$

where D_t and S_t are aggregate demand and supply respectively.

Previous condition could be rewritten as follows:

$$D_t = \sum_{i=1}^{N} d_i(p_t) = N(A - Bp_t) = S_t = \sum_{i=1}^{N} s_i(p_{i,t}^e) = b(\sum_{i=1}^{N}(p_{i,t}^e)).$$

We can define average expected price $\overline{p_t^e} = \frac{\sum_{i=1}^{N}(p_{i,t}^e)}{N}$. Using this notation allows to rewrite the market clearing condition as follows:

$$A - Bp_t = b\overline{p_t^e}.$$

In such a case expectations about future market price play crucial role in determining the systems' dynamics. Model presumes there are two predictors available to agents: rational and naive expectations. Costly rational expectation provides perfect foresight prediction and costless naive expectation only uses last observed market price as a prediction of price for the upcoming period.

$$p^e_{r,t+1} = p_{t+1},$$

$$p^e_{n,t+1} = p_t,$$

where $p^e_{r,t+1}$ is rational prediction and $p^e_{n,t+1}$ is naive prediction.

Denote a number of agents choosing rational expectation in time t to be $N_{r,t}$ and number of agents choosing naive expectation in time t to be $N_{n,t} = N - N_{r,t}$, where N is total number of producers. Given such composition of producers in the market, average expected price can be defined as a function of fractions fo different predictors:

$$\overline{p^e_{t+1}} = p^e_{r,t+1}\frac{N_{r,t}}{N} + p^e_{n,+1}\frac{N_{n,t}}{N} = p_{t+1}(1 - n_{n,t}) + p_t n_{n,t},$$

where $n_{n,t}$ is a fraction of naive predictor users.

We can substitute this expression for average price to the market clearing condition to obtain:

$$\frac{A - Bp_{t+1}}{b} = p_{t+1}(1 - n_{n,t}) + p_t n_{n,t}.$$

By rearranging items in preceding expression the following equation for the price dynamics as a function of the previous period price and the fraction of the naive prediction users can be written down:

$$p_{t+1} = \frac{A - p_t b n_{n,t}}{B + b(1 - n_{n,t})} = f(p_t, n_{n,t}).$$

To define how agents choose prediction rule we firstly calculate each predictors corresponding realized profits at time t:

$$\pi_{r,t} = \frac{b}{2} p_t^2 - C,$$

$$\pi_{n,t} = \frac{b}{2} p_{t-1}(2p_t - p_{t-1}),$$

where C represents fixed exogenously given cost of acquiring perfect foresight prediction as a rational expectation.

Agents' decision on particular prediction rule is determined by the discrete choice model, where higher realized profit from using particular prediction, leads to a higher probability for the prediction rule to be chosen for the next period:

$$P_{n,t} = \frac{e^{\beta \pi_{n,p_{t+1}}}}{e^{\beta \pi_{r,p_t}} + e^{\beta \pi_{n,p_t}}} = \frac{1}{1 + e^{\beta(\pi_{r,p_t} - \pi_{n,p_t})}}$$

$$P_{r,t} = \frac{e^{\beta \pi_{r,p_t}}}{e^{\beta \pi_{r,p_t}} + e^{\beta \pi_{n,p_t}}} = 1 - P_{n,t}$$

where β represents agents' intensity of choice.

In case of continuum of agents, the Law of Large Numbers guarantees that the total fraction of users of each rule is exactly equal to the corresponding probability. Then dynamics of fractions are described by following equations:

$$n_{n,t} = \frac{1}{1 + e^{\beta(\pi_{r,p_t} - \pi_{n,p_t})}},$$

$$n_{r,t} = 1 - n_{n,t}.$$

In this case main determinant of the distribution of fractions is difference in profits:

$$\pi_{r,t} - \pi_{n,t} = \frac{b}{2}p_t^2 - C - \frac{b}{2}p_{t-1}(2p_t - p_{t-1}) = \frac{b}{2}(p_t - p_{t-1})^2 - C.$$

We can substitute obtained expression for profits difference in the expression determining $n_{n,t}$ to obtain:

$$n_{n,t} = \frac{1}{1 + e^{\beta(\frac{b}{2}(p_t - p_{t-1})^2 - C)}}.$$

To obtain equation for fractions dynamics previous expression for the period $t + 1$ is rewritten:

$$n_{n,t+1} = \frac{1}{1 + e^{\beta(\frac{b}{2}(p_{t+1} - p_t)^2 - C)}}.$$

Now replace price of p_{t+1} with equation of it's dynamics to get expression for the fraction of naive prediction users as a function of previous period fraction and price:

$$n_{n,t+1} = \frac{1}{1 + e^{\beta(\frac{b}{2}(\frac{A - p_t b n_{n,t}}{B + b(1 - n_{n,t})} - p_t)^2 - C)}} = \frac{1}{1 + e^{\beta(\frac{b}{2}(\frac{A - p_t(B+b)}{B + b(1 - n_{n,t})})^2 - C)}} = g(p_t, n_{n,t}).$$

Now by combining two equations of dynamics for price and naive prediction fraction system of non-linear equations can be obtained.

B.5 Extension of the self-tuning model for N alternatives

Self-tuning model in case of binary-choice problems have both simple form and simple intuitive behavioural interpretation. In this section I would like to discuss how model can be applied in the settings where more than two alternatives are available for the choice. As I introduce general N-alternatives choice model, I will discuss how intuition still applies in particular case of 3 alternatives as an example.

We start by introducing logit choice model for N alternatives. Probability to pick alternative a is defined as follows:

$$P_{a,t} = \frac{\exp(\beta\pi_{a,t-1})}{\sum_{n=1}^{N}\exp(\beta\pi_{n,t-1})}. \tag{B.1}$$

Analogously to the binary case I show that choice can be equivalently presented as a function of payoff differences, by dividing both numerator and denominator by $exp(\beta\pi_{a,t-1})$, which gives the following expression:

$$P_{a,t} = \frac{1}{1+\sum_{n=1,n\neq a}^{N}\exp(\beta(\pi_{n,t-1}-\pi_{a,t-1}))}. \tag{B.2}$$

Self-tuning version of the model as before modifies the logit parameter β:

$$P_{a,t} = \frac{1}{1+\sum_{n=1,n\neq a}^{N}\exp(\frac{\beta}{F}(\pi_{n,t-1}-\pi_{a,t-1}))}. \tag{B.3}$$

We now introduce common normalising factor F for number of previous observations used for normalisation L:

$$F = \max_{k=1..N, j=1..N, T=1..L} |\pi_{k,t-T} - \pi_{j,t-T}|.$$

Lets consider the case with N=3. For the sake of simplicity I also fix $L = 1$. Normalising factor is then can be written as follows:

$$F = \max[|\pi_{1,t-1} - \pi_{2,t-1}|, |\pi_{1,t-1} - \pi_{3,t-1}|, |\pi_{2,t-1} - \pi_{3,t-1}|].$$

We can interpret decision making process as follows: if all profit differences are similar in absolute values then F has also similar values, therefore, logit parameter is constant over time. Once profits experience transition from period where large differences was observed to small differences, then logit parameter is decreasing. Later means that agents become indifferent between options, because they observe indistinguishably different profits in term *relative* to the previously observed differences. If transition takes place from small differences to large differences, then logit parameter will increase because current decision is more important as profit difference is large *relative* to previously observed.

B.6 Simulations

As some heterogeneity among estimations is still present in our parameter values, I run simulation to check that predictions of the augmented switching model are robust to identified differences in estimations of the IoC parameter. I now study the dynamics of modified model for previously estimated parameter values. I include the lowest estimated value $\beta = 0.87$, the largest estimation $\beta = 3.83$ and the best fit $\beta = 1.53$ suggested for calibration.

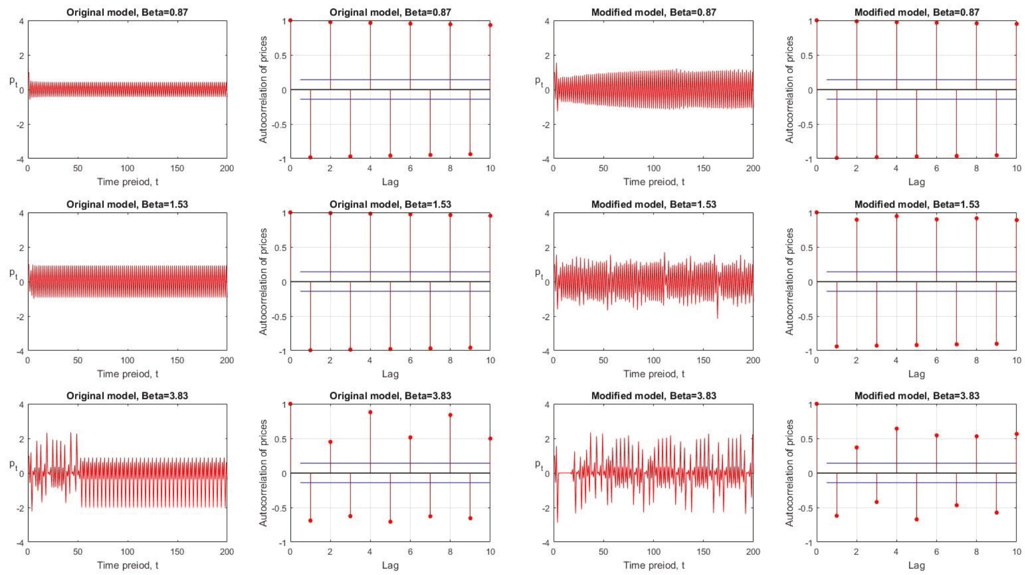

Figure B.1: Simulations of the price dynamics and corresponding sample autocorrelation. Parameters are fixed to $A = 0$, $B = 0.5$, $b = 1.35$ and $C = 1$. *Top row*: $\beta = 0.87$. *Middle row*: $\beta = 1.53$. *Bottom row*: $\beta = 3.83$. *Left panels*: Predictions of the standard model. *Right panels*: Predictions of the modified model.

Figure B.1 provides graphs of both modified and original model for the values attached. We can observe that opposite to the original model dynamics, which as predicted settles on the two-period cycle, modified model induce non-stable dynamics, especially for a large value of β. This suggest, that according to experimentally observed range of modified IoC parameter values, only unstable price dynamics is produced, although autocorrelation structure very close to two-period cycles.

Next exercise is devoted to studying effects of changes in rational expectation cost C. I vary the value of C from 0.1 to 10, fixing the rest of parameters to be the same as before and setting $\beta = 1.53$. Simulation are presented in Figure B.2. Similarly to original model, decrease in information cost brings stability to price dynamics, which can be observed in the middle row graphs. On the contrary, increase in the cost destabilises price dynam-

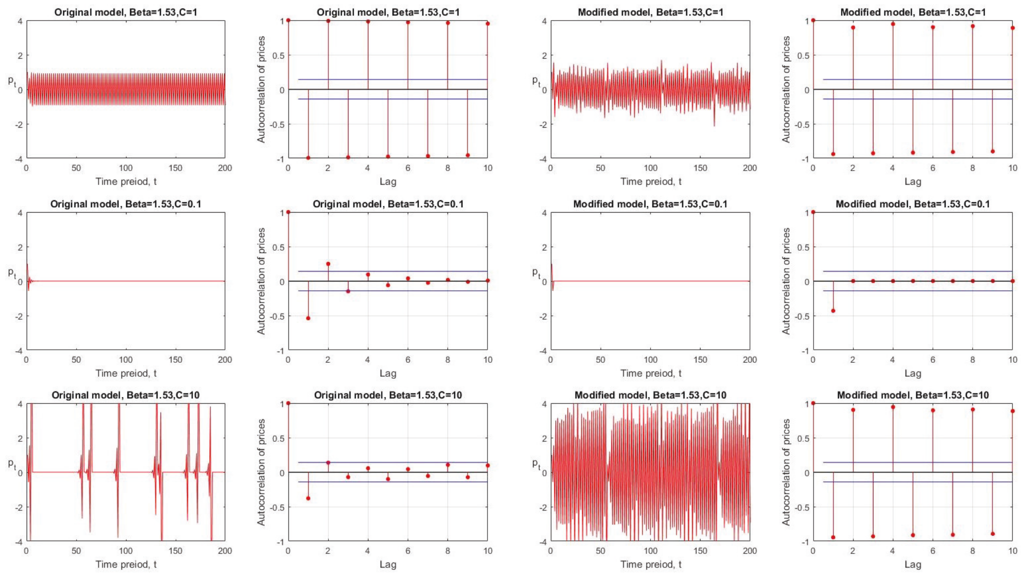

Figure B.2: Simulations of the price dynamics and corresponding sample autocorrelation. Parameters are fixed to $A = 0$, $B = 0.5$, $b = 1.35$ and $\beta = 1.53$. *Top row*: $C = 1$. *Middle row*: $C = 0.1$. *Bottom row*: $C = 10$. *Left panels*: Predictions of the standard model. *Right panels*: Predictions of the modified model.

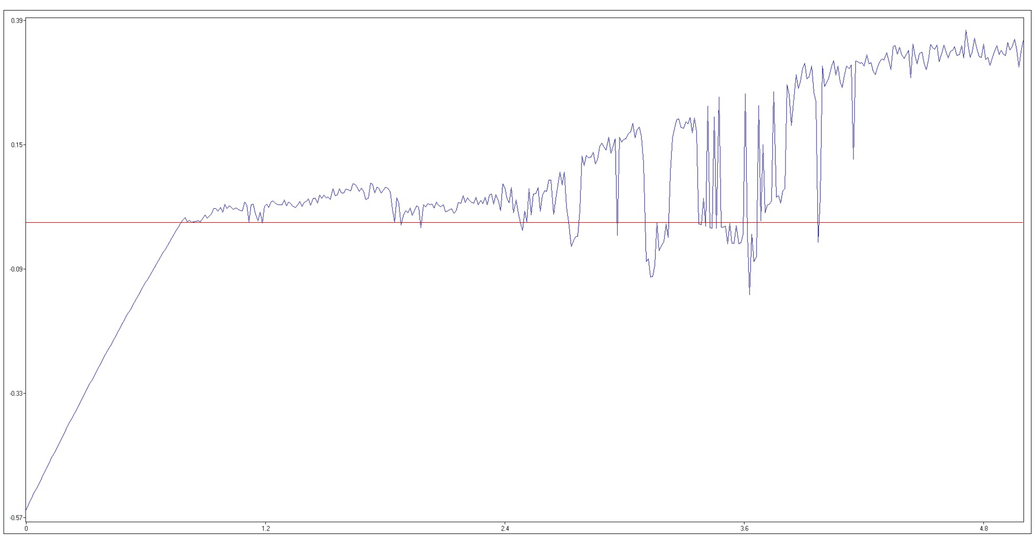

Figure B.3: Lyapunov exponent calculated for different values of the IoC parameter in the augmented model with self-tuning switching mechanism. Positive values indicate chaotic price dynamics for corresponding values of the IoC parameter.

ics in the bottom row graphs. Therefore, I can conclude that predictions with respect to relation between cost and stability are valid for the modified model as well.

Appendix C

Appendix to Chapter 4

C.1 Decomposition

$$MSE = \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{N}(\sum_{i=1}^{N}(p_{t,i}^{e} - p_t)^2)\right) = \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{N}(\sum_{i=1}^{N}(p_{t,i}^{e} - \bar{p_t^e} + \bar{p_t^e} - p_t)^2)\right) =$$

$$= \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{N}(\sum_{i=1}^{N}((p_{t,i}^{e} - \bar{p_t^e}) + (\bar{p_t^e} - p_t))^2)\right) =$$

$$= \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{N}(\sum_{i=1}^{N}(p_{t,i}^{e} - \bar{p_t^e})^2 + 2\frac{1}{N}\sum_{i=1}^{N}(p_{t,i}^{e} - \bar{p_t^e})(\bar{p_t^e} - p_t) + \frac{1}{N}(\sum_{i=1}^{N}(\bar{p_t^e} - p_t))^2)\right) =$$

$$= \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{N}(\sum_{i=1}^{N}(p_{t,i}^{e} - \bar{p_t^e})^2 + 2(\bar{p_t^e} - p_t)\frac{1}{N}\sum_{i=1}^{N}(p_{t,i}^{e} - \bar{p_t^e}) + (\bar{p_t^e} - p_t)^2)\right) =$$

$$= \frac{1}{T}\sum_{t=1}^{T}\left(\frac{1}{N}(\sum_{i=1}^{N}(p_{t,i}^{e} - \bar{p_t^e})^2)\right) + \frac{1}{T}\sum_{t=1}^{T}\left((\bar{p_t^e} - p_t)^2\right) =$$

$$= Dispersion + CommonError,$$

C.2 Decomposition Calculation

			All 50 forecasts					1-17 forecasts					18-34 forecasts					34-50 forecasts				
Price history	**Horizon length**	**Market**	**Individual Error**	**Dispersion**	**Fraction**	**Common Error**	**Fraction**	**Individual Error**	**Dispersion**	**Fraction**	**Common Error**	**Fraction**	**Individual Error**	**Dispersion**	**Fraction**	**Common Error**	**Fraction**	**Individual Error**	**Dispersion**	**Fraction**	**Common Error**	**Fraction**
Stable	H=1	S1G1	3.04	0.80	26%	2.24	74%	2.98	0.63	21%	2.35	79%	3.85	1.36	35%	2.48	65%	2.24	0.37	16%	1.88	84%
		S1G2	1.04	0.11	11%	0.93	89%	1.33	0.17	13%	1.16	87%	1.07	0.12	11%	0.95	89%	0.71	0.05	7%	0.66	93%
		S1G3	158.47	106.77	67%	51.69	33%	166.55	128.53	77%	38.02	23%	185.07	115.98	63%	69.09	37%	121.61	73.87	61%	47.73	39%
		S1G4	7.17	6.04	84%	1.13	16%	19.02	17.16	90%	1.86	10%	1.38	0.41	30%	0.96	70%	0.73	0.20	28%	0.53	72%
	H=2	S2G1	3.81	1.83	48%	1.98	52%	6.81	3.70	54%	3.11	46%	3.25	1.17	36%	2.08	64%	1.23	0.55	45%	0.68	55%
		S2G2	1498.64	1056.03	70%	442.61	30%	10.05	3.23	32%	6.82	68%	12.75	2.34	18%	10.41	82%	4659.03	3294.17	71%	1364.86	29%
		S2G3	8.28	5.50	66%	2.78	34%	22.66	15.44	68%	7.23	32%	1.12	0.49	44%	0.63	56%	0.62	0.27	44%	0.34	56%
		S2G4	2.48	1.07	43%	1.41	57%	3.94	1.43	36%	2.51	64%	1.36	0.45	33%	0.91	67%	2.12	1.35	64%	0.77	36%
	H=3	S3G1	18.03	10.73	60%	7.30	40%	42.75	26.43	62%	16.33	38%	7.71	3.47	45%	4.23	55%	2.72	1.76	65%	0.96	35%
		S3G2	13.17	7.18	55%	5.99	45%	33.72	19.06	57%	14.66	43%	2.90	1.49	51%	1.40	49%	2.25	0.61	27%	1.65	73%
		S3G3	1.80	0.76	42%	1.05	58%	4.43	1.79	40%	2.64	60%	0.59	0.30	51%	0.29	49%	0.30	0.14	48%	0.16	52%
Unstable	H=1	U1G1	1.28	0.54	42%	0.75	58%	1.87	1.17	63%	0.70	37%	0.86	0.11	13%	0.75	87%	1.10	0.32	29%	0.78	71%
		U1G2	3.47	2.31	67%	1.16	33%	6.64	5.56	84%	1.07	16%	1.69	0.61	36%	1.07	64%	1.99	0.66	33%	1.33	67%
		U1G3	251.67	44.64	18%	207.03	82%	11.46	3.66	32%	7.80	68%	250.19	34.75	14%	215.44	86%	508.47	98.69	19%	409.78	81%
		U1G4	23.94	5.36	22%	18.58	78%	2.86	0.88	31%	1.99	69%	9.24	1.68	18%	7.56	82%	61.95	14.02	23%	47.93	77%
		U1G5	1.32	0.69	52%	0.63	48%	2.00	1.22	61%	0.78	39%	0.71	0.08	11%	0.64	89%	1.26	0.79	63%	0.46	37%
		U1G6	423.07	98.27	23%	324.80	77%	40.37	5.75	14%	34.62	86%	310.19	67.71	22%	242.48	78%	949.64	229.05	24%	720.59	76%
		U1G7	13200.67	1644.70	12%	11555.96	88%	5.90	2.28	39%	3.61	61%	6638.38	924.65	14%	5713.74	86%	34192.54	4154.83	12%	30037.71	88%
		U1G8	373.91	37.54	10%	336.37	90%	2.67	1.50	56%	1.16	44%	6.98	1.04	15%	5.94	85%	1158.23	114.62	10%	1043.61	90%
	H=2	U2G1	127.29	39.64	31%	87.65	69%	188.08	76.25	41%	111.84	59%	138.79	26.32	19%	112.47	81%	50.46	14.89	30%	35.56	70%
		U2G2	684.72	53.76	8%	630.96	92%	27.73	8.62	31%	19.11	69%	621.54	43.39	7%	578.15	93%	1449.90	112.74	8%	1337.15	92%
		U2G3	410.45	50.60	12%	359.85	88%	136.97	15.29	11%	121.69	89%	696.52	56.24	8%	640.27	92%	397.09	82.12	21%	314.96	79%
		U2G4	331.19	41.35	12%	289.84	88%	163.23	20.92	13%	142.31	87%	436.81	29.27	7%	407.54	93%	397.43	75.89	19%	321.54	81%
		U2G5	466.10	117.50	25%	348.60	75%	520.33	116.89	22%	403.45	78%	519.19	110.16	21%	409.04	79%	352.07	125.97	36%	226.11	64%
		U2G6	555.16	180.24	32%	374.93	68%	94.79	59.46	63%	35.33	37%	976.48	282.54	29%	693.94	71%	596.66	199.87	33%	396.80	67%
	H=3	U3G1	35.98	23.97	67%	12.01	33%	51.78	30.70	59%	21.08	41%	49.34	36.32	74%	13.02	26%	4.99	3.70	74%	1.29	26%
		U3G2	1410.13	998.19	71%	411.94	29%	68.23	22.46	33%	45.78	67%	4062.27	2904.23	71%	1158.04	29%	18.01	9.73	54%	8.27	46%
		U3G3	23.63	5.75	24%	17.88	76%	30.97	11.61	37%	19.36	63%	25.19	2.85	11%	22.34	89%	14.16	2.59	18%	11.57	82%
		U3G4	472.73	111.42	24%	361.30	76%	985.95	163.91	17%	822.04	83%	322.30	111.89	35%	210.41	65%	87.25	55.16	63%	32.09	37%
		U3G5	87.04	23.43	27%	63.61	73%	225.08	50.67	23%	174.41	77%	27.15	15.40	57%	11.75	43%	4.01	3.03	75%	0.98	25%
		U3G6	20.02	2.86	14%	17.16	86%	43.67	5.01	11%	38.65	89%	12.67	2.22	18%	10.45	82%	2.69	1.24	46%	1.46	54%
Unstable with strong feedback	H=3	SU3G1	11122.26	9749.27	88%	1372.99	12%	1.63	0.62	38%	1.01	62%	51.85	26.40	51%	25.46	49%	34700.23	30437.77	88%	4262.46	12%
		SU3G2	4.04	1.21	30%	2.83	70%	8.40	2.24	27%	6.16	73%	2.73	1.00	37%	1.73	63%	0.81	0.34	42%	0.47	58%
		SU3G3	677.85	102.41	15%	575.44	85%	652.51	32.17	5%	620.34	95%	1247.12	222.81	18%	1024.31	82%	99.92	49.11	49%	50.81	51%
		SU3G4	21.34	12.93	61%	8.41	39%	51.42	30.51	59%	20.91	41%	7.23	4.09	57%	3.13	43%	4.39	3.64	83%	0.75	17%
		SU3G5	1744.78	919.55	53%	825.23	47%	979.60	78.18	8%	901.42	92%	455.58	177.25	39%	278.33	61%	3927.56	2602.19	66%	1325.37	34%
		SU3G6	633.61	73.58	12%	560.03	88%	653.38	16.67	3%	636.71	97%	1049.05	115.04	11%	934.01	89%	171.22	9	53%	81.21	47%

Table C.1: Decomposition of forecasting errors calculated across markets and forecasting periods within 3 subsamples.

C.3 Individual forecasts

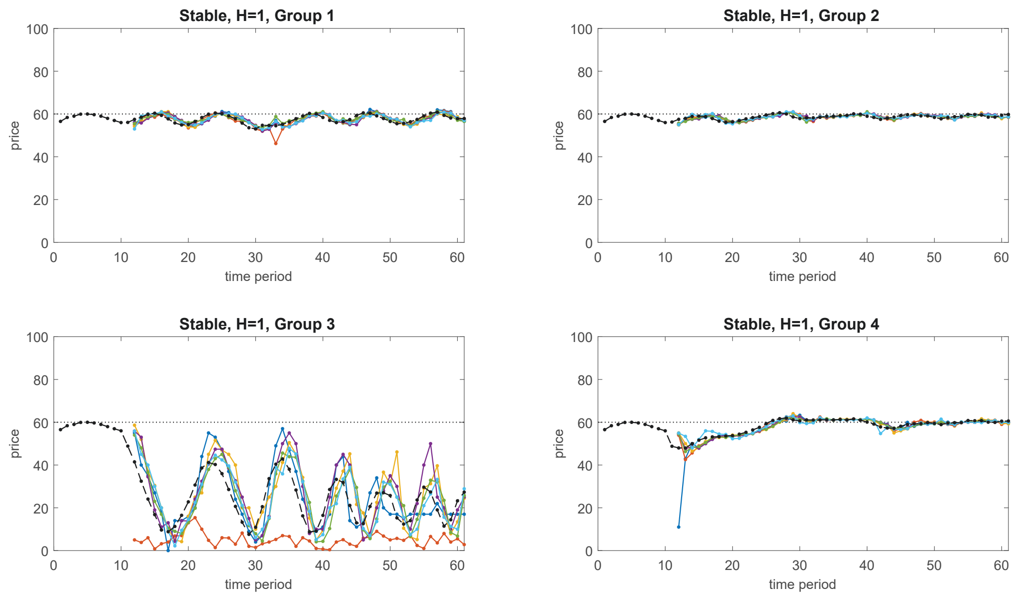

Figure C.1: Forecasts and price dynamics in 4 groups of 1-period ahead forecasting with stable price history.

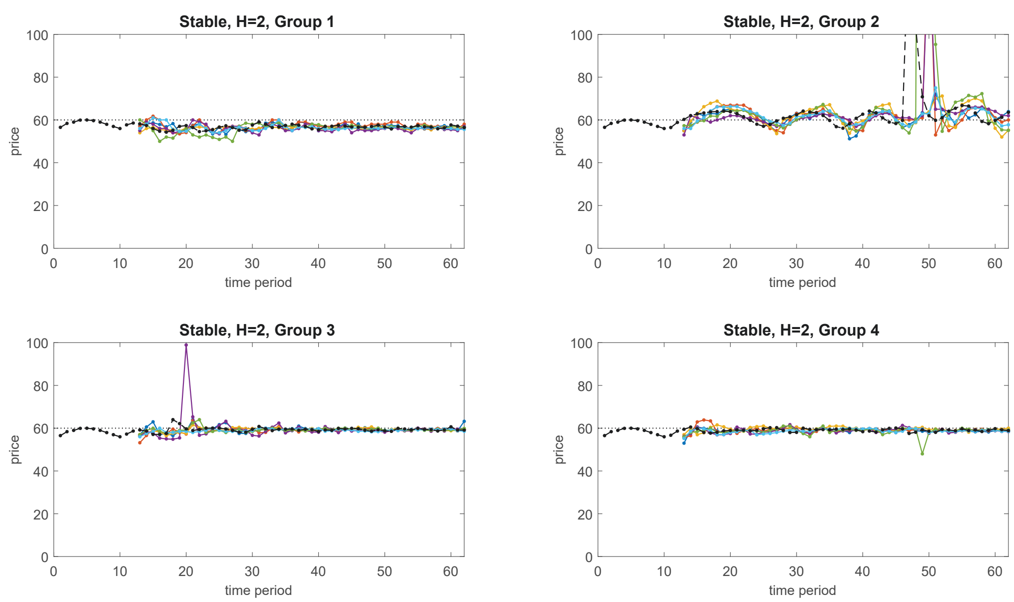

Figure C.2: Forecasts and price dynamics in 4 groups of 2-periods ahead forecasting with stable price history.

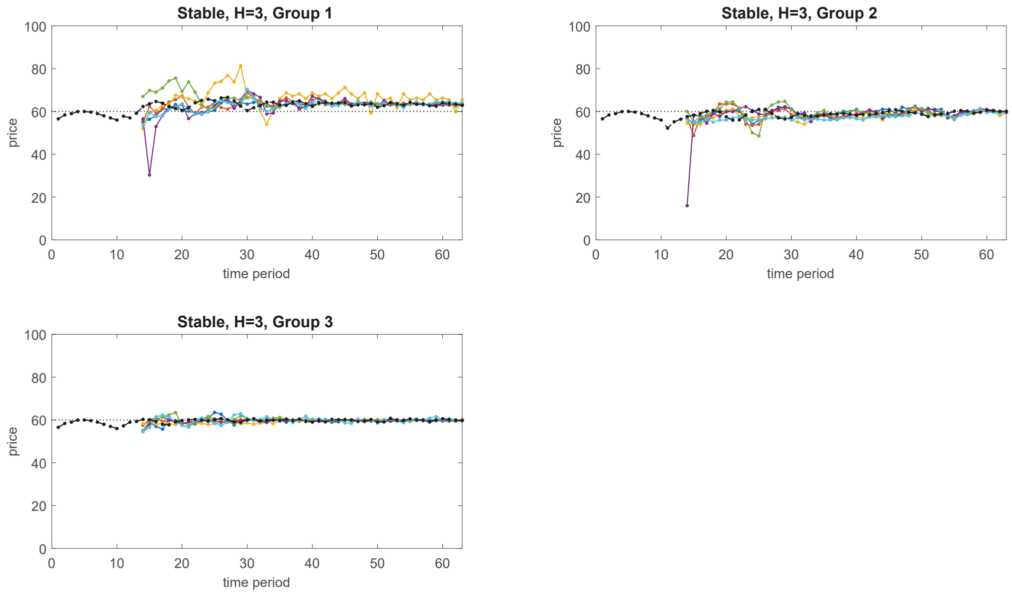

Figure C.3: Forecasts and price dynamics in 3 groups of 3-periods ahead forecasting with stable price history.

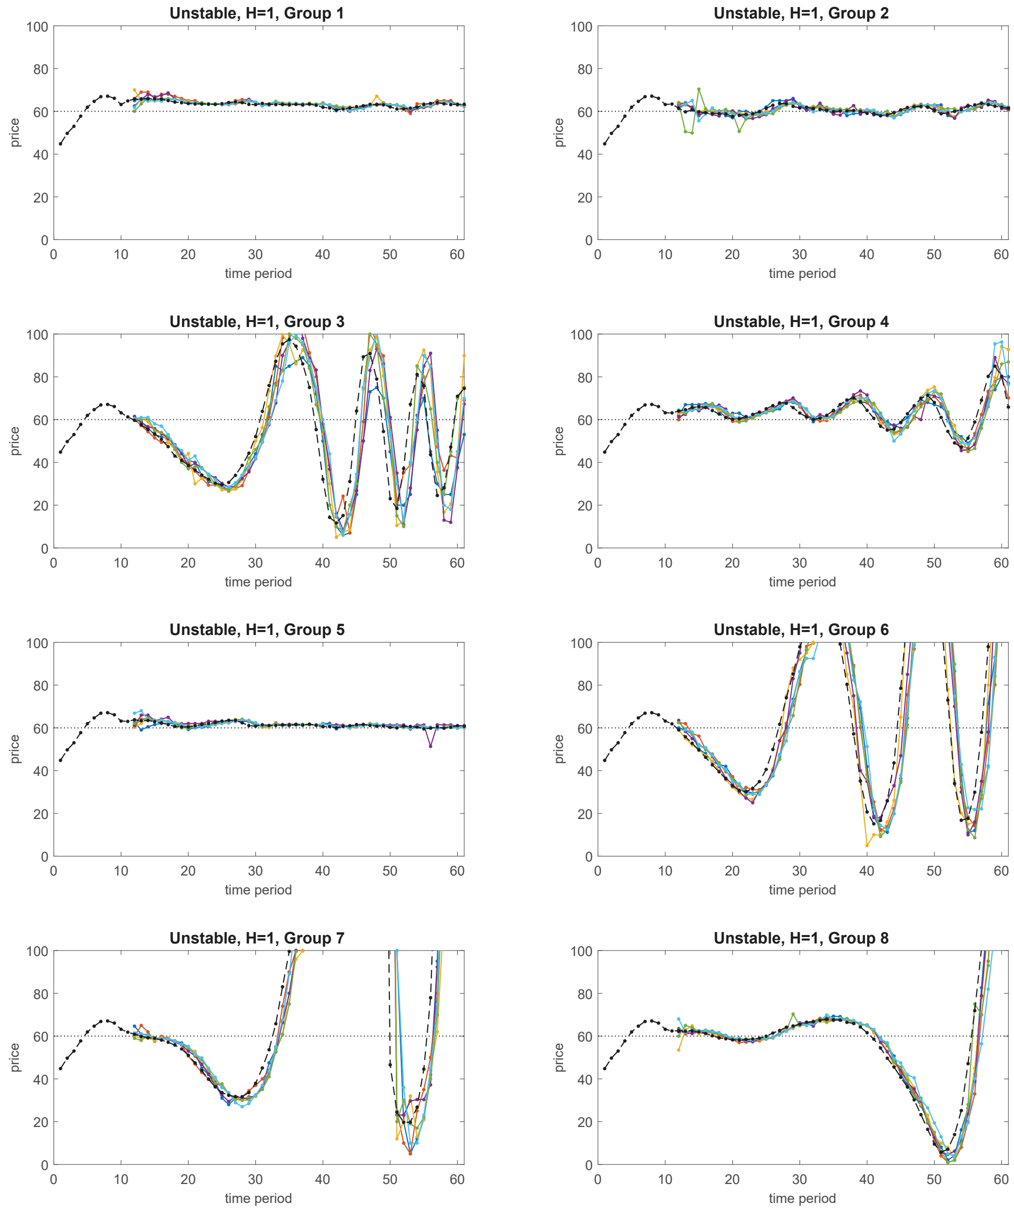

Figure C.4: Forecasts and price dynamics in 8 groups of 1-period ahead forecasting with unstable price history.

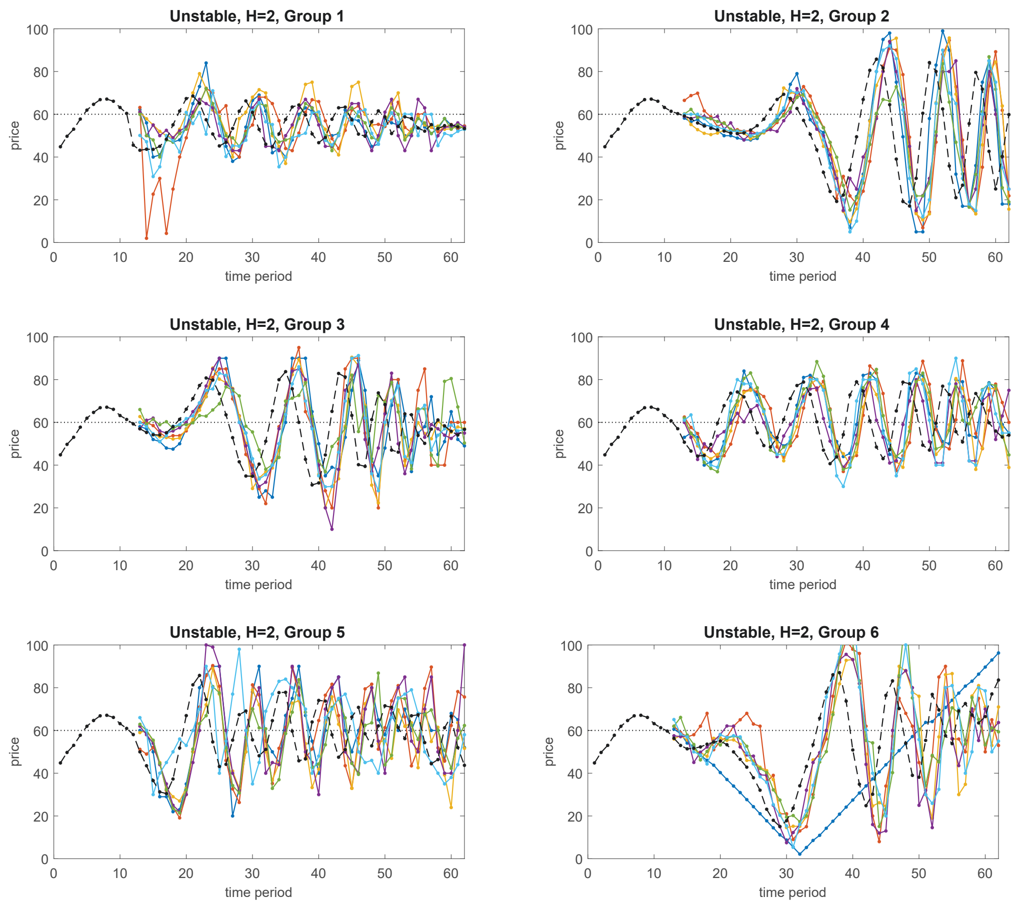

Figure C.5: Forecasts and price dynamics in 6 groups of 2-periods ahead forecasting with unstable price history.

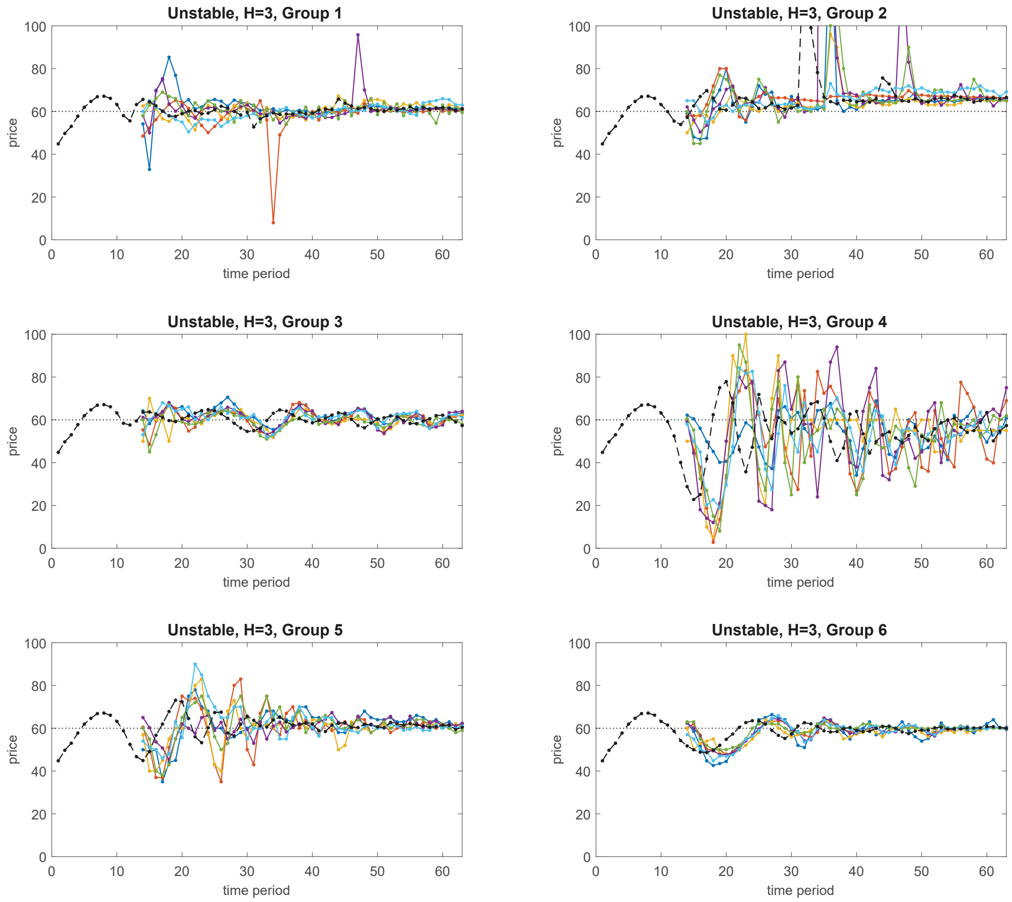

Figure C.6: Forecasts and price dynamics in 6 groups of 3-periods ahead forecasting with unstable price history.

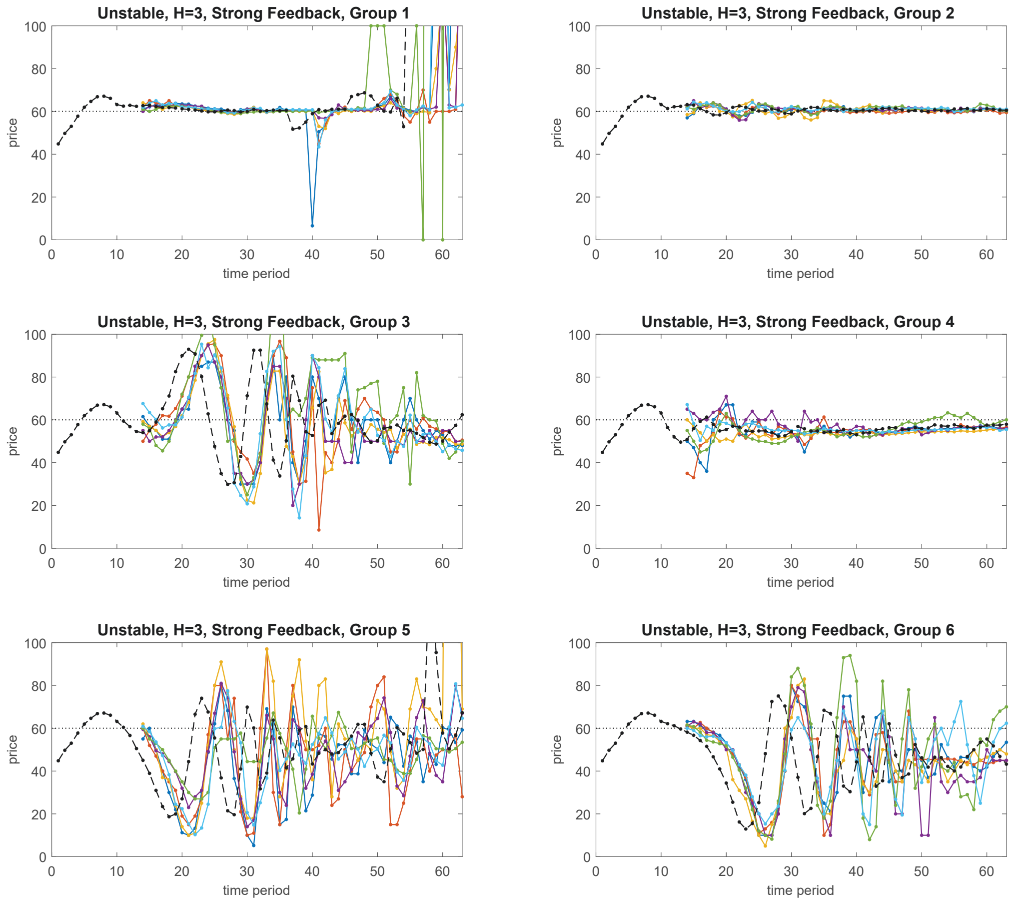

Figure C.7: Forecasts and price dynamics in 6 groups of 3-periods ahead forecasting with unstable price history and strong feedback.

C.4 Non-parametric tests

	S1	S2	S3	U1	U2	U3
S1		0.40	0.40	0.38	0.02	0.14
S2	0.40		1.00	0.13	0.02	0.04
S3	0.40	1.00		0.13	0.02	0.04
U1	0.38	0.13	0.13		0.95	0.62
U2	0.02	0.02	0.02	0.95		0.05
U3	0.14	0.04	0.04	0.62	0.05	

Table C.2: P-values of a two-sided Wilcoxon rank sum test of the null hypothesis of equal medians of average prices in sessions of different treatments are reported.

	S1	S2	S3	U1	U2	U3
S1		0.40	0.40	0.28	0.02	0.25
S2	0.40		0.70	0.02	0.02	0.04
S3	0.40	0.70		0.08	0.02	0.04
U1	0.28	0.02	0.08		0.95	0.62
U2	0.02	0.02	0.02	0.95		0.01
U3	0.25	0.04	0.04	0.62	0.01	

Table C.3: P-values of a two-sided Wilcoxon rank sum test of the null hypothesis of equal medians of variances of prices in sessions of different treatments are reported.

	S1	S2	S3	U1	U2	U3
S1		0.70	0.70	0.19	0.02	0.79
S2	0.70		1.00	0.08	0.02	0.39
S3	0.70	1.00		0.08	0.02	0.39
U1	0.19	0.08	0.08		1.00	0.28
U2	0.02	0.02	0.02	1.00		0.01
U3	0.79	0.39	0.39	0.28	0.01	

Table C.4: P-values of a two-sided Wilcoxon rank sum test of the null hypothesis of equal medians of RAD of prices in sessions of different treatments are reported.

	S1	S2	S3	U1	U2	U3
S1		0.70	0.10	0.08	0.55	0.39
S2	0.70		0.40	0.08	0.38	0.57
S3	0.10	0.40		0.28	0.10	0.79
U1	0.08	0.08	0.28		0.01	0.07
U2	0.55	0.38	0.10	0.01		0.08
U3	0.39	0.57	0.79	0.07	0.08	

Table C.5: P-values of a two-sided Wilcoxon rank sum test of the null hypothesis of equal medians of RD of prices in sessions of different treatments are reported.

	S1	S2	S3	U1	U2	U3
S1		0.40	0.40	0.38	0.02	0.14
S2	0.40		1.00	0.13	0.02	0.04
S3	0.40	1.00		0.13	0.02	0.04
U1	0.38	0.13	0.13		0.95	0.62
U2	0.02	0.02	0.02	0.95		0.05
U3	0.14	0.04	0.04	0.62	0.05	

Table C.6: P-values of a two-sided Wilcoxon rank sum test of the null hypothesis of equal medians of PA in sessions of different treatments are reported.

C.5 Individual forecasting strategies

Market	Participant	α	β_1	β_2	β_3	β_4	γ_1	γ_2	γ_3	γ_4	Autocorrelation
'S1G1'	1		2.14	-1.67							
'S1G1'	2		1.52								
'S1G1'	3		1.58	-0.96							
'S1G1'	4	0.92	1.70	-0.90							
'S1G1'	5		1.83	-2.05				0.30	0.26		
'S1G1'	6		1.82	-0.99		0.45				-0.29	
'S1G2'	1		1.49		-0.84			0.36			
'S1G2'	2		1.54								
'S1G2'	3		1.78	-0.85							
'S1G2'	4		1.22				-0.35				
'S1G2'	5		1.60	-0.94							
'S1G2'	6		1.56	-0.77			0.35	-0.48			
'S1G3'	1		1.45				0.74	-0.54			
'S1G3'	2						0.51				
'S1G3'	3		2.22								
'S1G3'	4	8.85	1.70								
'S1G3'	5		1.95	-0.85							
'S1G3'	6		2.22				-0.44				
'S1G4'	1		1.23								
'S1G4'	2		1.96	-1.38							
'S1G4'	3		1.53								
'S1G4'	4		1.40	-0.67							
'S1G4'	5		1.36								
'S1G4'	6		2.27	-1.11	-1.19	0.71					
'S2G1'	1		1.24			0.92					
'S2G1'	2		1.17	-1.07			0.64	-0.48			
'S2G1'	3					0.32	0.46				
'S2G1'	4		1.41	-1.63							
'S2G1'	5		0.67				0.67				
'S2G1'	6		0.80				0.45				
'S2G2'	1			0.11	-0.12		1.00			0.30	
'S2G2'	2		1.02	-1.63	0.72		1.08		-0.19		
'S2G2'	3						1.14				
'S2G2'	4		0.90	-1.18	0.49		0.89				
'S2G2'	5										
'S2G2'	6	4.62		0.15	-0.26	0.17	1.55	-0.57	-0.40	0.31	
'S2G3'	1		0.95			0.79			-0.53		
'S2G3'	2		0.78						0.33		
'S2G3'	3		0.67	-0.39	0.32		0.35				
'S2G3'	4					2.90					
'S2G3'	5		0.54		-0.43	0.33	0.57				
'S2G3'	6						0.60				
'S2G4'	1		0.80	-0.44							
'S2G4'	2		0.98	-0.47							5
'S2G4'	3		0.69				0.40		-0.33		
'S2G4'	4		0.80	-0.56		0.61	0.39				
'S2G4'	5		1.48								
'S2G4'	6		0.43				0.79				
'S3G1'	1		0.32		-0.42		0.63				
'S3G1'	2		0.94	-0.71			0.52	-0.40			
'S3G1'	3										
'S3G1'	4		1.39	-1.09							
'S3G1'	5		0.97				0.39		0.32		
'S3G1'	6		1.43	-1.08						0.33	
'S3G2'	1		1.66	-1.54			0.35				
'S3G2'	2		0.93	-0.95			0.69		-0.32	-0.32	
'S3G2'	3		1.07	-1.39			0.81				
'S3G2'	4		1.04								
'S3G2'	5		2.12	-2.07			0.52				
'S3G2'	6		0.43				0.50				
'S3G3'	1		1.01				0.44				
'S3G3'	2		0.21	-0.28		0.24	0.84	0.46	-0.27		1
'S3G3'	3										
'S3G3'	4		0.29			0.26	0.58	0.25	-0.44	0.30	1
'S3G3'	5		0.99					-0.39			
'S3G3'	6										

Table C.7: Estimates for individual forecasting strategies in markets with initially stable price histories. Only significant coefficients are reported.

Market	Participant	α	β_1	β_2	β_3	β_4	γ_1	γ_2	γ_3	γ_4	Autocorrelation
'U1G1'	1		1.74	-0.85							
'U1G1'	2		2.15	-1.74			0.33				
'U1G1'	3		1.42								
'U1G1'	4		1.51	-0.90			0.37			-0.18	
'U1G1'	5		1.15	-0.54	-0.75	0.49	0.42	0.34			
'U1G1'	6		1.27	-0.66			0.45				
'U1G2'	1		1.77	-1.45			0.61				
'U1G2'	2		1.26								
'U1G2'	3		1.34	-0.57			0.39			0.24	
'U1G2'	4		1.98								
'U1G2'	5		1.17								
'U1G2'	6		1.41				0.28	-0.31			
'U1G3'	1	8.30	1.05	-1.38			0.89				
'U1G3'	2		2.59	-2.96	1.64						
'U1G3'	3		1.96							0.70	
'U1G3'	4		1.81	-1.11							
'U1G3'	5	11.37	2.17	-3.15	2.90	-0.93			-0.77		
'U1G3'	6	7.99	1.72	-1.14							
'U1G4'	1	13.03	1.04		-2.40	1.36					
'U1G4'	2	6.02	1.04	-1.35		0.84	0.95				
'U1G4'	3		2.20								
'U1G4'	4		2.61								
'U1G4'	5	4.66	1.70	-1.05							
'U1G4'	6		3.87	-4.09	2.84			-0.46			
'U1G5'	1		1.20								
'U1G5'	2		1.32								
'U1G5'	3	1.36	1.07	-0.58			0.49				
'U1G5'	4		1.41			1.29					
'U1G5'	5		1.42	-0.83	0.60		0.47	-0.51			
'U1G5'	6		1.00								
'U1G6'	1		2.21	-2.35					-0.47		
'U1G6'	2		2.04					-0.48		-0.72	
'U1G6'	3		3.32	-3.51				-0.74			
'U1G6'	4		2.69	-2.53							
'U1G6'	5		2.25				-0.71				
'U1G6'	6		2.17	-1.46				-0.47			
'U1G7'	1		2.32	-1.55	1.92				-1.07		
'U1G7'	2										
'U1G7'	3		2.37				-1.08	-0.74			
'U1G7'	4	34.23			-2.40		3.20	-3.19	2.65		
'U1G7'	5		2.21	-1.04	2.16		-0.71		-1.22		
'U1G7'	6						1.71	-1.57			
'U1G8'	1	13.90	1.75	-1.53		-2.13	0.58	-0.78	0.66	0.74	
'U1G8'	2	-12.42	2.36	-1.74	1.79				-1.35		
'U1G8'	3	24.28				-2.86			2.18		
'U1G8'	4	-6.63	2.74		-1.51	1.60	-0.43	-0.86		-0.92	
'U1G8'	5	-23.88	5.09			3.10	-0.83	-1.58		-1.67	
'U1G8'	6		2.18			4.43	-0.98	-1.38	-1.44	-1.83	

Table C.8: Estimates for individual forecasting strategies in markets with initially unstable price histories and one-period ahead horizons. Only significant coefficients are reported.

Market	**Participant**	α	β_1	β_2	β_3	β_4	γ_1	γ_2	γ_3	γ_4	**Autocorrelation**
'U2G1'	1		1.75	-2.03	1.21						
'U2G1'	2		1.70	-1.53			0.32		0.38		
'U2G1'	3		2.03	-1.49							
'U2G1'	4		2.23	-1.75		0.64			0.39		
'U2G1'	5		1.86	-1.40							
'U2G1'	6		0.93								
'U2G2'	1	19.95	2.71	-4.45	4.78	-1.39	0.64	-1.01		-0.85	
'U2G2'	2		2.02	-1.09							
'U2G2'	3		2.11	-1.47							
'U2G2'	4	20.88	1.92	-2.36	1.80	-1.35				0.60	
'U2G2'	5		2.51	-2.98	1.29						
'U2G2'	6		3.65	-5.46	3.06			0.76			
'U2G3'	1		1.59	-1.26				0.41			
'U2G3'	2		2.71	-3.57	2.39						
'U2G3'	3		2.98	-2.16	2.09	-1.01					
'U2G3'	4		2.26	-2.21	1.70	-0.70		-0.66			
'U2G3'	5		1.72	-1.76							
'U2G3'	6	10.56	2.55	-3.24	1.44			0.68			
'U2G4'	1		1.88	-1.97							
'U2G4'	2		1.53								
'U2G4'	3		1.70	-1.00							
'U2G4'	4	69.27	2.73	-2.29			-0.57	-1.00	0.57		
'U2G4'	5		1.65								
'U2G4'	6		3.31	-2.20				-0.97			
'U2G5'	1		1.98	-1.78	1.82						
'U2G5'	2		2.02	-2.70	2.14		0.52				
'U2G5'	3		2.33	-1.58			-0.46				
'U2G5'	4		3.68	-2.50			-0.67				
'U2G5'	5		1.66	-0.91							
'U2G5'	6						0.98	-0.55			
'U2G6'	1	1.37					1.53				
'U2G6'	2		1.51				0.45				
'U2G6'	3		2.60	-1.57				-0.36	0.43		
'U2G6'	4		3.08	-3.06							
'U2G6'	5		2.02	-1.74	1.33						
'U2G6'	6	11.97	2.09	-1.41							
'U3G1'	1		0.58				0.74		-0.24		5.00
'U3G1'	2										
'U3G1'	3		0.50	-0.69			0.58			0.32	
'U3G1'	4										
'U3G1'	5		0.65					0.33			
'U3G1'	6						1.00				
'U3G2'	1		1.09	-0.78		0.33	0.36				
'U3G2'	2						1.55	-0.97			
'U3G2'	3		0.40		-0.27		0.67				
'U3G2'	4										
'U3G2'	5		0.50				0.52				
'U3G2'	6		0.13	-0.09			0.83				
'U3G3'	1		1.56	-1.42			0.52				
'U3G3'	2		2.17	-1.83			0.56	-0.56			
'U3G3'	3						0.35		-0.39	0.30	1.00
'U3G3'	4		2.52	-1.66							
'U3G3'	5		1.95	-2.34	1.30		0.43				
'U3G3'	6		1.78	-1.83							
'U3G4'	1		0.72	-0.72			0.43				
'U3G4'	2		1.48	-1.49					-0.35		
'U3G4'	3	44.42									
'U3G4'	4		1.92	-1.88							
'U3G4'	5		1.88	-1.86	2.04						
'U3G4'	6		1.46								
'U3G5'	1		0.57				0.65		-0.44		
'U3G5'	2		2.65	-2.23							
'U3G5'	3		2.59			-0.90		-0.53			
'U3G5'	4		0.85						0.30		
'U3G5'	5		1.76	-1.63							
'U3G5'	6		1.54	-1.57			0.43		0.46		
'U3G6'	1		2.75	-3.00	1.74		0.38				
'U3G6'	2		1.62	-1.02							
'U3G6'	3		1.86	-1.90			0.36		-0.34		
'U3G6'	4		1.73								
'U3G6'	5		0.80	-1.12			1.03				
'U3G6'	6		1.66								

Table C.9: Estimates for individual forecasting strategies in markets with initially unstable price histories and both two- and three-period ahead horizons. Only significant coefficients are reported.

Market	Participant	α	β_1	β_2	β_3	β_4	γ_1	γ_2	γ_3	γ_4	Autocorrelation
'WU3G1'	1		0.71	0.40	-0.48	0.53	-0.37				0
'WU3G1'	2			0.31	-0.45	0.68	0.46	0.48	-0.41		5
'WU3G1'	3	-154.08	0.22	0.21	-0.46	-0.66		2.32	2.43		4
'WU3G1'	4	-33.98	0.08	0.32		-0.15	-0.77		1.43	0.77	3
'WU3G1'	5			-23.86				2.85			0
'WU3G1'	6	83.16	1.06	0.47	-0.75		-0.35	0.19	-0.73	-0.35	1
'WU3G2'	1		1.37							0.06	0
'WU3G2'	2		1.07	-0.54			0.38			0.38	0
'WU3G2'	3		1.21	-1.31			0.50		-0.36	0.36	0
'WU3G2'	4		1.19	-0.95			0.69	-0.41			0
'WU3G2'	5		1.24	-1.03			0.66				0
'WU3G2'	6					0.53	0.65				0
'WU3G3'	1		1.71	-1.22							0
'WU3G3'	2		1.55	-1.73				0.53			0
'WU3G3'	3		1.62	-1.47					0.43	-0.47	0
'WU3G3'	4		1.03	-0.97			0.60				0
'WU3G3'	5		1.94	-2.02							0
'WU3G3'	6		1.44	-1.70							0
'WU3G4'	1		2.09	-1.28		0.83				-0.31	4
'WU3G4'	2		1.85	-1.03							0
'WU3G4'	3						0.49				0
'WU3G4'	4										0
'WU3G4'	5				-1.05	0.75	1.45	-0.57			0
'WU3G4'	6		0.47				0.55				4
'WU3G5'	1						0.90	-0.87	0.51	-0.50	0
'WU3G5'	2	21.94					0.55				0
'WU3G5'	3										0
'WU3G5'	4		0.32	-0.38			1.00	-0.56			0
'WU3G5'	5						0.99	-0.57			0
'WU3G5'	6		0.53				0.55				0
'WU3G6'	1		1.26					-0.51			0
'WU3G6'	2		1.31					-0.59			0
'WU3G6'	3	13.88					0.94				0
'WU3G6'	4		1.27						0.50		1
'WU3G6'	5		1.20								0
'WU3G6'	6		1.09	-1.11							0

Table C.10: Estimates for individual forecasting strategies in markets from extra sessions. Only significant coefficients are reported.

C.6 Data from previous LtF experiments

Table C.11: Descriptive statistics of price observed in previous 0-period ahead and 1-periods ahead LtF experiments.

Data	Mean price'	'Std Dev'
'1 period ahead. Group 1'	54.79	28.41
'1 period ahead. Group 2'	56.09	23.33
'1 period ahead. Group 3'	56.37	24.26
'1 period ahead. Group 4'	58.67	4.56
'0 periods ahead. Group 1'	57.15	9.36
'0 periods ahead. Group 2'	58.97	6.53
'0 periods ahead. Group 3'	58.58	4.63
'0 periods ahead. Group 4'	59.14	13.07
'0 periods ahead. Group 5'	94.23	142.09
'0 periods ahead. Group 6'	58.89	7.76
'0 periods ahead. Group 7'	62.13	0.87

We can observe that, generally, standard deviation is substantially higher in 1-periods ahead forecasting markets with two notable exceptions. Group 4 in 1-periods ahead markets has no bubbles and crushes patterns and Group

5 has outlier previously attributed to the participant?s possible typo. The rest of observation suggests that in positive feedback markets increase of forecasting horizon leads to higher volatility of the market price.

The crucial difference between 0-period and 1-period ahead treatments lies in price setting mechanism. Market maker in one case and market clearing in other could be a potential source of identified differences.

Instructions

General information

Today you will participate in an experiment which will require you to predict the future price of a risky asset. During the experiment you will be able to earn a number of points. The better your predictions are, the more points will you earn. These points will be converted into Australian dollars after the experiment.

Information about your task

You are a **financial forecaster** working for a pension fund that wants to optimally invest a large amount of money for 3 periods. The pension fund has two investment options: a risk-free investment and a risky investment. The risk-free investment is putting money in **a savings account**, which pays a fixed and constant interest rate over 3 periods. The alternative for the pension fund is to invest its money in **a risky asset**, where risk comes from the uncertain future price of that asset.

In each period the pension fund has to decide which fraction of its money to put in the savings account and which fraction of its money to invest in the risky asset. To make the optimal investment decision, the pension fund needs an accurate prediction of the future price of the asset. The pension fund is only interested in the price of the risky asset after 3 periods.

As the **financial forecaster** of the fund, you have to predict the price for the risky asset 3 periods ahead during 53 subsequent periods. Your earnings during the experiment depend upon the accuracy of your predictions. **The smaller your errors in each period are, the higher your total earnings will be.**

Information about the asset market

The market price of the risky asset in each period is determined by demand and supply. The total supply of assets is fixed during the experiment. The demand for assets is mainly determined by the aggregate demand of several large pension funds active in the asset market. There is also some uncertain, small demand for assets by private investors but the effect of private investors upon the asset price is small.

Information about the investment strategies of the pension funds

The precise investment strategy of the pension fund that you are advising and the investment strategies of the other pension funds are unknown. The savings account, that provides the risk-free investment, pays a fixed interest rate of 5% per period. The owner of the risky asset receives an uncertain payment in each period, but economic experts have computed that this payment is 3 dollars per period on average. The return of the asset market per period depends upon these payments as well as upon price changes of the asset.

As the **financial forecaster** of a pension fund you are only asked to predict, in each period, the 3 periods ahead price of the asset. Based upon your future price predictions, your pension fund will make an optimal investment decision and hold the asset for 3 periods. The higher your predicted future price is, the larger will be the fraction of money invested by your pension fund in the asset market in the current period, so the larger will be its demand for assets.

Information during the experiment

At the beginning of the experiment, you have the history of the asset price in the first 10 periods, and you start in period 11 by giving your prediction of the price in period 14. After all participants have given their predictions, the realized asset price for period 11 will be revealed. Then you (as all other participants) will need to make a new prediction, now for the price in period 15, so that the asset price for period 12 can be defined. And so on. This process continues until period 63, where the last prediction, for the price in period 66, will be given.

To predict the asset price for period $t+3$ in period t, the available information consists of

- past prices up to period $t-1$,
- your previous predictions up to period $t+2$,
- your past earnings up to period $t-1$.

Starting from period 14, your earnings in each period will be based upon your prediction error, that is, the difference between the price you predicted for that period and the realized price in that period. The last period for which you will be paid is period 63.

The better you predict the asset price in each period, the higher your aggregate earnings will be. Earnings for each period in points will be automatically computed according to the following earnings table, where "error" denotes the absolute value of the difference between your prediction and price in that period. Information on your earnings in the current period and cumulative earnings will be reported to you during the experiment.

After the experiment your earned points will be converted into Australian dollars, with 1300 points equal to 50 cents. You will be paid the sum of show-up fee and all your earnings in AUD.

Earnings table									
Error	**Points**	**Error**	**Points**	**Error**	**Points**	**Error**	**Points**	**Error**	**Points**
0.1	1300	**1.5**	1240	**2.9**	1077	**4.3**	809	**5.7**	438
0.2	1299	**1.6**	1232	**3**	1061	**4.4**	786	**5.8**	408
0.3	1298	**1.7**	1223	**3.1**	1045	**4.5**	763	**5.9**	376
0.4	1296	**1.8**	1214	**3.2**	1028	**4.6**	739	**6**	345
0.5	1293	**1.9**	1204	**3.3**	1011	**4.7**	714	**6.1**	313
0.6	1290	**2**	1194	**3.4**	993	**4.8**	689	**6.2**	280
0.7	1287	**2.1**	1183	**3.5**	975	**4.9**	663	**6.3**	247
0.8	1283	**2.2**	1172	**3.6**	956	**5**	637	**6.4**	213
0.9	1279	**2.3**	1160	**3.7**	937	**5.1**	610	**6.5**	179
1	1273	**2.4**	1147	**3.8**	917	**5.2**	583	**6.6**	144
1.1	1268	**2.5**	1134	**3.9**	896	**5.3**	555	**6.7**	109
1.2	1262	**2.6**	1121	**4**	876	**5.4**	526	**6.8**	73
1.3	1255	**2.7**	1107	**4.1**	854	**5.5**	497	**6.9**	37
1.4	1248	**2.8**	1092	**4.2**	832	**5.6**	468	**7 or more**	0

Additional information

- By the end of the experiment, you will be paid privately. Before the payment you will be asked to answer a questionnaire. Inserted data will be processed in nameless form only. Please fill in the correct information.

- During the experiment any communication with other participants, whether verbal or written, is forbidden. The use of phones, tablets or any other gadgets is not allowed. Violation of the rules can result in exclusion from the experiment without any remuneration.

- Please follow the instructions carefully at all the stages of the experiment. If you have any questions or encounter any problems during the experiment, please raise your hand and the experimenter will come to help you.

Please ask any questions you have now!

Bibliography

[1] Helen Allen and Mark P. Taylor. Charts, noise and fundamentals in the london foreign exchange market. *The Economic Journal*, 100:49–59, 1990.

[2] Mikhail Anufriev, Te Bao, and Jan Tuinstra. Microfoundations for switching behavior in heterogeneous agent models: An experiment. *Journal of Economic Behavior and Organization*, 129:74–99, 2016.

[3] Mikhail Anufriev, Aleksei Chernulich, and Jan Tuinstra. A laboratory experiment on the heuristic switching model. *Journal of Economic Dynamics & Control*, 91:21–42, 2018.

[4] Mikhail Anufriev and Cars Hommes. Evolution of market heuristics. *The Knowledge Engineering Review*, 27(2):255–271, 2012.

[5] Mikhail Anufriev and Cars Hommes. Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. *American Economic Journal: Microeconomics*, 4(4):35—64, 2012.

[6] Mikhail Anufriev, Cars Hommes, and Tomasz Makarewicz. Simple forecasting heuristics that make us smart: Evidence from different market experiments. *Journal of the European Economic Association*, 2018.

[7] Mikhail Anufriev, Cars Hommes, and Raoul H.S. Philipse. Evolutionary selection of expectations in positive and negaive feedback markets. *Journal of Evolutionary Economics*, 23:663–688, 2013.

[8] Mikhail Anufriev and Valentyn Panchenko. Asset prices, traders' behavior and market design. *Journal of Economic Dynamics and Control*, 33(5):1073–1090, 2009.

[9] Mikhail Anufriev and Jan Tuinstra. The impact of short-selling constraints on financial market stability in a heterogeneous agents model. *Journal of Economic Dynamics and Control*, 37(8):1523–1543, 2013.

[10] J Arifovic, CH Hommes, Isabelle Salle, et al. Learning to believe in simple equilibria in a complex olg economy-evidence from the lab.

Technical report, Universiteit van Amsterdam, Center for Nonlinear Dynamics in Economics and Finance, 2016.

[11] Jasmina Arifovic and John Ledyard. Scaling up learning models in public good games. *Journal of Public Economic Theory*, 6(2):203–238, 2004.

[12] Te Bao, Cars Hommes, and Tomasz Makarewicz. Bubble formation and (in) efficient markets in learning-to-forecast and optimise experiments. *The Economic Journal*, 127(605), 2017.

[13] H. Peter Boswijk, Cars H. Hommes, and Sebastiano Manzan. Behavioral heterogeneity in stock prices. *Journal of Economic Dynamics and Control*, 31:1938—-1970, 2007.

[14] H Peter Boswijk, Cars H Hommes, and Sebastiano Manzan. Behavioral heterogeneity in stock prices. *Journal of Economic dynamics and control*, 31(6):1938–1970, 2007.

[15] William A. Branch. The theory of rationally heterogeneous expectations: Evidence from survey data on inflation expectations. *The Economic Journal*, 114(497):592–621, 2004.

[16] William A Branch and George W Evans. Intrinsic heterogeneity in expectation formation. *Journal of Economic theory*, 127(1):264–295, 2006.

[17] William A Brock and Cars H Hommes. A rational route to randomness. *Econometrica*, 65(5):1059–1095, 1997.

[18] William A Brock and Cars H Hommes. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8):1235–1274, 1998.

[19] William A Brock, Cars H Hommes, and Florian OO Wagener. Evolutionary dynamics in markets with many trader types. *Journal of Mathematical Economics*, 41(1):7–42, 2005.

[20] William A Brock, Carsien Harm Hommes, and Florian Oskar Ottokar Wagener. More hedging instruments may destabilize markets. *Journal of Economic Dynamics and Control*, 33(11):1912–1928, 2009.

[21] Christoph Brunner, Colin F Camerer, and Jacob K Goeree. Stationary concepts for experimental 2 x 2 games: Comment. *American Economic Review*, 101(2):1029–40, 2011.

[22] Colin Camerer and Teck Ho. Experienced-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874, 1999.

[23] Colin F Camerer, Teck-Hua Ho, and Juin-Kuan Chong. Sophisticated experience-weighted attraction learning and strategic teaching in repeated games. *Journal of Economic theory*, 104(1):137–188, 2002.

[24] Anton Cheremukhin, Anna Popova, and Antonella Tutino. A theory of discrete choice with information costs. *Journal of Economic Behavior & Organization*, 113:34–50, 2015.

[25] Carl Chiarella. The dynamics of speculative behaviour. *Annals of operations research*, 37(1):101–123, 1992.

[26] Carl Chiarella, Roberto Dieci, and Laura Gardini. Asset price dynamics in a financial market with fundamentalists and chartists. *Discrete Dynamics in Nature and Society*, 6(2):69–99, 2001.

[27] Carl Chiarella, Roberto Dieci, and Laura Gardini. Asset price and wealth dynamics in a financial market with heterogeneous agents. *Journal of Economic Dynamics and Control*, 30(9):1755–1786, 2006.

[28] Carl Chiarella and Xue-Zhong He. Heterogeneous beliefs, risk and learning in a simple asset pricing model. *Computational Economics*, 19:95–132, 2002.

[29] Carl Chiarella and Xue-Zhong He. Dynamics of beliefs and learning under a_L-processes – the heterogeneous case. *Journal of Economic Dynamics and Control*, 27:503–531, 2003.

[30] Carl Chiarella, Xue-Zhong He, and Cars Hommes. A dynamic analysis of moving average rules. *Journal of Economic Dynamics and Control*, 30:1729–1753, 2006.

[31] Carl Chiarella, Xue-Zhong He, and Remco C.J. Zwinkels. Heterogeneous expectations in asset pricing:empirical evidence from the s&p500. *Journal of Economic Behavior and Organization*, 105:1–16, 2014.

[32] Thorsten Chmura, Sebastian J Goerg, and Reinhard Selten. Learning in experimental 2× 2 games. *Games and Economic Behavior*, 76(1):44–73, 2012.

[33] Annarita Colasante, Simone Alfarano, Eva Camacho, and Mauro Gallegati. Long-run expectations in a learning-to-forecast experiment. *Applied Economics Letters*, 25(10):681–687, 2018.

[34] Adriana Cornea-Madeira, Cars Hommes, and Domenico Massaro. Behavioral heterogeneity in us inflation dynamics. *Journal of Business and Economic Statistics*, 2017. forthcoming.

[35] Richard H. Day and Weihong Huang. Bulls, bears and market sheep. *Journal of Economic Behavior and Organization*, 14:299–219, 1990.

[36] Cees Diks, Cars Hommes, Valentyn Panchenko, and Roy Van Der Weide. E&f chaos: a user friendly software package for nonlinear economic dynamics. *Computational Economics*, 32(1-2):221–244, 2008.

[37] Cees Diks, Valentyn Panchenko, Oleg Sokolinskiy, and Dick van Dijk. Comparing the accuracy of multivariate density forecasts in selected regions of the copula support. *Journal of Economic Dynamics & Control*, 48:79–94, 2014.

[38] Pietro Dindo and Jan Tuinstra. A class of evolutionary models for participation games with negative feedback. *Computational Economics*, 37(3):267–300, 2011.

[39] Gerald P. Dwyer, Arlington W. Williams, Raymond C. Battalio, and Timothy I. Mason. Tests of rational expectations in a stark setting. *The Economic Journal*, 103(418):586–601, 1993.

[40] Ido Erev, Yoella Bereby-Meyer, and Alvin E Roth. The effect of adding a constant to all payoffs: experimental investigation, and implications for reinforcement learning models. *Journal of Economic Behavior & Organization*, 39(1):111–128, 1999.

[41] Ido Erev and Ernan Haruvy. Learning and the economics of small decisions. *The handbook of experimental economics*, 2, 2013.

[42] Ido Erev and Amnon Rapoport. Coordination,“magic,” and reinforcement learning in a market entry game. *Games and economic behavior*, 23(2):146–175, 1998.

[43] George W. Evans and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.

[44] Mordecai Ezekiel. The cobweb theorem. *The Quarterly Journal of Economics*, 52(2):255–280, 1938.

[45] Nick Feltovich. Reinforcement-based vs. belief-based learning models in experimental asymmetric-information games. *Econometrica*, 68(3):605–641, 2000.

[46] Urs Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171—-178, 2007.

[47] Reiner Franke and Frank Westerhoff. Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *Journal of Economic Dynamics & Control*, 36(8):1193–1211, 2012.

[48] Jeffrey A. Frankel and Kenneth A. Froot. Chartists, fundamentalists, and trading in the foreign exchange market. *The American Economic Review*, 80(2):181–185, 1990.

[49] Daniel Friedman, Steffen Huck, Ryan Oprea, and Simon Weidenholzer. From imitation to collusion: Long-run learning in a low-information environment. *Journal of Economic Theory*, 155:185–205, 2015.

[50] Andrea Gaunersdorfer. Endogenous fluctuations in a simple asset pricing model with heterogeneous agents. *Journal of Economic Dynamics and Control*, 24(5):799–831, 2000.

[51] Andrea Gaunersdorfer, Cars H Hommes, and Florian OO Wagener. Bifurcation routes to volatility clustering under evolutionary learning. *Journal of Economic Behavior & Organization*, 67(1):27–47, 2008.

[52] David Goldbaum and Bruce Mizrach. Estimating the intensity of choice in a dynamic mutual fund allocation decision. *Journal of Economic Dynamics and Control*, 32:3866–3876, 2008.

[53] David Goldbaum and Remco C.J. Zwinkels. An empirical examination of heterogeneity and switching in foreign exchange markets. *Journal of Economic Behavior and Organization*, 107(2):667–684, 2014.

[54] Ben Greiner. Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1):114–125, 2015.

[55] Peter R Hansen, Asger Lunde, and James M Nason. The model confidence set. *Econometrica*, 79(2):453–497, 2011.

[56] David W Harless and Colin F Camerer. The predictive utility of generalized expected utility theories. *Econometrica*, pages 1251–1289, 1994.

[57] Peter Heemeijer, Cars Hommes, Joep Sonnemans, and Jan Tuinstra. Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic Dynamics and Control*, 33:1052 – 1072, 2009.

[58] Peter Heemeijer, Cars Hommes, Joep Sonnemans, and Jan Tuinstra. 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, 2009.

[59] John D. Hey. Expectations formation: Rational or adaptive or...? *Journal of Economic Behavior and Organization*, 25(3):329–349, 1994.

[60] John D Hey. Expectations formation: Rational or adaptive or...? *Journal of Economic Behavior & Organization*, 25(3):329–349, 1994.

[61] Teck H Ho, Colin F Camerer, and Juin-Kuan Chong. Self-tuning experience weighted attraction learning in games. *Journal of Economic Theory*, 133(1):177–198, 2007.

[62] Teck H Ho, Xin Wang, and Colin F Camerer. Individual differences in ewa learning with partial payoff information. *The Economic Journal*, 118(525):37–59, 2008.

[63] Cars Hommes. Heterogeneous agent models: two simple examples. In *Nonlinear Dynamical Systems in Economics*, pages 131–164. Springer, 2005.

[64] Cars Hommes. The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics and Control*, 35:1–24, 2011.

[65] Cars Hommes. The heterogeneous expectations hypothesis: Some evidence from the lab. *Journal of Economic Dynamics & Control*, 35(1):1–24, 2011.

[66] Cars Hommes. *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press, 2013.

[67] Cars Hommes, Tatiana Kiseleva, Yuri Kuznetsov, and Miroslav Verbic. Is more memory in evolutionary selection (de) stabilizing? *Macroeconomic Dynamics*, 16(3):335–357, 2012.

[68] Cars Hommes, Joep Sonnemans, Jan Tuinstra, and Henk Van de Velden. Coordination of expectations in asset pricing experiments. *The Review of Financial Studies*, 18(3):955–980, 2004.

[69] Cars Hommes, Joep Sonnemans, Jan Tuinstra, and Henk van de Velden. Coordination of expectations in asset pricing experiments. *The Review of Financial Studies*, 18(3):955–980, 2005.

[70] Cars Hommes, Joep Sonnemans, Jan Tuinstra, and Henk van de Velden. Learning in cobweb experiments. *Macroeconomic Dynamics*, 11(1):8–33, 2007.

[71] Cars Hommes, Joep Sonnemans, Jan Tuinstra, and Henk Van de Velden. Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization*, 67(1):116–133, 2008.

[72] Reshmaan N Hussam, David Porter, and Vernon L Smith. Thar she blows: Can bubbles be rekindled with experienced subjects? *American Economic Review*, 98(3):924–37, 2008.

[73] Hugh Kelley and Daniel Friedman. Learning to forecast price. *Economic Inquiry*, 40(4):556–573, 2002.

[74] Michael Kirchler, Jürgen Huber, and Thomas Stöckl. Thar she bursts: Reducing confusion reduces bubbles. *American Economic Review*, 102(2):865–83, 2012.

[75] Bernhardt Lieberman. Human behavior in a strictly determined 3× 3 matrix game. *Behavioral Science*, 5(4):317–322, 1960.

[76] Lily Y Liu, Andrew J Patton, and Kevin Sheppard. Does anything beat 5-minute rv? a comparison of realized measures across multiple asset classes. *Journal of Econometrics*, 187(1):293–311, 2015.

[77] Michael C. Lovell. Test of the rational expectation hypothesis. *The American Economic Review*, 76(1):110–124, 1986.

[78] Robert Lucas. Adaptive behavior and economic theory. *The Journal of Business*, 59(4):401–426, 1986.

[79] Ramon Marimon and Shyam Sunder. Indeterminacy of equilibria in a hyperinflationary world: experimental evidence. *Econometrica: Journal of the Econometric Society*, pages 1073–1107, 1993.

[80] Jacob Marschak. Binary choice constraints and random utility indicators. Technical report, Yale Univ New Heaven CT Cowles Foundation For Research in Economics, 1959.

[81] Filip Matějka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.

[82] Daniel McFadden. Conditional logit analysis of qualitative choice behavior. In P Zarembka, editor, *Frontiers in Econometrics*, chapter 4, pages 105–142. Academic Press, 1973.

[83] Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995.

[84] Heinrich H Nax, Maxwell N Burton-Chellew, Stuart A West, and H Peyton Young. Learning in a black box. *Journal of Economic Behavior & Organization*, 127:1–15, 2016.

[85] Charles Noussair, Stephane Robin, and Bernard Ruffieux. Price bubbles in laboratory asset markets with constant fundamental values. *Experimental Economics*, 4(1):87–105, 2001.

[86] Jack Ochs. Games with unique, mixed strategy equilibria: An experimental study. *Games and Economic Behavior*, 10(1):202–217, 1995.

[87] Thomas R Palfrey, Charles A Holt, and Jacob Goeree. *Quantal Response Equilibrium-a Stochastic Theory of Games*. Princeton University Press, 2016.

[88] Valentyn Panchenko, Sergiy Gerasymchuk, and Oleg V Pavlov. Asset price dynamics with heterogeneous beliefs and local network interactions. *Journal of Economic Dynamics and Control*, 37(12):2623–2642, 2013.

[89] Dimitris N Politis and Joseph P Romano. The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313, 1994.

[90] Matthew Rabin. An approach to incorporating psychology into economics. *American Economic Review*, 103(3):617–22, 2013.

[91] Thomas J. Sargent. *Bounded Rationality in Macroeconomics*. Oxford University Press Inc., New York, 1993.

[92] Richard Schmalensee. An experimental study of expectation formation. *Econometrica*, 44(1):17–41, 1976.

[93] Richard Schmalensee. An experimental study of expectation formation. *Econometrica: journal of the Econometric Society*, pages 17–41, 1976.

[94] Noemi Schmitt, Jan Tuinstra, and Frank Westerhoff. Side effects of nonlinear profit taxes in an evolutionary market entry model: abrupt changes, coexisting attractors and hysteresis problems. *Journal of Economic Behavior & Organization*, 135:15–38, 2017.

[95] Noemi Schmitt, Jan Tuinstra, and Frank Westerhoff. Side effects of nonlinear profit taxes in an evolutionary market entry model: abrupt changes, coexisting attractors and hysteresis problems. *Journal of Economic Behavior & Organization*, 135:15–38, 2017.

[96] Reinhard Selten and Thorsten Chmura. Stationary concepts for experimental 2x2-games. *American Economic Review*, 98(3):938–66, 2008.

[97] Robert J Shiller. Do stock prices move too much to be justified by subsequent changes in dividends?, 1980.

[98] Herbert A. Simon. *Models of man*. John Wiley and Sons,Inc., New York, 1957.

[99] Christopher A. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980.

[100] Vernon L. Smith, Gerry L. Suchanek, and Alington W. Williams. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56(5):1119–1151, 1988.

[101] Jakub Steiner, Colin Stewart, and Filip Matějka. Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica*, 85(2):521–553, 2017.

[102] Thomas Stöckl, Jürgen Huber, and Michael Kirchler. Bubble measures in experimental asset markets. *Experimental Economics*, 13(3):284–298, 2010.

[103] Saska ter Ellen, Cars Hommes, and Remco Zwinkels. Comparing behavioural heterogeneity across asset classes. Technical report, 2017. SSRN paper.

[104] Saskia ter Ellen, Cars H. Hommes, and Remco C. J. Zwinkels. Comparing behavioural heterogeneity across asset classes. *https://ssrn.com/abstract=2944022 or http://dx.doi.org/10.2139/ssrn.2944022*, 2017.

[105] Saskia ter Ellen and Remco CJ Zwinkels. Oil price dynamics: A behavioral finance approach with heterogeneous agents. *Energy Economics*, 32(6):1427–1434, 2010.

[106] Louis L Thurstone. A law of comparative judgment. *Psychological review*, 34(4):273, 1927.

[107] Kenneth E Train. *Discrete choice methods with simulation*. Cambridge University Press, 2009.

[108] Jan Tuinstra, Michael Wegener, and Franc Westerhoff. Positive welfare effects of trade barriers in a dynamic partial equilibrium model. *Journal of Economic Dynamics and Control*, 48:246–264, 2014.

[109] Eric Van Damme. *Stability and perfection of Nash equilibria*. Springer-Verlag Berlin Heidelberg, 1991.

[110] Eric Van Damme and Jörgen W Weibull. Evolution in games with endogenous mistake probabilities. *Journal of Economic Theory*, 106(2):296–315, 2002.

[111] Frank Westerhoff. Heterogeneous traders and the tobin tax. *Journal of Evolutionary Economics*, 13(1):53–70, 2003.

[112] Frank H Westerhoff and Roberto Dieci. The effectiveness of keynes–tobin transaction taxes when heterogeneous agents can trade in different markets: a behavioral finance approach. *Journal of Economic Dynamics and Control*, 30(2):293–322, 2006.

[113] Frank H Westerhoff and Roberto Dieci. The effectiveness of keynes–tobin transaction taxes when heterogeneous agents can trade in different markets: a behavioral finance approach. *Journal of Economic Dynamics and Control*, 30(2):293–322, 2006.

[114] Nathaniel T Wilcox. Theories of learning in games and heterogeneity bias. *Econometrica*, 74(5):1271–1292, 2006.

[115] Michael Woodford. Stochastic choice: An optimizing neuroeconomic model. *American Economic Review*, 104(5):495–500, 2014.